

A COVARIANCE-ENHANCED APPROACH TO MULTI-TISSUE JOINT EQTL MAPPING WITH APPLICATION TO TRANSCRIPTOME-WIDE ASSOCIATION STUDIES

BY AARON J. MOLSTAD^{*}, WEI SUN^{†,‡,§} AND LI HSU^{†,‡}

University of Florida^{}, Fred Hutchinson Cancer Research Center[†], University of Washington[‡], and University of North Carolina-Chapel Hill[§]*

Transcriptome-wide association studies based on genetically predicted gene expression have the potential to identify novel regions associated with various complex traits. It has been shown that incorporating expression quantitative trait loci (eQTLs) corresponding to multiple tissue types can improve power for association studies involving complex etiology. In this article, we propose a new multivariate response linear regression model and method for predicting gene expression in multiple tissues simultaneously. Unlike existing methods for multi-tissue joint eQTL mapping, our approach incorporates tissue-tissue expression correlation, which allows us to more efficiently handle missing expression measurements and more accurately predict gene expression using a weighted summation of eQTL genotypes. We show through simulation studies that our approach performs better than the existing methods in many scenarios. We use our method to estimate eQTL weights for 29 tissues collected by GTEx, and show that our approach significantly improves expression prediction accuracy compared to competitors. Using our eQTL weights, we perform a multi-tissue-based S-MultiXcan [2] transcriptome-wide association study and show that our method leads to more discoveries in novel regions and more discoveries overall than the existing methods. Estimated eQTL weights and code for implementing the method are available for download online at github.com/ajmolstad/MTeQTLResults

1. Introduction. Genome-wide association studies (GWAS) have identified tens of thousands of reproducible trait associated single-nucleotide polymorphisms (SNPs) through agnostic SNP-by-SNP association analysis (see <https://ebi.ac.uk/gwas/>) [4]. Though most of these associated SNPs lie outside of any gene, they are enriched for expression quantitative trait loci (eQTL)[12, 22, 30], which are genetic loci that affect the expression of one or more genes [6, 27, 34, 39, 42]. Machine learning methods have been used to infer eQTL regulation of a gene using all nearby genetic variants [8, 21, 51]. Using genetically predicted gene expression from these models, researchers have performed transcriptome-wide association studies (TWAS) and reported many novel regions associated with various complex traits [11, 32, 20, 1], many of which have no GWAS association within 1Mb. There are several advantages to such analyses: leveraging gene expression enriches potential trait associated SNPs, aggregating signals through joint eQTL modeling enhances the overall association strength, and the number of tests is substantially reduced from testing millions of SNPs to about 20,000 genes. Because of the tissue-dependent nature of gene expression, these

Keywords and phrases: Transcriptome-wide association studies, multivariate regression, expression quantitative trait loci, multi-tissue integrative analysis, GTEx

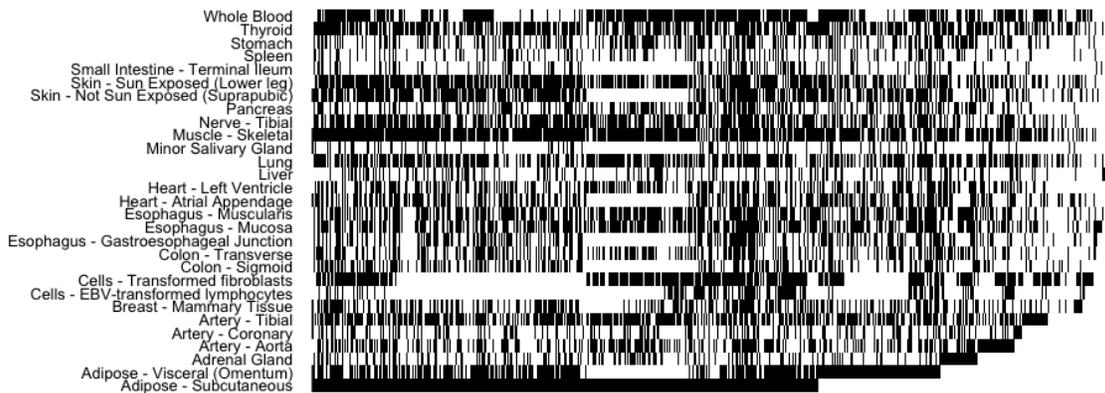


FIG 1. A heatmap of missingness in GTEx gene expression data for the 29 tissues we analyzed. Rows correspond to the 29 tissues and columns correspond to the 613 subjects with expression measured in at least one of the corresponding tissues. White spaces denote missing measured expression, whereas black denote observed measured expression.

analyses typically use gene expression from a single trait-relevant tissue. However, recent works have shown that eQTLs are often shared across multiple tissues and assessing the association of genetically predicted gene expression using multiple tissues improves power for genetic association with complex traits [1, 2].

Leveraging shared eQTLs across tissues improves power for eQTL discoveries and gene expression imputation accuracy, which can, in turn, further improve power for transcriptome-wide association analysis. Flutre et al. [7] and Li et al. [19] proposed multi-tissue eQTL mapping approaches that identified more eQTLs than tissue-by-tissue approaches. More recently, Hu et al. [14] proposed a penalized regression approach for joint modeling of eQTLs using a penalty which encourages shared eQTLs across tissues. Using the genotype and expression data for various tissues from the Genotype-Tissue Expression (GTEx) project [9, 10], they showed multi-tissue eQTL models improve imputation accuracy and gene association detection substantially compared to single-tissue approaches.

However, the method of Hu et al. [14] does not take into account tissue-tissue correlation of gene expression in joint eQTL modeling. In the recent statistical literature, it has been shown that in high-dimensional penalized multivariate response linear regression, accounting for error correlation (here, tissue-tissue correlation) often leads to improved variable selection and prediction accuracy [38, 48, 45, 18]. This phenomenon can be partly explained by a seemingly unrelated regression interpretation of high-dimensional sparse multivariate response linear regression [50, 41]. Further, owing to biological and cost constraints, some tissue types are harder to obtain than others. For example, of the 29 GTEx tissues we focus on in our analysis, there were 613 individuals with expression measured in at least one tissue of interest. Some individuals have as few as one tissue with measured expression, whereas others have expression measured in as many as 28 tissues. No individual has expression measured in all 29 tissues. The sample size per tissue varies from 85 (Minor Salivary Gland) to 491 (Muscle - Skeletal) (Figure 1). We show that leveraging tissue-tissue correlation can substantially improve gene expression prediction accuracy, especially for tissues with small

sample sizes. We do so by imputing the missing gene expression in a way which simultaneously estimates tissue-tissue correlation and joint eQTL weights. The mechanism through which this operates is straightforward. For example, if two tissues' expression is highly correlated, measuring expression in only one of these tissues allows one to reasonably estimate expression in the unmeasured tissue. If tissue-tissue correlation is ignored, substantial gains in gene expression prediction accuracy may be lost.

In this article, we propose a new method for multi-tissue joint eQTL mapping which can be used when individuals have missing gene expression measurements in many tissues. We develop an efficient penalized expectation conditional maximization (ECM) algorithm to solve the optimization problem. Our penalties allow us to identify both tissue specific and shared eQTLs while simultaneously modeling cross-tissue expression covariance. Compared to existing methods for multi-tissue joint eQTL mapping, our approach has several advantages:

1. our method explicitly models tissue-tissue correlation, providing new insights about cross-tissue expression dependence which cannot be explained by eQTLs;
2. by modeling cross-tissue correlation, our method more efficiently makes use of the available expression measurements for eQTL weight estimation;
3. in both simulations and our analysis of the GTEx data, our approach leads to improved gene expression prediction accuracy compared to (a) tissue-by-tissue approaches which estimate eQTL weights separately; (b) two-step approaches where imputation and prediction are performed in sequential steps; and (c) approaches which ignore cross-tissue correlation;
4. our computational approach can be used to more efficiently compute special cases of our method, e.g., the method of Hu et al. [14].

The implication of 3) on association analyses of genetically predicted gene expression is clear: with more accurate expression prediction models, one can perform more reliable tests, and thus, can expect more novel regions to be discovered. In particular, with the weights estimated from the GTEx data using our method, we performed a multi-tissue-based S-MultiXcan[2] TWAS analysis of four complex traits. We found that our weights led to more total discoveries and more novel discoveries beyond single variant analyses than existing methods in four traits we studied. Focusing on genes associated with the occurrence of a heart attack, we identified multiple regions which were not attributable to any GWAS associated SNP, but have been associated with traits related to heart function, e.g., coronary heart disease, triglyceride levels, and LDL cholesterol. These findings demonstrate the potential power gain using an integrative analysis with multiple tissues and also offer insight into potential factors associated with the occurrence of a heart attack. Discovered genes from our S-MultiXcan analysis, our estimated eQTL weights, and R code to reproduce the simulation study results are available for download at github.com/ajmolstad/MTeQTLResults.

2. Method.

2.1. *Penalized maximum likelihood estimator.* Throughout, we will assume a general model for cross-tissue expression given SNP genotypes. For a particular gene, let $\mathbf{y}_i \in \mathbb{R}^q$ denote the vector of centered and normalized measured expression in q tissues for the i th subject and let $\mathbf{x}_i \in \mathbb{R}^p$ denote the genotypes of p SNPs (also centered and normalized)

within a certain distance (e.g., 500kb) of the gene of interest. We assume that for the i th subject, log-transformed expression (e.g., see Section 2 of the Supplementary Material [26]) is a realization of the random vector

$$(2.1) \quad \mathbf{y}_i = \boldsymbol{\beta}'_* \mathbf{x}_i + \boldsymbol{\epsilon}_i, \quad \boldsymbol{\epsilon}_i \sim N_q(0, \boldsymbol{\Omega}_*^{-1}), \quad (i = 1, \dots, n)$$

where N_q denotes the q -dimensional multivariate normal distribution, $\boldsymbol{\beta}_* \in \mathbb{R}^{p \times q}$ is the unknown regression coefficient matrix (i.e., eQTL weights), and $\boldsymbol{\Omega}_* \in \mathbb{S}_+^q$ is the cross-tissue error precision (inverse covariance) matrix. We further assume that $\boldsymbol{\epsilon}_i$ is independent of $\boldsymbol{\epsilon}_j$ for all $i \neq j$. Throughout, we let \mathbb{S}_+^q denote the set of $q \times q$ symmetric and positive definite matrices. To simplify notation, we let $[n] \equiv \{1, 2, \dots, n\}$ throughout for all $n \in \mathbb{N}$.

We estimate $\boldsymbol{\beta}_*$ and $\boldsymbol{\Omega}_*$ jointly by maximizing the observed-data penalized log-likelihood with respect to $\boldsymbol{\beta}$ and $\boldsymbol{\Omega}$, the optimization variables corresponding to $\boldsymbol{\beta}_*$ and $\boldsymbol{\Omega}_*$, respectively. Let o_i and m_i be sets of indices corresponding the components of \mathbf{y}_i (tissues) which are observed and missing, respectively for $i = [n]$. Without loss of generality, we can then write $\mathbf{y}_i = (\mathbf{y}'_{i,o_i}, \mathbf{y}'_{i,m_i})' \in \mathbb{R}^q$ where A' denotes the transpose of matrix or vector A , and \mathbf{y}_{i,o_i} denotes the subvector of \mathbf{y}_i consisting of elements indexed by o_i . Let \mathbf{Y}_O denote the collection of the \mathbf{y}_{i,o_i} 's for $i = [n]$, i.e., \mathbf{Y}_O is the collection of all observed gene expression for the n subjects. Thus, the observed-data log-likelihood, $\log L(\boldsymbol{\beta}, \boldsymbol{\Omega} \mid \mathbf{Y}_O)$, (ignoring constants) for the n subjects is proportional to

$$(2.2) \quad \log L(\boldsymbol{\beta}, \boldsymbol{\Omega} \mid \mathbf{Y}_O) = -\frac{1}{n} \sum_{i=1}^n \{(\mathbf{y}'_{i,o_i} - \mathbf{x}'_i \boldsymbol{\beta}_{\cdot, o_i}) \boldsymbol{\Sigma}_{o_i}^{-1} (\mathbf{y}'_{i,o_i} - \mathbf{x}'_i \boldsymbol{\beta}_{\cdot, o_i})' + \log \det(\boldsymbol{\Sigma}_{o_i})\},$$

where $\boldsymbol{\Sigma}_{o_i} \in \mathbb{S}_+^{|o_i|}$ is the submatrix of $\boldsymbol{\Sigma} \equiv \boldsymbol{\Omega}^{-1}$ consisting of rows and columns indexed by o_i , $\boldsymbol{\beta}_{\cdot, o_i} \in \mathbb{R}^{p \times |o_i|}$ is the submatrix of $\boldsymbol{\beta}$ containing columns indexed by o_i , and $|o_i|$ denotes the cardinality of o_i . Throughout, tr and \det denote the trace and determinant operators, respectively. Note that we have assumed \mathbf{y}_i and \mathbf{x}_i have columnwise average zero so that we can write (2.1) without an intercept for ease of display.

Unfortunately, maximizing (2.2) directly is computationally difficult since missingness patterns differ across subjects. Instead, we use a variation of the expectation-maximization (EM) algorithm, which allows us to operate on a function of the complete-data log-likelihood. Analogously to \mathbf{Y}_O , let \mathbf{Y}_M denote the collection of all unmeasured (missing) gene expression for the n subjects. If we observed \mathbf{Y}_M , the complete data log-likelihood would be $\sum_{i=1}^n \log f_i(\mathbf{y}_{i,o_i}, \mathbf{y}_{i,m_i} \mid \boldsymbol{\beta}, \boldsymbol{\Omega})$ where f_i is the density corresponding to (2.1) for the i th subject.

To estimate $\boldsymbol{\beta}_*$ and $\boldsymbol{\Omega}_*$ while accounting for missingness, we propose to maximize a penalized version of $\log L(\boldsymbol{\beta}, \boldsymbol{\Omega} \mid \mathbf{Y}_O)$ with respect to $\boldsymbol{\beta}$ and $\boldsymbol{\Omega}$ jointly. Specifically, we propose to estimate $(\boldsymbol{\beta}_*, \boldsymbol{\Omega}_*)$ with

$$(2.3) \quad \arg \max_{\boldsymbol{\beta} \in \mathbb{R}^{p \times q}, \boldsymbol{\Omega} \in \mathbb{S}_+^q} \{ \log L(\boldsymbol{\beta}, \boldsymbol{\Omega} \mid \mathbf{Y}_O) - g_{\lambda_\beta, \alpha}(\boldsymbol{\beta}) - h_{\lambda_\Omega}(\boldsymbol{\Omega}) \},$$

where $\lambda_\beta > 0$, $\lambda_\Omega > 0$, and $\alpha \in [0, 1]$ are tuning parameters; and $g_{\lambda_\beta, \alpha} : \mathbb{R}^{p \times q} \rightarrow \mathbb{R}$ and $h_{\lambda_\Omega} : \mathbb{S}_+^q \rightarrow \mathbb{R}$ are convex penalty functions applied to $\boldsymbol{\beta}$ and $\boldsymbol{\Omega}$, respectively. These penalties are chosen based on the biology underlying multi-tissue joint eQTL mapping. In particular,

it is believed that large proportion of eQTLs are shared across multiple tissues in most genes [7, 19, 14], which would imply that nonzero entries of β_* are likely to occur in a subset of rows since each row of β_* corresponds to a particular SNP's regression coefficients (eQTL weights) for the q tissues. To exploit this assumption, we use a penalty which balances row-wise sparsity with element-wise sparsity:

$$(2.4) \quad g_{\lambda_\beta, \alpha}(\beta) = \lambda_\beta \sum_{j=1}^p \left\{ \alpha \left(\sum_{k=1}^q \sqrt{\frac{n_{\max}}{n_k}} |\beta_{j,k}| \right) + (1 - \alpha) \|\beta_{j,\cdot}\|_2 \right\},$$

where $\|\cdot\|_2$ denotes the Euclidean norm of a vector, $\beta_{j,\cdot} \in \mathbb{R}^q$ is the j th row of β for $j \in [p]$, $n_k = |\{j : j \in o_k\}|$ for each $k \in [q]$, and $n_{\max} = \max_{l \in [q]} n_l$. If $\alpha = 0$, estimates of β_* would only have rows which are entirely zero or nonzero. Conversely, when $\alpha = 1$, this penalty does not encourage eQTL sharing explicitly. A variation of the penalty in (2.4) was also used by Peng et al. [33] in the context of the multivariate response linear regression model of gene expression on copy number alterations.

To estimate the cross-tissue error precision matrix Ω_* , we use an L_1 -norm penalty on the entries of the corresponding optimization variable

$$h_{\lambda_\Omega}(\Omega) = \lambda_\Omega \sum_{j=1}^q \sum_{k=1}^q |\Omega_{j,k}|.$$

For sufficiently large values of the tuning parameter λ_Ω , this penalty leads to estimates of the precision matrix which will have all off-diagonal entries equal to zero [49, 37]. Hence, this penalty implicitly assumes that some entries of Ω_* equal zero. This assumption is also biologically reasonable in the context of multi-tissue joint eQTL mapping: it is well known that under (2.1), a zero in the (j, k) th entry of Ω_* implies that expression in the j th and k th tissues are independent given expression in all other tissues and all p SNP genotypes. Note that we penalize the diagonals of Ω . We do so for two reasons: first, when $p > n$ (which is the case for nearly every gene we fit the model to), without penalizing the diagonals, a perfect fit can occur in the M-step. This point was mentioned, for example, in the R package documentation for MRCE, which implements the method described in Rothman et al. [37]. Second, when fitting predictive models, additional shrinkage imposed on the diagonals can lead to improved prediction accuracy. This was observed in the recent literature on precision matrix estimators used in predictive models, e.g., Molstad and Rothman [24].

2.2. Penalized expectation conditional maximization algorithm. To obtain the penalized maximum likelihood estimator, i.e., to solve (2.3) with the penalties g and h , we propose a penalized expectation conditional maximization (ECM) algorithm [23]. The (k) th iterate of the standard EM algorithm is computed in two steps: the E-step and the M-step. The E-step requires computing the expectation of complete data negative log-likelihood conditional on \mathbf{Y}_O and the previous iterates $\beta^{(k-1)}$ and $\Omega^{(k-1)}$

$$\mathcal{Q}(\beta, \Omega \mid \beta^{(k-1)}, \Omega^{(k-1)}) = -\mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n \log f_i(\mathbf{y}_{i,o_i}, \mathbf{y}_{i,m_i} \mid \beta, \Omega) \mid \mathbf{Y}_O, \beta^{(k-1)}, \Omega^{(k-1)} \right]$$

and the subsequent M-step updates

$$(2.5) \quad (\boldsymbol{\beta}^{(k)}, \boldsymbol{\Omega}^{(k)}) \in \arg \min_{\substack{\boldsymbol{\beta} \in \mathbb{R}^{p \times q}, \boldsymbol{\Omega} \in \mathbb{S}_+^q}} \left\{ \mathcal{Q}(\boldsymbol{\beta}, \boldsymbol{\Omega} \mid \boldsymbol{\beta}^{(k-1)}, \boldsymbol{\Omega}^{(k-1)}) + g_{\lambda_{\beta}, \alpha}(\boldsymbol{\beta}) + h_{\lambda_{\Omega}}(\boldsymbol{\Omega}) \right\}.$$

However, existing methods for solving (2.5) use a blockwise coordinate descent algorithm iterating between updating $\boldsymbol{\Omega}$ and $\boldsymbol{\beta}$. Instead, we propose to update each variable once (with the other held fixed) and then return to the E-step. Hence, our algorithm is a *generalized* expectation conditional maximization algorithm in the sense that we do not solve the M-step exactly at each iteration, but are guaranteed that the objective function is non-increasing. An outline of the complete algorithm is given in Algorithm 1.

Algorithm 1. Penalized ECM algorithm for computing (2.3)

1. Initialize $\boldsymbol{\beta}^{(1)} \in \mathbb{R}^{p \times q}$ and $\boldsymbol{\Omega}^{(1)} \in \mathbb{S}_+^q$. Set $k = 1$.
 2. Compute $\mathcal{Q}(\boldsymbol{\beta}, \boldsymbol{\Omega} \mid \boldsymbol{\beta}^{(k)}, \boldsymbol{\Omega}^{(k)})$
 3. Update $\boldsymbol{\Omega}^{(k+1)} \leftarrow \arg \min_{\boldsymbol{\Omega} \in \mathbb{S}_+^q} \{ \mathcal{Q}(\boldsymbol{\beta}^{(k)}, \boldsymbol{\Omega} \mid \boldsymbol{\beta}^{(k)}, \boldsymbol{\Omega}^{(k)}) + h_{\lambda_{\Omega}}(\boldsymbol{\Omega}) \}$.
 4. Update $\boldsymbol{\beta}^{(k+1)} \leftarrow \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^{p \times q}} \{ \mathcal{Q}(\boldsymbol{\beta}, \boldsymbol{\Omega}^{(k+1)} \mid \boldsymbol{\beta}^{(k)}, \boldsymbol{\Omega}^{(k)}) + g_{\lambda_{\beta}, \alpha}(\boldsymbol{\beta}) \}$.
 5. If the objective value from (2.3) has not converged, update $k \leftarrow k + 1$ and return to Step 2.
-

Because the updates for $\boldsymbol{\Omega}$ and $\boldsymbol{\beta}$ (with the other held fixed) are both convex optimization problems, the M-step is an instance of a *biconvex* optimization problem. In the following subsection, we describe how to solve Steps 2-4 of Algorithm 1.

2.3. *Algorithm details.* To compute the Q function in Step 2 (E-Step) of the penalized ECM algorithm, we use the conditional multivariate normal distribution for \boldsymbol{y}_{i, m_i} given \boldsymbol{y}_{i, o_i} and the current iterates of $\boldsymbol{\beta}$ and $\boldsymbol{\Omega}$. Specifically,

$$(2.6) \quad \boldsymbol{y}_{i, m_i} \mid \boldsymbol{y}_{i, o_i}, \boldsymbol{\beta}^{(k)}, \boldsymbol{\Omega}^{(k)} \sim N_{|m_i|} \left(\boldsymbol{\mu}_i^{(k)}, \boldsymbol{V}_i^{(k)} \right), \quad (i = 1, \dots, n),$$

where the mean and variance from (2.6) are

$$\begin{aligned} \boldsymbol{\mu}_i^{(k)} &= \boldsymbol{\beta}_{\cdot, m_i}^{(k)'} \boldsymbol{x}_i + \boldsymbol{\Sigma}_{m_i, o_i}^{(k)} [\boldsymbol{\Sigma}_{o_i}^{(k)}]^{-1} (\boldsymbol{y}_{i, o_i} - \boldsymbol{\beta}_{\cdot, o_i}^{(k)'} \boldsymbol{x}_i); \\ \boldsymbol{V}_i^{(k)} &= \boldsymbol{\Sigma}_{m_i}^{(k)} - \boldsymbol{\Sigma}_{m_i, o_i}^{(k)} [\boldsymbol{\Sigma}_{o_i}^{(k)}]^{-1} \boldsymbol{\Sigma}_{o_i, m_i}^{(k)} \end{aligned}$$

with $\boldsymbol{\Sigma}_{m_i, o_i}^{(k)}$ denoting the submatrix of $\boldsymbol{\Sigma}^{(k)}$ whose rows correspond to the index set m_i and whose columns correspond to index set o_i , i.e., $\boldsymbol{\Sigma}_{o_i}^{(k)} \equiv \boldsymbol{\Sigma}_{o_i, o_i}^{(k)}$ and $\boldsymbol{\Sigma}_{m_i}^{(k)} \equiv \boldsymbol{\Sigma}_{m_i, m_i}^{(k)}$. With $\boldsymbol{\mu}_i^{(k)}$ and $\boldsymbol{V}_i^{(k)}$ computed and stored for each $i \in [n]$, we can then express Step 3 of the penalized ECM algorithm as a familiar convex optimization problem.

REMARK 1. *Step 3 of Algorithm 1 can be expressed*

$$(2.7) \quad \boldsymbol{\Omega}^{(k+1)} = \arg \min_{\boldsymbol{\Omega} \in \mathbb{S}_+^q} \left[\text{tr} \left\{ \mathbf{S}(\boldsymbol{\beta}^{(k)}, \boldsymbol{\Sigma}^{(k)}) \boldsymbol{\Omega} \right\} - \log \det(\boldsymbol{\Omega}) + \lambda_{\Omega} \sum_{j=1}^q \sum_{k=1}^q |\boldsymbol{\Omega}_{j, k}| \right]$$

where $\mathbf{S}(\boldsymbol{\beta}^{(k)}, \boldsymbol{\Sigma}^{(k)}) = n^{-1} \sum_{i=1}^n \boldsymbol{\Gamma}_i^{(k)}$ with submatrices of $\boldsymbol{\Gamma}_i^{(k)}$ for each $i \in [n]$ equal to

$$\begin{aligned} [\boldsymbol{\Gamma}_i^{(k)}]_{o_i} &= (\mathbf{y}'_{i,o_i} - \mathbf{x}'_i \boldsymbol{\beta}_{\cdot,o_i}^{(k)})' (\mathbf{y}'_{i,o_i} - \mathbf{x}'_i \boldsymbol{\beta}_{\cdot,o_i}^{(k)}), \\ [\boldsymbol{\Gamma}_i^{(k)}]_{m_i} &= (\boldsymbol{\mu}_i^{(k)'} - \mathbf{x}'_i \boldsymbol{\beta}_{\cdot,m_i}^{(k)})' (\boldsymbol{\mu}_i^{(k)'} - \mathbf{x}'_i \boldsymbol{\beta}_{\cdot,m_i}^{(k)}) + \mathbf{V}_i^{(k)}, \\ [\boldsymbol{\Gamma}_i^{(k)}]_{o_i,m_i} &= (\mathbf{y}'_{i,o_i} - \mathbf{x}'_i \boldsymbol{\beta}_{\cdot,o_i}^{(k)})' (\boldsymbol{\mu}_i^{(k)'} - \mathbf{x}'_i \boldsymbol{\beta}_{\cdot,m_i}^{(k)}). \end{aligned}$$

Conveniently, (2.7) is exactly the optimization problem for computing the L_1 -penalized normal log-likelihood precision matrix estimator with input sample covariance matrix $\mathbf{S}(\boldsymbol{\beta}^{(k)}, \boldsymbol{\Sigma}^{(k)})$. Many efficient algorithms and software packages exist for computing (2.7). In our implementation, we used the software QUIC to solve (2.7) [13].

Step 4 of Algorithm 1, the update for $\boldsymbol{\beta}$ with $\boldsymbol{\Omega}^{(k+1)}$ fixed, can be expressed as a minimizer of a penalized weighted residual sum of squares criterion.

REMARK 2. *Step 4 of Algorithm 1 can be expressed*

$$(2.8) \quad \boldsymbol{\beta}^{(k+1)} = \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^{p \times q}} \left\{ \frac{1}{n} \sum_{i=1}^n (\tilde{\mathbf{y}}_i^{(k)} - \boldsymbol{\beta}' \mathbf{x}_i)' \boldsymbol{\Omega}^{(k+1)} (\tilde{\mathbf{y}}_i^{(k)} - \boldsymbol{\beta}' \mathbf{x}_i) + g_{\lambda_\beta, \alpha}(\boldsymbol{\beta}) \right\}$$

where $\tilde{\mathbf{y}}_i^{(k)} \in \mathbb{R}^q$ with $\tilde{\mathbf{y}}_{i,o_i}^{(k)} = \mathbf{y}_{i,o_i}$ and $\tilde{\mathbf{y}}_{i,m_i}^{(k)} = \boldsymbol{\mu}_i^{(k)}$ for $i = [n]$.

To solve (2.8), we use an accelerated proximal gradient descent algorithm [31]. We briefly motivate this iterative procedure from a majorize-minimize perspective [17]. Let $\ell_{\lambda_\beta} : \mathbb{R}^{p \times q} \rightarrow \mathbb{R}$ denote the objective function from (2.8) (ignoring dependence on α for ease of display) so that ℓ_0 denotes the unpenalized weighted residual sum of squares (i.e., the objective from (2.8) with $\lambda_\beta = 0$). Let $\|\mathbf{A}\|_F^2 = \text{tr}(\mathbf{A}'\mathbf{A})$ denote the squared Frobenius norm of a matrix \mathbf{A} . Let $\boldsymbol{\beta}^{(r)}$ be the r th iterate of $\boldsymbol{\beta}$ for the sub-algorithm to solve Step 4. Given $\boldsymbol{\beta}^{(r)}$, because ℓ_0 is convex and has Lipschitz continuous gradient,

$$\ell_0(\boldsymbol{\beta}) \leq \ell_0(\boldsymbol{\beta}^{(r)}) + \text{tr}\{\nabla \ell_0(\boldsymbol{\beta}^{(r)})'(\boldsymbol{\beta} - \boldsymbol{\beta}^{(r)})\} + \frac{1}{2\gamma} \|\boldsymbol{\beta} - \boldsymbol{\beta}^{(r)}\|_F^2$$

for all $\boldsymbol{\beta}$ with step size $\gamma > 0$ sufficiently small. It follows that

$$(2.9) \quad \ell_{\lambda_\beta}(\boldsymbol{\beta}) \leq \ell_0(\boldsymbol{\beta}^{(r)}) + \text{tr}\{\nabla \ell_0(\boldsymbol{\beta}^{(r)})'(\boldsymbol{\beta} - \boldsymbol{\beta}^{(r)})\} + \frac{1}{2\gamma} \|\boldsymbol{\beta} - \boldsymbol{\beta}^{(r)}\|_F^2 + g_{\lambda_\beta, \alpha}(\boldsymbol{\beta})$$

for all $\boldsymbol{\beta}$ with equality when $\boldsymbol{\beta} = \boldsymbol{\beta}^{(r)}$. Thus, if we minimize the right hand side of (2.9) with respect to $\boldsymbol{\beta}$, we are guaranteed that $\ell_{\lambda_\beta}(\boldsymbol{\beta}^{(r+1)}) \leq \ell_{\lambda_\beta}(\boldsymbol{\beta}^{(r)})$. This suggests an iterative procedure to solve (2.8): in the first step, we construct the *majorizing function* [17] from the right hand side of (2.9) at the current iterate; in the second step, we minimize this majorizing function to obtain our new iterate; and then we repeat these two steps until the objective function from (2.8) converges.

This approach is computationally efficient because (2.10) can be solved in (essentially) closed form by computing the so-called *proximal operator* of g [36]. After some algebra, we see to minimize the right hand side of (2.9), we solve

$$(2.10) \quad \boldsymbol{\beta}^{(r+1)} = \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^{p \times q}} \left\{ \frac{1}{2} \|\boldsymbol{\beta} - \boldsymbol{\beta}^{(r)}\|_F^2 + \gamma \nabla \ell_0(\boldsymbol{\beta}^{(r)})' \boldsymbol{\beta} + \gamma g_{\lambda_\beta, \alpha}(\boldsymbol{\beta}) \right\},$$

which is, by definition, the proximal operator of γg evaluated at $\boldsymbol{\beta}^{(r)} - \gamma \nabla \ell_0(\boldsymbol{\beta}^{(r)})$. Define $\text{soft}(x, \tau) = \max(|x| - \tau, 0)\text{sign}(x)$ elementwise. Then, $\boldsymbol{\beta}^{(r+1)}$ can be obtained using the updating equations

1. $\boldsymbol{\Delta} \leftarrow \boldsymbol{\beta}^{(r)} - \gamma \nabla \ell_0(\boldsymbol{\beta}^{(r)})$,
2. $\bar{\boldsymbol{\Delta}}_{j,k} \leftarrow \text{soft}\left(\boldsymbol{\Delta}_{j,k}, \gamma \lambda_\beta \alpha \sqrt{\frac{n_{\max}}{n_k}}\right)$ for all $(j, k) \in [p] \times [q]$,
3. $\boldsymbol{\beta}_{j,\cdot}^{(r+1)} \leftarrow \begin{cases} \max\left(1 - \frac{\gamma \lambda_\beta (1-\alpha)}{\|\boldsymbol{\Delta}_{j,\cdot}\|_2}, 0\right) \bar{\boldsymbol{\Delta}}_{j,\cdot} & : \bar{\boldsymbol{\Delta}}_{j,\cdot} \neq 0 \\ 0 & : \bar{\boldsymbol{\Delta}}_{j,\cdot} = 0 \end{cases}$ for all $j \in [p]$.

For a derivation, see Simon et al. [40]. For the complete algorithm to solve (2.8), steps 1–3 above are repeated until the objective function converges.

In Section 4 of the Supplementary Material [26], we detail exactly the steps of an accelerated version this algorithm which we use in our implementation, and discuss selecting the step size γ . In Section 1 of the Supplementary Material, we provide additional details about our implementation of the penalized ECM algorithm. To summarize, we use *warm-starts* and employ a simple heuristic to determine where to stop computing (2.3) along its solution path. In this same section, we propose a procedure for constructing sets of reasonable candidate tuning parameters based on the Karush-Kuhn-Tucker (KKT) conditions for (2.7) and (2.8). Computing time statistics and a discussion of approaches for parallelization can also be found in the Supplementary Material.

2.4. Related methods. A special case of our method was proposed by Hu et al. [14]. In particular, the objective function they propose to minimize to estimate eQTL weights is equivalent to (2.3) with the constraint that $\boldsymbol{\Omega} = I_q$ in the case that $n_j = n_k$ for all $j \neq k$ (although they use scaling factor n_{\max}/n in their g). That is, their method implicitly assume that gene expression is uncorrelated with equal variance after conditioning on SNP genotypes. However, Hu et al. [14] do not solve the optimization problem they posed directly. Instead, they devised an efficient coordinate descent scheme which approximated their estimator. While their approximation performed well in terms of predicting expression, comparing our method to theirs directly is difficult: it is not clear when to terminate their iterative procedure because their algorithm does not minimize an objective function whose value can be computed. Hence, we compare our approach, which assumes only $\boldsymbol{\Omega} \in \mathbb{S}_+^q$, to what we refer to as the exact version of the Hu et al. [14] (i.e., (2.3) with the constraint $\boldsymbol{\Omega} = I_q$) in our simulation studies, the GTEx data analysis, and the S-MultiXcan TWAS. To compute the estimator of Hu et al. [14], we use a proximal gradient descent algorithm similar to that used for the optimization problem in (2.8). Additional details about this algorithm are also provided in the Supplementary Material.

3. Simulation studies.

3.1. Data generating models and competing methods. We performed extensive numerical experiments to study how the number of shared eQTLs, the population R^2 (equivalently, narrow-sense heritability, [44]), and tissue-tissue correlation structure affect the performance of various methods for estimating multi-tissue eQTL weights.

To closely mimic the settings of the joint eQTL mapping in the GTEx data, we obtained whole genome sequencing SNP genotype data for all SNPs within 500kb of the

gene SPATC1L [28] for 620 subjects from the GTEx dataset. After pruning highly correlated SNPs (see Data Preparation section), we are left with $p = 1178$ SNP genotypes. For each replication, we then generated $n = 620$ subjects' expression in $q = 29$ tissues: letting $\mathbf{x}_i \in \mathbb{R}^p$ be the SNP genotypes for the i th subject, we generated $\mathbf{y}_i \in \mathbb{R}^q$, as a realization of the random vector $\beta_*' \mathbf{x}_i + \epsilon_i$ for $i = [n]$, where $\beta_* \in \mathbb{R}^{p \times q}$ are the eQTL weights and $\epsilon_i \sim N_q(0, \Omega_*^{-1})$ are independent and identically distributed errors. For five hundred independent replications in each setting, we randomly split the $n = 620$ subjects into a training set of size $n_{\text{train}} = 400$, a validation set of size 110, and a testing set of size 110.

Independently in each replication, we generated β_* as follows: first, we generated $B \in \mathbb{R}^{p \times q}$ to have entries which were independent $N(0, 1)$. Then, we generated $S \in \mathbb{R}^{p \times q}$ to be a matrix whose rows are either all zero or all one: we randomly select s rows to be nonzero, where $s \in [20]$. Given S , we then generated $U \in \mathbb{R}^{p \times q}$ so that each of the q columns has 20- s randomly selected entries equal to one only from entries which are zero in S and all others equal to zero. With these, we set $\beta_* = B \circ S + B \circ U$ where \circ denotes the elementwise product. By constructing β_* in this way, each tissue has twenty total eQTLs, s of which are shared across all tissue types. We consider $s = \{5, 10, 15, 18, 20\}$ in the simulations we present in this section. Note that since many SNPs are highly correlated (linkage disequilibrium), marginally, there are far more SNPs associated with gene expression.

We construct Ω_*^{-1} to have a block-diagonal structure and to control the R^2 . Specifically, we set $\Omega_*^{-1} = \mathbf{D} \tilde{\Sigma} \mathbf{D}$ where $\mathbf{D} \in \mathbb{R}^{q \times q}$ is a diagonal matrix with positive entries and $\tilde{\Sigma} \in \mathbb{S}_+^q$ is a correlation matrix. In our analysis of the GTEx data, we found that on average, the estimated correlation matrix had an approximately 20×20 correlated block, of which a 10×10 sub-block was more highly correlated. Thus, we set

$$\tilde{\Sigma}_{j,k} = \begin{cases} \rho & (j, k) \in [20] \times [20]; (j, k) \notin [10] \times [10]; j \neq k \\ \rho + 0.2 & (j, k) \in [10] \times [10]; j \neq k \\ 1 & j = k \\ 0 & \text{otherwise} \end{cases}$$

In the results presented here, we considered $\rho \in \{0, 0.1, 0.3, 0.5, 0.7\}$. Given $\tilde{\Sigma}$ and β_* , we then generated entries of \mathbf{D} to determine the R^2 for all $q = 29$ tissues: we considered $R^2 \in \{.01, .05, .10, .20, .40\}$.

Finally, we also randomly assigned missingness to both the training set and validation set responses with missing probability equal to 0.55, i.e., the missing rate in the GTEx gene expression data we analyzed. For each method, we fit the model to the training data, selected tuning parameters using the validation data, and measured the prediction and variable selection accuracy on the testing data.

We considered three different methods: two of which can be considered ‘‘complete-case’’ estimators. We define the missingness matrix $\mathbf{M} \in \mathbb{R}^{n_{\text{train}} \times q}$

$$\mathbf{M}_{i,k} = \begin{cases} n_k^{-1/2} & : \mathbf{y}_{i,k} \text{ was observed} \\ 0 & : \mathbf{y}_{i,k} \text{ was missing} \end{cases} \quad (i, k) \in [n_{\text{train}}] \times [q],$$

where n_k is the number of subjects with expression observed for the k th tissue in the training data. Similarly, let $\mathbf{Y} \in \mathbb{R}^{n_{\text{train}} \times q}$ be the fully-observed gene expression training data matrix, and $\mathbf{X} \in \mathbb{R}^{n_{\text{train}} \times p}$ the SNP genotypes for the training data. We consider three alternative methods.

– **EN**. The tissue-by-tissue elastic net defined as

$$\arg \min_{\boldsymbol{\beta} \in \mathbb{R}^{p \times q}} \left\{ \frac{1}{2} \|(\mathbf{Y} - \mathbf{X}\boldsymbol{\beta}) \circ \mathbf{M}\|_F^2 + \sum_{k=1}^q \lambda_k \left(\alpha_k \sum_{j=1}^p |\boldsymbol{\beta}_{j,k}| + \frac{(1 - \alpha_k)}{2} \sum_{j=1}^p \boldsymbol{\beta}_{j,k}^2 \right) \right\},$$

where for each $k \in [q]$, $(\lambda_k, \alpha_k) \in \mathbb{R}_+ \times [0, 1]$ is chosen to maximize the validation set R^2 for the k th tissue. This is the default method for estimating eQTL weights in Gamazon et al. [8].

– **MT**. The exact version of the estimator proposed by Hu et al. [14]

$$(3.1) \quad \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^{p \times q}} \left\{ \frac{1}{2} \|(\mathbf{Y} - \mathbf{X}\boldsymbol{\beta}) \circ \mathbf{M}\|_F^2 + g_{\lambda_\beta, \alpha}(\boldsymbol{\beta}) \right\},$$

where $\lambda_\beta > 0$ and $\alpha \in [0, 1]$ are chosen to maximize validation set R^2 averaged over all 29 tissues.

– **Cov-MT**. The model-based approach we proposed in (2.3) with tuning parameters $\lambda_\beta > 0, \lambda_\Omega > 0$, and $\alpha \in [0, 1]$ chosen to maximize the validation set R^2 averaged over all 29 tissues.

We also obtained oracle estimators **Or-EN** and **Or-MT**. These are equivalent to **EN** and **MT** without any missingness. Here **Or** stands for ‘‘Oracle’’, i.e., an estimator which has information not available in practice. These estimators replace \mathbf{M} with an $\mathbb{R}^{n_{\text{train}} \times q}$ matrix with each entry equal to $n_{\text{train}}^{-1/2}$. These estimators are meant to compare to the idealized setting where there is no missingness in the response to distinguish between the effects of missingness and the effects of ignoring tissue-tissue correlation.

Finally, we also considered **KNN(20)-MT**, which first imputes the missing responses using a weighted mean based on the 20-nearest neighbors (subjects) and then used the same criterion as **Or-MT** to estimate $\boldsymbol{\beta}_*$. We also tried imputation via k -nearest neighbors with $k \in \{2, 5, 10, 50\}$, but results did not differ substantially across these choices of k , so additional results are omitted.

We measured performance using three metrics. The metric we used to measure prediction accuracy was test-set R^2 averaged over the q responses.

Note that testing set R^2 can be less than zero, which would occur when the training sample mean predicts the testing data better than the estimate of $\boldsymbol{\beta}_*$. We also measured the linkage disequilibrium (LD)-adjusted true positive rate

$$\frac{\sum_{k=1}^q \left\{ \sum_{j=1}^p \mathbf{1} \left(\boldsymbol{\beta}_{*j,k} \neq 0 \cap \hat{\boldsymbol{\beta}}_{l,k} \neq 0 \text{ for any } l \in \{l : |\boldsymbol{\Sigma}_{l,j}^X| > 0.60\} \right) \right\}}{\sum_{k=1}^q \sum_{j=1}^p \mathbf{1}(\boldsymbol{\beta}_{*j,k} \neq 0)}$$

where $\boldsymbol{\Sigma}_{l,j}^X$ denotes the correlation between the l th and j th SNPs in the complete dataset. The LD-adjustment accounts for the fact that SNP genotypes are often highly correlated. Under our data generating model, we consider an eQTL discovery ‘‘true’’ if the selected SNP genotype is moderately correlated with a true eQTL (i.e., predictor whose corresponding regression coefficient is nonzero).

In addition to true positive rate, we also measured model size, i.e., the proportion of nonzero entries of the estimated regression coefficient matrix. Note that under our data generating models, the true model size is approximately .017, i.e., $(20/p)$, but because SNPs are so highly correlated, much larger estimated models could be expected.

3.2. Simulation study results. We present complete simulation study results in Figure 2. In the top row, (a), we present results with tissue-tissue correlation $\rho \in \{0.0, 0.1, 0.3, 0.5, 0.7\}$ varying, the population R^2 fixed at 0.10, and the proportion of shared eQTLs fixed at 0.75 (i.e., $s = 15$). In this setting, we observe that our method, **Cov-MT** performs better than all realistic competitors: only **Or-MT**, the version of MT which does not have any responses missing, outperforms our method when ρ is less than 0.5. As one would expect, when expression is nearly uncorrelated ($\rho = 0$), our method **Cov-MT** performs similarly to the exact version of the method of Hu et al. [14], which implicitly assumes no tissue-tissue correlation. Remarkably, when ρ is greater than or equal to 0.50, **Cov-MT** outperforms even the “Oracle” methods which have no missingness. In fact, the prediction accuracy of **Cov-MT** increases as ρ increases, whereas all other methods, which do not explicitly model tissue-tissue correlation, have prediction accuracy remaining constant or slightly decreasing as ρ increases. This demonstrates the benefit of accounting for tissue-tissue correlation in multi-tissue joint eQTL mapping when expression across tissue types can be reasonably assumed to be conditionally dependent. It is also notable that **EN** performs very poorly: this is partly attributable to the fact that this approach does not leverage potential eQTLs shared across tissues, and thus, has relatively poor variable selection of eQTLs, which is apparent from the results displayed in the middle figure of row (a).

As ρ increases, the LD-adjusted true positive rate (TPR) of our method tends towards one, whereas for many of the competitors, the true positive rate decreases as ρ increases. This may partly be due to the fact that these methods tend to estimate fewer eQTLs as ρ increases, which is demonstrated in the rightmost figure of row (a). Notably, all methods tend to yield much larger models than the true model. Finally, we also observe that our method significantly outperforms the 20-nearest neighbor two-step imputation approach (**KNN(20)-MT**), which first imputes missing values via k -nearest neighbor and then fits the **Or-MT** model to the imputed dataset. This demonstrates the importance of jointly estimating the model parameters and performing expression imputation.

In the middle row, (b), of Figure 2, we present results with $R^2 \in \{0.01, 0.05, 0.1, 0.2, 0.4\}$, the proportion of shared eQTLs fixed at 0.75, and $\rho = 0.50$ fixed. We observe that our method performs as well or better than competitors across all settings in terms of average test set R^2 , except for when $R^2 = 0.40$, in which case **Or-MT** performs best. Interestingly, our method also performs best in terms of LD-adjusted TPR for small R^2 , but as R^2 approaches 0.40, many methods tend to perform similarly, though **KNN(20)-MT** nearly doubled the model size relative to other methods. It is also notable that even **EN**, which performed very poorly in the settings displayed in row (a), actually performs better than **KNN(20)-MT** in terms of prediction accuracy when the population $R^2 = 0.40$.

Finally, in the bottom row, (c), of Figure 2, results are displayed letting the proportion of shared eQTLs vary with $R^2 = 0.10$ fixed and $\rho = 0.50$ fixed. All methods which can exploit shared eQTLs (i.e., all methods other than **EN**, tissue-by-tissue elastic net) improve in prediction accuracy and LD-adjusted TPR as the proportion of shared eQTLs approaches

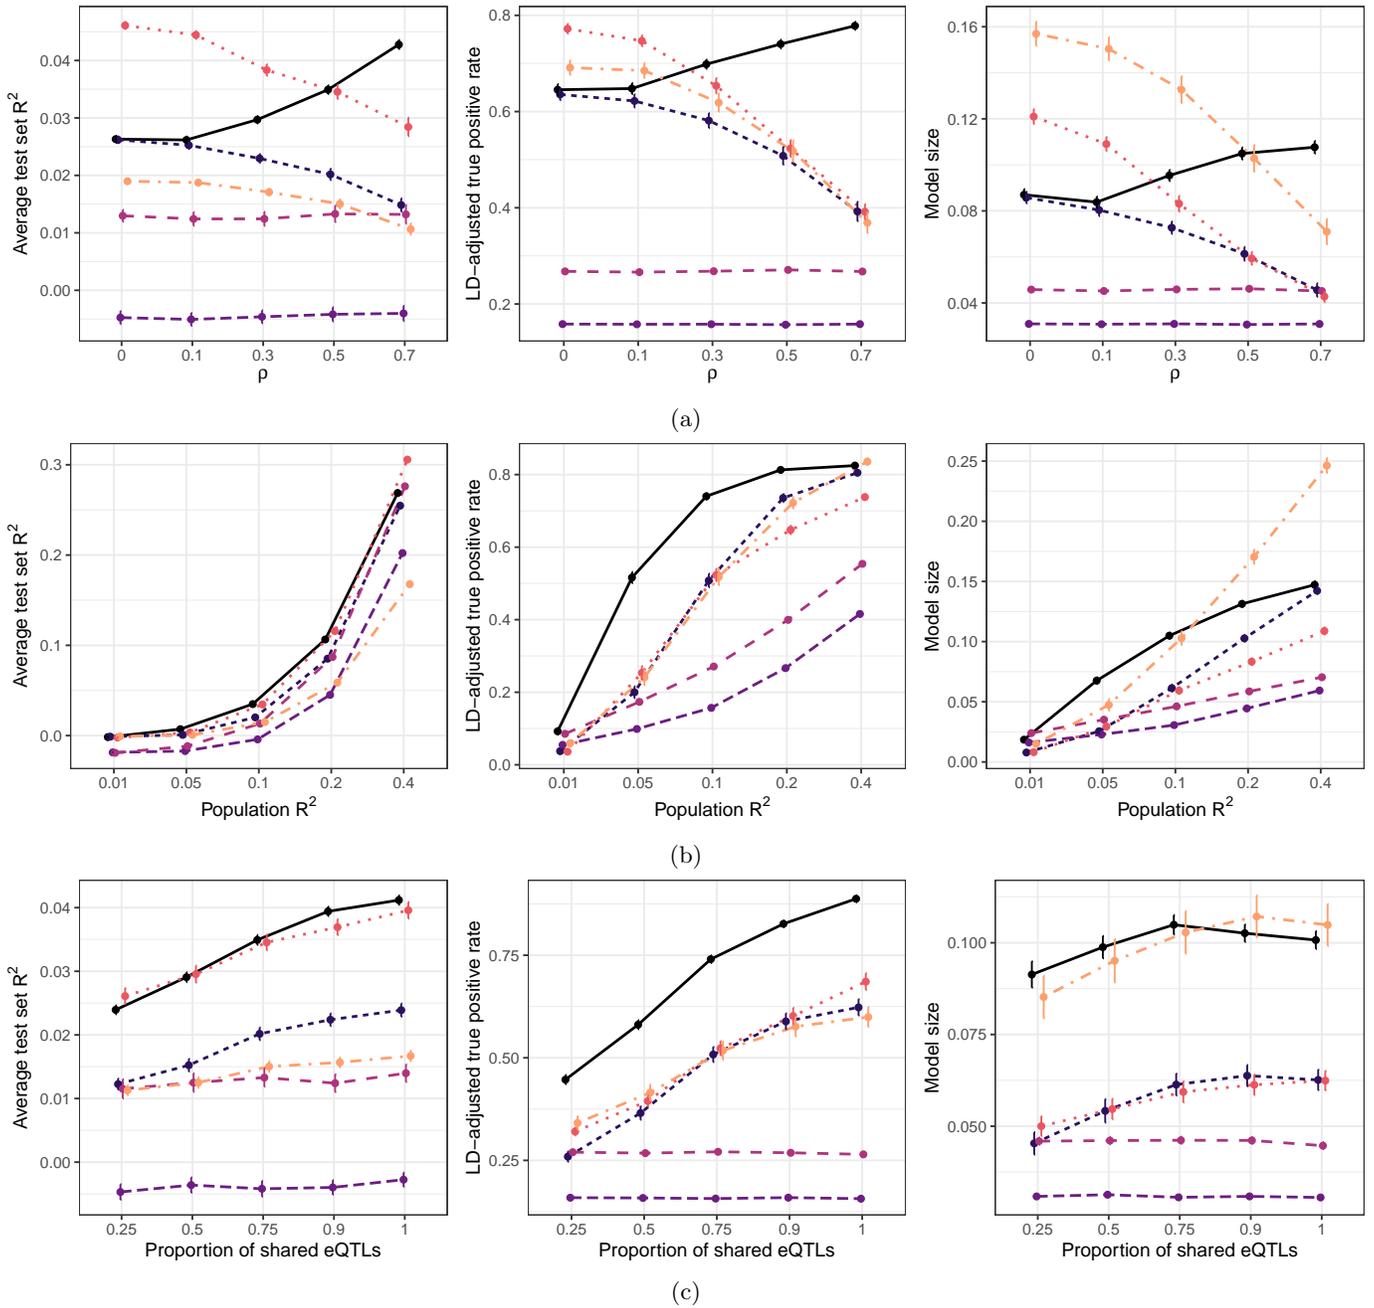


FIG 2. Average test set R^2 , average LD-adjusted true positive rate, and average model size (proportion) for six competing methods where: (a) ρ , the correlation of the errors varies; (b) the population R^2 varies; and (c) the proportion of the twenty eQTLs which are shared across all 29 tissues varies. Error bars denote two times the standard error for each method. Note that the spacing on the horizontal axes is not proportional to the difference in values represented on the horizontal axes. In the leftmost plot of the (b), error bars are not visible due to the range of the vertical axis. Note that throughout, the default settings were $R^2 = 0.1$, the proportional of shared eQTLs was equal to 0.75, and $\rho = 0.50$.

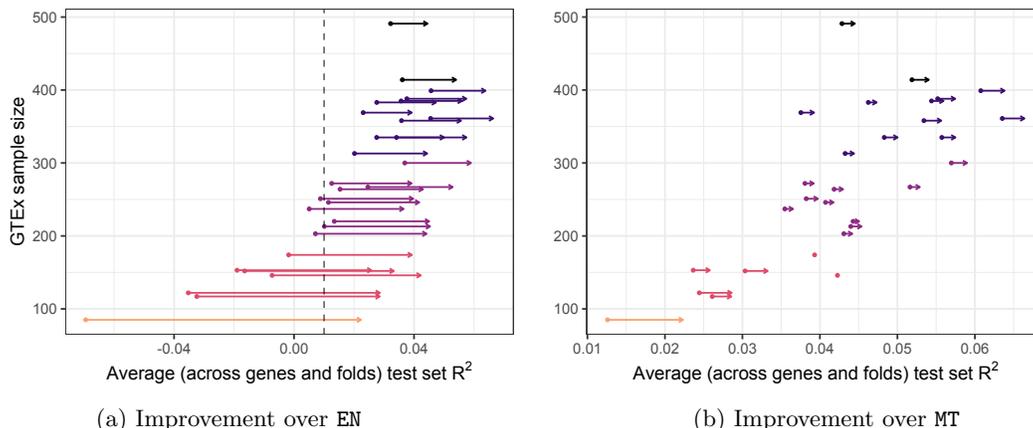


FIG 3. **Five-fold cross-validation prediction accuracy results from 29 tissues collected by GTEx.** (a, b): Left-out fold R^2 averaged across five folds and all genes. In (a), dots correspond to the average R^2 for one tissue estimated by EN. In (b), dots correspond to the average R^2 for one tissue estimated using MT. In both, arrow heads point to the average R^2 of the same tissue using Cov-MT, and colors correspond to the sample sizes partitioned into intervals of size 100. In (a), the vertical line denotes where the horizontal axis begins in the (b). The points with no visible arrow in (b) had differences less than 3×10^{-4} in magnitude.

one. Notably, our method performs similarly in terms of prediction accuracy to Or-MT, which is not applicable in practice. In addition, our method, Cov-MT, has higher LD-adjusted TPR than all other methods across all proportions of shared eQTLs. This suggests that accounting for cross-tissue dependence may also improve variable selection accuracy.

In Tables 1–3 of the Supplementary Material [26], we display summary statistics of the computing times for each setting considered in our simulation studies. For example, with ρ varying, median computing times for our method (including constructing candidate tuning parameter sets and validation) were 47.48, 52.61, 55.70, 60.81, and 80.30 minutes for $\rho \in \{0.0, 0.1, 0.3, 0.5, 0.7\}$, respectively, using six candidate values for λ_Ω , six for α , and 30 for λ_β for each unique pair (λ_Ω, α) . These experiments were run on HiperGator 2.0 at the University of Florida using a single CPU per replication. Of course, were computing times a concern for a practitioner, computing (2.3) over a grid of candidate tuning parameters could be trivially parallelized. More details on our implementation, including computing times for other simulation scenarios, can be found in Section 1 of the Supplementary Material.

4. Genome-wide multi-tissue joint eQTL mapping in GTEx. In our analysis of the GTEx data, we focus on 29 types of human tissues (see Figure 1). These are tissues which (i) are not sex-specific, (ii) had PEER factors and other covariates available from GTEx, and (iii) are not brain or pituitary gland tissues. Brain and pituitary tissues were omitted because subjects which had expression measured in brain tissue often had no expression measured in the other tissues and vice-versa. The gene expression we model is log-transformed RNA-seq counts which have been adjusted for PEER factors. Further details about our data processing are provided in the Supplementary Materials [26].

For each gene, we considered only local eQTLs (cis-SNPs with $\text{MAF} \geq .05$), which we defined as SNPs within 500kb of the transcription start or end site of the gene. For each gene, we further prune cis-SNPs until no two SNPs have absolute correlation greater than

0.95.

We performed joint eQTL mapping using the three approaches described in the Simulation studies subsection. Following a similar approach to Hu et al. [14], we measured prediction accuracy using a five-fold cross validation procedure. That is, each of the five folds was once treated as a testing fold. Of the remaining four folds, three were used to train the model and one was used as a validation set to select tuning parameters. Tuning parameters were selected to maximize the average R^2 on the validation set. We repeat this procedure for each of the five folds, with each fold once serving as testing fold and once serving as validation fold.

Results are displayed in Figure 3, (a) and (b). In Figure 3(a), our method significantly improves on the tissue-by-tissue elastic net in terms of prediction accuracy, especially for those tissue types with small sample sizes. For example, in the tissue with the smallest sample size (Minor Salivary Gland), the average testing-fold R^2 for tissue-by-tissue elastic net was less than -0.05 , indicating that this approach performed significantly worse than the null model in terms of expression prediction. This suggests that for small sample sizes, the tissue-by-tissue elastic net likely overfits to the training data. Our method, on the other hand, had average testing fold R^2 greater than 0.02 in all tissues, including Minor Salivary Gland. In Figure 3(b), we also compare Cov-MT to MT and see that incorporating an estimate of the precision matrix Ω_* improves testing fold R^2 averaged over genes and folds by 6.78% on the tissues we analyzed. Both MT and Cov-MT select tuning parameters to maximize validation set R^2 averaged over all tissues, which may partly explain why Cov-MT improves the R^2 on tissues with larger sample sizes – given that these tissues have the highest frequency in the validation folds, they will play the largest role in computing R^2 averaged over all tissues.

In Figure 4 of the Supplementary Material, we display differences in testing-fold R^2 averaged across all tissues and folds for each gene we analyzed. Our method improved on tissue-by-tissue elastic net for nearly every genes; whereas we improved over MT for a majority of genes. In particular, the summary statistics for the difference displayed in Figure 3(d) (Cov-MT average minus MT average) are: $\min = -.0244$, $Q_1 = 8.03 \times 10^{-5}$, $Q_2 = 0.00174$, $\text{Mean} = 0.00197$, $Q_3 = 0.00370$, and $\max = 0.0245$.

In Figure 5 of the Supplementary Material, we display a heatmap of how frequently our method estimated two tissues to be conditionally dependent. As one would expect, biologically related tissues often had nonzero estimated conditional correlations. For example, the three Esophagus tissues have some of the largest numbers of genes with nonzero tissue-tissue conditional correlations. Similarly, Adipose - Subcutaneous and Skin - Not Sun Exposed were often conditionally dependent, as were Skin - Not Sun Exposed and Skin - Sun Exposed. Interestingly, many tissues rarely had nonzero estimated conditional correlations with any other tissue: for example, see Small Intestine - Terminal Ileum, Liver, Cells - EBV-transformed lymphocytes, and Minor Salivary Gland.

Another important point of scientific interest is the frequency with which two tissue types share eQTLs. In Figure 4, we display a heatmap displaying the proportion of eQTLs shared between pairs of tissues. The most notable result is that our method (and the method of Hu et al. [14]) tend to estimate that the majority of eQTLs are shared across tissue types. For example, the first row of Figure 4 indicates that of all estimated Adipose - Subcutaneous eQTLs, approximately 96.3% were also estimated to be eQTLs for Adipose - Visceral and

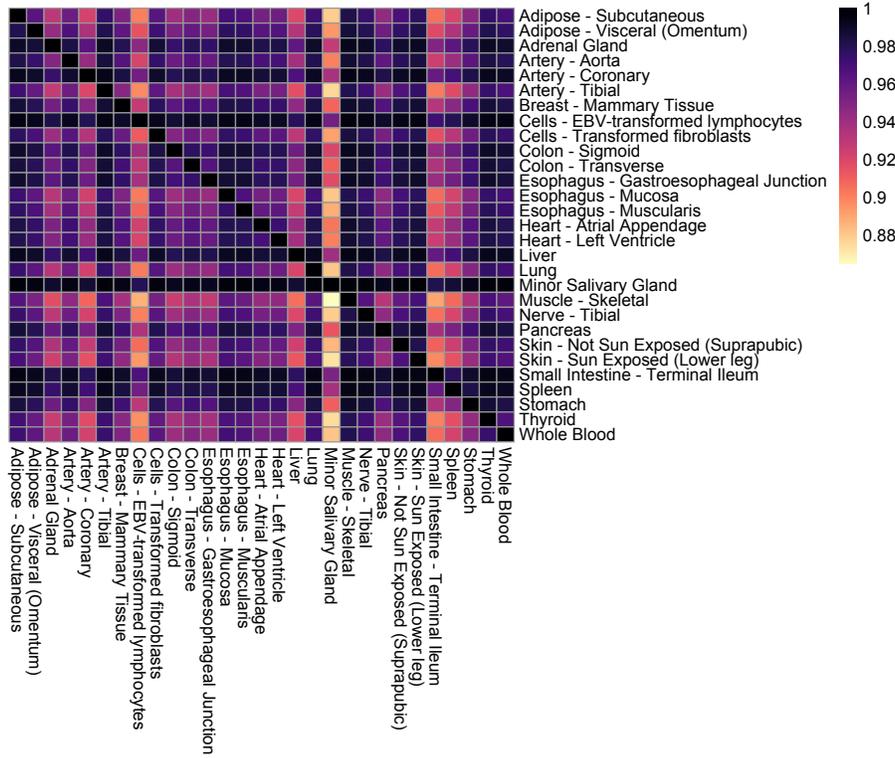


FIG 4. A heatmap of the proportions of eQTLs shared across tissue types. For a particular row, the intensities correspond to the proportions of total eQTLs for that tissue which are also an eQTL for the column-tissue. For example, the first row indicates that of all Adipose - Subcutaneous eQTLs, approximately 96.3% are also eQTLs for Adipose - Visceral and approximately 93.0% are also eQTLs for Adrenal Gland.

approximately 93.0% were also estimated to be eQTLs for Adrenal Gland. Notably, Minor Salivary Gland and Cells-EBV-transformed lymphocytes were the two tissues which had relatively low proportions of eQTLs shared with each of the other tissues (indicated by the light vertical bands). This may be partly due to the small sample sizes for both of these tissue types, which often led to fewer estimated eQTLs overall.

5. S-MultiXcan analysis of UKBiobank data.

5.1. *Multi-tissue transcriptome-wide association studies.* To demonstrate that our improved multi-tissue joint eQTL mapping method can lead to a higher number of novel gene-level TWAS discoveries, we performed a summary-MultiXcan (S-MultiXcan) analysis following the method proposed in Barbeira et al. [2]. To obtain eQTL weights from the full GTEx dataset, we re-estimated these coefficients using the complete dataset based on the tuning parameters which had the largest average left-out fold R^2 . We then downloaded UKBiobank GWAS summary statistics from the Neale [29] lab for four complex traits: two binary (asthma and heart attack) and two continuous (haemoglobin and platelet count). These traits were selected as we thought no single tissue-type was an obvious candidate for a single-tissue summary-PrediXcan analysis. To compute the LD-matrices needed to perform

the S-MultiXcan analysis, we used genome-wide genotype data obtained from the Genetics and Epidemiology of Colorectal Cancer Consortium (GECCO), which was imputed using the Haplotype Reference Consortium (HRC) reference panel. For each of the three eQTL weight set estimation methods, we tested only those genes which had test set R^2 averaged over all folds and tissues greater than zero. The number of genes tested were 10364, 9393, 4608 for Cov-MT, MT, and EN respectively. When testing for association with each phenotype, we adjust for multiple tests using a Bonferonni correction. Thus, EN, which has a smaller number of genes tested, also has a more liberal p-value cutoff.

For each gene, we recorded the S-MultiXcan p-value based on the weights computed using our method, the exact version of the method proposed by Hu et al. [14], and tissue-by-tissue elastic net (Cov-MT, MT, and EN, respectively). To validate our findings, we also recorded (a) the minimum GWAS p-value amongst the SNPs which had nonzero weights across the methods considered for each gene, and (b) the minimum GWAS p-value of any SNP with 500kb of the gene transcription start or end site. This way, for any S-MultiXcan discovery, we verified whether this discovery could be attributed to a genome-wide significant eQTL or SNP (based on (a) or (b), respectively) defined as p-value $< 5 \times 10^{-8}$ or represents a potentially novel finding. In Table 1, we display the total number of discoveries and the number of novel discoveries for two binary traits and two continuous traits.

The weight set obtained using MT, the exact version of the method propose by Hu et al. [14], tended to include a larger number of SNPs. Conversely, our method, which yielded a slightly smaller set of SNPs, has a similar or larger number of significant discoveries than MT. Further, our method also led to more discoveries which could not be attributed to a GWAS associated eQTL or SNP genotype. Focusing on the phenotype heart attack, our method identified 16 significant genes: nine did not have an estimated eQTL reaching genome-wide significance, and one had no cis-SNPs reaching genome-wide significance. In comparison, MT identified 14 genes, among which six and zero gene(s) had no estimated eQTLs or cis-SNPs reaching genome-wide significance, respectively. These genes are listed in Table 2, along with phenotypes with which these genes have been associated in previous GWAS studies (p-value $< 5 \times 10^{-8}$) [3]. Notably, most of these genes are associated with phenotypes related with heart diseases such as coronary artery disease, cholesterol, high-density lipoproteins (HDL), low-density lipoproteins (LDL), and systolic blood pressure. When using the more stringent definition of a novel finding based on the p-value of all cis-SNP genotypes, our method still identifies a higher number of novel discoveries for heart attack, platelet count, and haemoglobin count. In the following section, we further discuss the genes associated with heart attack which could not be attributed to a genome-wide significant eQTL.

5.2. S-MultiXcan results. We have identified multiple genes associated with the occurrence of a heart attack that would be missed in a SNP-by-SNP association analysis. Interestingly, many of these genes are associated with phenotypes related to heart attack in previous GWAS, or are known to have biological functions associated with the occurrence of a heart attack and coronary artery disease. A heart attack occurs when an artery supplying the heart with blood and oxygen becomes blocked. This is closely related to more broadly defined coronary artery disease, which is the narrowing or blockage of the coronary arteries that leads to reduction of the amount of oxygen and nutrients delivered to the heart. Coronary artery disease tends to develop when cholesterol or fatty deposits builds along the

TABLE 1

Number of significant discoveries using different eQTL weights for two binary and two continuous phenotypes. The bottommost two rows (with superscript *) denote the number of discoveries which did not have an estimated eQTL reaching genome-wide significance (left) and did not have any cis-SNP reaching genome-wide significance (right).

	Asthma	Heart Attack	Haemoglobin	Platelet
Cov-MT	7	16	751	1255
MT	8	14	734	1232
EN	11	9	586	934
Cov-MT*	4/2	9/1	150/56	182/45
MT*	3/2	6/0	146/54	162/39
EN*	3/1	3/1	81/50	78/39

TABLE 2

Genes associated with Heart Attack discovered in the S-MultiXcan analysis which could not be attributed to a GWAS associated eQTL. In the rightmost two columns, we provide the phenotypes with which the particular gene has been associated before according to the GWAS Catalog as of November 20th, 2020 [3] for all the associations with p -value $< 5 \times 10^{-8}$. The phenotypes for each gene are ranked by the number of times they are associated with this gene (indicated by the numbers in the parenthesis next to the phenotype) and only top two phenotypes are shown for each gene. The bold/underlined gene was that wherein no SNP within 500kb of TSS or TES was genome-wide significant (including those not identified as eQTLs).

Gene	Chr	Region	Associated phenotype 1	Associated phenotype 2
<i>AIDA</i>	1	222668013-222713210	Coronary artery disease (1)	
<i>FAM117B</i>	2	202634969-202769757	Total cholesterol (2)	LDL (2)
<i>ICA1L</i>	2	202773150-202871985	White matter microstructure (7)	Urinary albumin-to-creatinine ratio (3)
<i>NBEAL1</i>	2	203014879-203226378	Coronary artery disease (4)	Body mass index (2)
<i>WDR12</i>	2	203738984-203879521	Coronary artery disease (3)	White matter microstructure (3)
<i>LPL</i>	8	19901717-19967259	Triglycerides (32)	HDL (18)
<i>ZBTB39</i>	12	56998836-57006546	Asthma (3)	Allergic disease (1)
<i>FURIN</i>	15	90868588-90883458	Systolic blood pressure (8)	Schizophrenia (6)
<i>FES</i>	15	90883695-90895776	Systolic blood pressure (7)	Diastolic blood pressure (4)

artery walls.

Endothelial dysfunction is a hallmark of coronary artery disease and one gene identified in our analysis, *AIDA*, was identified as a coronary artery disease candidate gene by integrative analysis of vascular endothelial cell genomic features [16]. Several genes we have identified are involved in different biological processes related with cancer development, implying some connections due to the hypoxia environment shared by narrowed or blocked coronary artery and tumor micro-environment. *FES*, for example, is a well-known cancer related gene involved in signaling for cell growth. A recent report shows that *FURIN* inhibits apoptosis [52]. Finally, there is evidence that *ZBTB39* is associated with red cell distribution width [35, 15], a prognostic marker in both acute and chronic heart failure [43, 47]. In Section 3 of the Supplementary Material [26], we discuss how these findings may be used for investigating disease mechanisms in future studies. The p-values for all genes we tested are available for download at github.com/ajmolstad/MTEQTLResults and as part of the Supplementary Material [25].

6. Discussion. In this article, we have proposed a new method for obtaining multi-tissue eQTL weights. While our method was motivated by the growing popularity of multi-tissue TWAS using genetically predicted gene expression, it is notable that compared to tissue-by-tissue elastic net, our method yielded eQTL weights which had higher prediction accuracy in every individual tissue we studied. This suggests that even single-tissue PrediXcan analyses could be improved using our estimated weight set. Of course, when to use a multi-tissue test versus a single-tissue test remains an unresolved and important question. Our analyses focused on phenotypes for which no individual tissue seemed an obvious candidate for analyses.

Another natural application of our method is for imputing unmeasured gene expression, e.g., as was the goal of Wang et al. [46]. Specifically, Wang et al. [46] focused on the case of imputing missing expression in GTEx in individuals with expression measured in a subset of tissues. Our method naturally applies to this problem as the conditional expectation of the missing tissues' expression from (2.6) is easy to compute given estimates of eQTL weights and the cross-tissue error covariance matrix. Furthermore, prediction (ellipsoids) intervals could be constructed using our estimates of the conditional (co)variance of gene expression. Unlike Wang [45], who preselect eQTLs to be used in their prediction model, our approach estimates eQTLs and fits the prediction model jointly.

Finally, it would be beneficial to extend our methodology to allow for more heavy-tailed error distributions. For example, one could relax the normality assumption in (2.1) and assume that errors have a multivariate t -distribution. One interesting direction along these lines would modify the method of Chen et al. [5] to handle missing data, which is highly nontrivial computationally.

Supplementary Material. Appendix. We provide more information about our GTEx data processing; the complete algorithms used to solve both (2.8) and the exact version of the estimator of Hu et al. [14]; timing results and additional details about our implementation; and a simulation study analyzing the the proportion of false positives in PrediXcan under varying prediction accuracy. **Code and results.** We provide a user-guide for accessing our estimated eQTL weights in order to perform multi-tissue MultiXcan analyses using our weight set. Comprehensive results presented in the article, as well as software to reproduce the simulation results, are available for download as part of the Supplementary Material and at github.com/ajmolstad/MTeQTLResults.

Acknowledgments. The work in this article is funded in part by the National Institutes of Health (CA189532, CA195789, GM105785). The Genotype-Tissue Expression (GTEx) Project was supported by the Common Fund of the Office of the Director of the National Institutes of Health, and by NCI, NHGRI, NHLBI, NIDA, NIMH, and NINDS. The data used for the analyses described in this manuscript were obtained from the GTEx Portal (gtexportal.org/home/, v7) on 04/25/2018.

References.

- [1] Barbeira, A. N., Dickinson, S. P., Bonazzola, R., Zheng, J., Wheeler, H. E., Torres, J. M., Torstenson, E. S., Shah, K. P., Garcia, T., and Edwards, T. L. (2018). Exploring the phenotypic consequences of tissue specific gene expression variation inferred from GWAS summary statistics. *Nature Communications*, 9(1):1825.

- [2] Barbeira, A. N., Pividori, M. D., Zheng, J., Wheeler, H. E., Nicolae, D. L., Im, H. K., et al. (2019). Integrating predicted transcriptome from multiple tissues improves association detection. *PLoS Genetics*, 15(1):e1007889.
- [3] Buniello, A., MacArthur, J. A. L., Cerezo, M., Harris, L. W., Hayhurst, J., Malangone, C., McMahon, A., Morales, J., Mountjoy, E., Sollis, E., et al. (2018). The NHGRI-EBI GWAS catalog of published genome-wide association studies, targeted arrays and summary statistics 2019. *Nucleic Acids Research*, 47(D1):D1005–D1012.
- [4] Buniello, A., MacArthur, J. A. L., Cerezo, M., Harris, L. W., Hayhurst, J., Malangone, C., McMahon, A., Morales, J., Mountjoy, E., Sollis, E., et al. (2019). The NHGRI-EBI GWAS Catalog of published genome-wide association studies, targeted arrays and summary statistics 2019. *Nucleic Acids Research*, 47(D1):D1005–D1012.
- [5] Chen, L., Pourahmadi, M., and Maadooliat, M. (2014). Regularized multivariate regression models with skew-t error distributions. *Journal of Statistical Planning and Inference*, 149:125–139.
- [6] Cheung, V. G., Conlin, L. K., Weber, T. M., Arcaro, M., Jen, K. Y., Morley, M., and Spielman, R. S. (2003). Natural variation in human gene expression assessed in lymphoblastoid cells. *Nature Genetics*, 33(3):422–425.
- [7] Flutre, T., Wen, X., Pritchard, J., and Stephens, M. (2013). A statistical framework for joint eQTL analysis in multiple tissues. *PLoS Genetics*, 9(5):e1003486.
- [8] Gamazon, E. R., Wheeler, H. E., Shah, K. P., Mozaffari, S. V., Aquino-Michaels, K., Carroll, R. J., Eyler, A. E., Denny, J. C., Nicolae, D. L., Cox, N. J., et al. (2015). A gene-based association method for mapping traits using reference transcriptome data. *Nature Genetics*, 47(9):1091.
- [9] GTEx Consortium (2015). The genotype-tissue expression (GTEx) pilot analysis: multitissue gene regulation in humans. *Science*, 348(6235):648–660.
- [10] GTEx Consortium (2017). Genetic effects on gene expression across human tissues. *Nature*, 550(7675):204.
- [11] Gusev, A., Ko, A., Shi, H., Bhatia, G., Chung, W., Penninx, B. W., Jansen, R., de Geus, E. J., et al. (2016). Integrative approaches for large-scale transcriptome-wide association studies. *Nature Genetics*, 48(3):245–252.
- [12] Hindorf, L. A., Sethupathy, P., Junkins, H. A., Ramos, E. M., Mehta, J. P., Collins, F. S., and Manolio, T. A. (2009). Potential etiologic and functional implications of genome-wide association loci for human diseases and traits. *Proc. Natl. Acad. Sci. U.S.A.*, 106(23):9362–9367.
- [13] Hsieh, C.-J., Sustik, M. A., Dhillon, I. S., and Ravikumar, P. (2014). Quic: quadratic approximation for sparse inverse covariance estimation. *The Journal of Machine Learning Research*, 15(1):2911–2947.
- [14] Hu, Y., Li, M., Lu, Q., Weng, H., Wang, J., Zekavat, S. M., Yu, Z., Li, B., Gu, J., Muchnik, S., et al. (2019). A statistical framework for cross-tissue transcriptome-wide association analysis. *Nature Genetics*, 51(3):568–576.
- [15] Kichaev, G., Bhatia, G., Loh, P.-R., Gazal, S., Burch, K., Freund, M. K., Schoech, A., Pasaniuc, B., and Price, A. L. (2019). Leveraging polygenic functional enrichment to improve gwas power. *The American Journal of Human Genetics*, 104(1):65–75.
- [16] Lalonde, S., Codina-Fauteux, V.-A., de Bellefon, S. M., Leblanc, F., Beaudoin, M., Simon, M.-M., Dali, R., Kwan, T., Lo, K. S., Pastinen, T., et al. (2019). Integrative analysis of vascular endothelial cell genomic features identifies AIDA as a coronary artery disease candidate gene. *Genome Biology*, 20(1):1–13.
- [17] Lange, K. (2016). *MM optimization algorithms*, volume 147. SIAM.
- [18] Lee, W. and Liu, Y. (2012). Simultaneous multiple response regression and inverse covariance matrix estimation via penalized Gaussian maximum likelihood. *Journal of Multivariate Analysis*, 111:241–255.
- [19] Li, G., Shabalin, A. A., Rusyn, I., Wright, F. A., and Nobel, A. B. (2017). An empirical bayes approach for multiple tissue eqtl analysis. *Biostatistics*, 19(3):391–406.
- [20] Mancuso, N., Shi, H., Goddard, P., Kichaev, G., Gusev, A., and Pasaniuc, B. (2017). Integrating gene expression with summary association statistics to identify genes associated with 30 complex traits. *The American Journal of Human Genetics*, 100(3):473–487.
- [21] Manor, O. and Segal, E. (2015). GenoExp: a web tool for predicting gene expression levels from single nucleotide polymorphisms. *Bioinformatics*, 31(11):1848–1850.
- [22] Maurano, M. T., Humbert, R., Rynes, E., Thurman, R. E., Haugen, E., Wang, H., Reynolds, A. P., Sandstrom, R., et al. (2012). Systematic localization of common disease-associated variation in regulatory DNA. *Science*, 337(6099):1190–1195.

- [23] Meng, X.-L. and Rubin, D. B. (1993). Maximum likelihood estimation via the ECM algorithm: A general framework. *Biometrika*, 80(2):267–278.
- [24] Molstad, A. J. and Rothman, A. J. (2018). Shrinking characteristics of precision matrix estimators. *Biometrika*, 105(3):563–574.
- [25] Molstad, A. J., Sun, W., and Hsu, L. (2021a). Code and results for “A covariance-enhanced approach to multi-tissue joint eQTL mapping with application to transcriptome-wide association studies”.
- [26] Molstad, A. J., Sun, W., and Hsu, L. (2021b). Supplement to “A covariance-enhanced approach to multi-tissue joint eQTL mapping with application to transcriptome-wide association studies”.
- [27] 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.
- [28] Morgan, A., Vuckovic, D., Krishnamoorthy, N., Rubinato, E., Ambrosetti, U., Castorina, P., Franzè, A., Vozzi, D., La Bianca, M., Cappellani, S., et al. (2019). Next-generation sequencing identified *spatc11* as a possible candidate gene for both early-onset and age-related hearing loss. *European Journal of Human Genetics*, 27(1):70–79.
- [29] Neale, B. (2019). UKBiobank GWAS summary statistics.
- [30] Nicolae, D. L., Gamazon, E., Zhang, W., Duan, S., Dolan, M. E., and Cox, N. J. (2010). Trait-associated SNPs are more likely to be eQTLs: annotation to enhance discovery from GWAS. *PLoS Genetics*, 6(4):e1000888.
- [31] Parikh, N. and Boyd, S. (2014). Proximal algorithms. *Foundations and Trends in Optimization*, 1(3):127–239.
- [32] Pavlides, J. M., Zhu, Z., Gratten, J., McRae, A. F., Wray, N. R., and Yang, J. (2016). Predicting gene targets from integrative analyses of summary data from GWAS and eQTL studies for 28 human complex traits. *Genome Medicine*, 8(1):84.
- [33] Peng, J., Zhu, J., Bergamaschi, A., Han, W., Noh, D.-Y., Pollack, J. R., Wang, P., et al. (2010). Regularized multivariate regression for identifying master predictors with application to integrative genomics study of breast cancer. *The Annals of Applied Statistics*, 4(1):53–77.
- [34] Pickrell, J. K., Marioni, J. C., Pai, A. A., Degner, J. F., Engelhardt, B. E., Nkadori, E., Veyrieras, J. B., Stephens, M., et al. (2010). Understanding mechanisms underlying human gene expression variation with RNA sequencing. *Nature*, 464(7289):768–772.
- [35] Pilling, L. C., Atkins, J. L., Duff, M. O., Beaumont, R. N., Jones, S. E., Tyrrell, J., Kuo, C.-L., Ruth, K. S., Tuke, M. A., Yaghootkar, H., et al. (2017). Red blood cell distribution width: Genetic evidence for aging pathways in 116,666 volunteers. *PLoS One*, 12(9):e0185083.
- [36] Polson, N. G., Scott, J. G., Willard, B. T., et al. (2015). Proximal algorithms in statistics and machine learning. *Statistical Science*, 30(4):559–581.
- [37] Rothman, A. J., Bickel, P. J., Levina, E., Zhu, J., et al. (2008). Sparse permutation invariant covariance estimation. *Electronic Journal of Statistics*, 2:494–515.
- [38] Rothman, A. J., Levina, E., and Zhu, J. (2010). Sparse multivariate regression with covariance estimation. *Journal of Computational and Graphical Statistics*, 19(4):947–962.
- [39] Schadt, E. E., Monks, S. A., Drake, T. A., Luskis, A. J., Che, N., Colinayo, V., Ruff, T. G., Milligan, S. B., et al. (2003). Genetics of gene expression surveyed in maize, mouse and man. *Nature*, 422(6929):297–302.
- [40] Simon, N., Friedman, J., Hastie, T., and Tibshirani, R. (2013). A sparse-group lasso. *Journal of Computational and Graphical Statistics*, 22(2):231–245.
- [41] Srivastava, V. K. and Giles, D. E. (1987). Seemingly unrelated regression models: Estimation and inference.
- [42] Stranger, B. E., Forrest, M. S., Clark, A. G., Minichiello, M. J., Deutsch, S., Lyle, R., Hunt, S., Kahl, B., Antonarakis, S. E., Tavare, S., Deloukas, P., and Dermitzakis, E. T. (2005). Genome-wide associations of gene expression variation in humans. *PLoS Genetics*, 1(6):e78.
- [43] Tseliou, E., Terrovitis, J. V., Kaldara, E. E., Ntalianis, A. S., Repasos, E., Katsaros, L., Margari, Z. J., Matsouka, C., Toumanidis, S., Nanas, S. N., et al. (2014). Red blood cell distribution width is a significant prognostic marker in advanced heart failure, independent of hemoglobin levels. *Hellenic J Cardiol*, 55(6):457–461.
- [44] Visscher, P. M., Hill, W. G., and Wray, N. R. (2008). Heritability in the genomics era—concepts and misconceptions. *Nature Reviews Genetics*, 9(4):255–266.
- [45] Wang, J. (2015). Joint estimation of sparse multivariate regression and conditional graphical models.

- Statistica Sinica*, pages 831–851.
- [46] Wang, J., Gamazon, E. R., Pierce, B. L., Stranger, B. E., Im, H. K., Gibbons, R. D., Cox, N. J., Nicolae, D. L., and Chen, L. S. (2016). Imputing gene expression in uncollected tissues within and beyond GTEx. *The American Journal of Human Genetics*, 98(4):697–708.
- [47] Xanthopoulos, A., Giamouzis, G., Melidonis, A., Kitai, T., Paraskevopoulou, E., Paraskevopoulou, P., Patsilina, S., Triposkiadis, F., and Skoularigis, J. (2017). Red blood cell distribution width as a prognostic marker in patients with heart failure and diabetes mellitus. *Cardiovascular Diabetology*, 16(1):81.
- [48] Yin, J. and Li, H. (2011). A sparse conditional gaussian graphical model for analysis of genetical genomics data. *The Annals of Applied Statistics*, 5(4):2630.
- [49] Yuan, M. and Lin, Y. (2007). Model selection and estimation in the Gaussian graphical model. *Biometrika*, 94(1):19–35.
- [50] Zellner, A. (1962). An efficient method of estimating seemingly unrelated regressions and tests for aggregation bias. *Journal of the American Statistical Association*, 57(298):348–368.
- [51] Zeng, P., Zhou, X., and Huang, S. (2017). Prediction of gene expression with cis-SNPs using mixed models and regularization methods. *BMC Genomics*, 18(1):368.
- [52] Zhao, G., Yang, W., Wu, J., Chen, B., Yang, X., Chen, J., McVey, D. G., Andreadi, C., Gong, P., Webb, T. R., et al. (2018). Influence of a coronary artery disease-associated genetic variant on FURIN expression and effect of FURIN on macrophage behavior. *Arteriosclerosis, Thrombosis, and Vascular Biology*, 38(8):1837–1844.

DEPARTMENT OF STATISTICS
GENETICS INSTITUTE
102 GRIFFIN-FLOYD HALL
P.O. BOX 118545
UNIVERSITY OF FLORIDA
GAINESVILLE, FL 32611
E-MAIL: amolstad@ufl.edu

BIostatistics PROGRAM
FRED HUTCHINSON CANCER RESEARCH CENTER
1100 FAIRVIEW AVE. N., M2-B500
SEATTLE, WA 98109
E-MAIL: wsun@fredhutch.org
lih@fredhutch.org