



Published in final edited form as:

Nat Genet. 2019 January ; 51(1): 187–195. doi:10.1038/s41588-018-0268-8.

Flexible statistical methods for estimating and testing effects in genomic studies with multiple conditions

Sarah M. Uebachs^{1,2}, Gao Wang², Peter Carbonetto^{2,4}, and Matthew Stephens^{2,3,*}

¹Pritzker School of Medicine, Growth & Development Training Program, University of Chicago, Chicago, IL, USA.

²Department of Human Genetics, University of Chicago, Chicago, IL, USA.

³Department of Statistics, University of Chicago, Chicago, IL, USA.

⁴Research Computing Center, University of Chicago, Chicago, IL, USA.

Abstract

We introduce new statistical methods for analyzing genomic datasets that measure many effects in many conditions (e.g., gene expression changes under many treatments). These new methods improve on existing methods by allowing for arbitrary correlations in effect sizes among conditions. This flexible approach increases power, improves effect estimates, and allows for more quantitative assessments of effect-size heterogeneity compared to simple “shared/condition-specific” assessments. We illustrate these features through an analysis of locally-acting variants associated with gene expression (“*cis* eQTLs”) in 44 human tissues. Our analysis identifies more eQTLs than existing approaches, consistent with improved power. We show that while genetic effects on expression are extensively shared among tissues, effect sizes can still vary greatly among tissues. Some shared eQTLs show stronger effects in subsets of biologically related tissues (e.g., brain-related tissues), or in only one tissue (e.g., testis). Our methods are widely applicable, computationally tractable for many conditions, and available online.

Users may view, print, copy, and download text and data-mine the content in such documents, for the purposes of academic research, subject always to the full Conditions of use:http://www.nature.com/authors/editorial_policies/license.html#terms

* mstephens@uchicago.edu.

Author contributions S.M.U. and M.S. conceived of the project and developed the statistical methods. S.M.U. implemented the comparisons with simulated data. S.M.U. and G.W. performed the analyses of the GTEx data, and additional analyses. S.M.U., G.W. and M.S. implemented the software, with contributions from P.C. S.M.U. and M.S. wrote the manuscript, with input from G.W. and P.C. P.C. and G.W. prepared the online code and data resources.

Data availability.

The GTEx study data are available through dbGaP under accession phs000424.v6.p1. The GTEx summary statistics used in the mash analysis have been deposited in Zenodo (doi:10.5281/zenodo.1296399).

Code availability.

Source code implementing our methods, including processing and analysis of the GTEx data, adaptive shrinkage (“ash”) and multivariate adaptive shrinkage (“mash”, “mash-bmalite”), is available online, and is distributed under GNU and BSD open source licenses (see URLs).

Competing interests The authors declare no competing financial interests.

Introduction

Genomic studies often involve estimating and comparing many effects across multiple conditions or outcomes. Examples include studying changes in expression of many genes under multiple treatments¹, differences in histone methylation at many genomic locations in multiple cell lines², effects of many genetic variants on risk of multiple diseases³, and effects of expression quantitative trait loci (eQTLs) in multiple cell-types or tissues^{4–6}. Analyses often aim to identify “significant” non-zero effects, and to identify differences in effects among conditions (e.g., tissue-specific effects), which may yield biological insights.

A common analysis strategy for such studies is to separately analyze each condition in turn, then compare the “significant” results among conditions. Although appealingly simple, this “condition-by-condition” approach is unsatisfactory in several ways: it under-represents sharing of effects among conditions because shared effects will be insignificant in some conditions by chance, and it misses the power gains that come from sharing information across conditions⁵.

To address deficiencies of condition-by-condition analyses, several groups have developed methods for *joint* analysis of multiple conditions^{2–3,5–15}. The simplest methods build on traditional meta-analysis methodology^{8,9}, and assume that non-zero effects are shared across all conditions. Other methods are more flexible, allowing for condition-specific effects, for sharing of effects among subsets of conditions, and for heterogeneity in the shared effects^{5,6,12}. The most flexible methods adapt themselves to each dataset by learning patterns of sharing using a hierarchical model⁵.

Nonetheless, existing methods have important limitations. First, they make restrictive assumptions about the correlations of non-zero effects among conditions. For example, Flutre *et al.*⁵ assume these correlations are non-negative and equal. Correlations may be negative in some applications; e.g., genetic variants that increase one trait may decrease another. And some conditions may be more correlated than others; for example, in our eQTL application (below), brain tissues are strongly correlated with one another. Second, the most flexible methods are computationally intractable for moderate numbers of conditions (e.g., 44 tissues in our eQTL application). Existing solutions to this problem substantially reduce flexibility. For example, Flutre *et al.*⁵ solve the computational problem by restricting effects to be shared in all conditions, or to be specific to a single condition. Alternatively, Wei *et al.*¹² allow for all possible patterns of sharing, but only under the restrictive assumption that non-zero effects are uncorrelated among conditions. Third, existing methods typically focus only on *testing* for significant effects in each condition, and not on *estimating effect sizes* which, as we illustrate, is important for assessing heterogeneity among conditions.

We introduce more flexible statistical methods that combine the most attractive features of existing approaches, while overcoming their major limitations. The methods, which we refer to as “multivariate adaptive shrinkage” (mash), build on recent approaches¹⁶ for testing and estimating effects in a *single* condition, extending these approaches to *multiple* conditions. Key features of mash include: (i) it is *flexible*, allowing for both shared and condition-specific effects, and arbitrary patterns of correlation among conditions; (ii) it is

computationally tractable for hundreds of thousands of tests in dozens of conditions, or more; (iii) it provides not only measures of significance, but also *estimates of effect sizes*, together with measures of uncertainty; (iv) it is *adaptive*, meaning that its behavior adapts to the patterns present in the data; and (v) it is *generic*, requiring only a matrix containing the observed effects in each condition and a matrix of the corresponding standard errors. (Alternatively, mash can be supplied with just matrix of Z scores, although this reduces the ability to estimate effect sizes.) Together, these features make mash the most flexible and widely applicable method available for estimating and testing multiple effects in multiple conditions.

To demonstrate the potential for mash to provide novel insights, we apply it to analyze *cis* eQTL effects in 16,069 genes across 44 human tissues. Focusing on the strongest *cis* eQTLs, we find that while most eQTLs are shared by many tissues, effect sizes can still vary considerably. Our results suggest that when assessing effects that are “tissue-specific” versus “tissue-consistent”, we should pay careful attention to the *sizes* of effects, and not only the tests for significance.

Results

Methods overview.

Our method, mash, estimates effects of many “units” (J) in many conditions (R), allowing that effects may be sparse (*i.e.*, many zero effects), and allowing for correlations among non-zero effects in different conditions. For example, in multi-tissue eQTL studies, the “units” are eQTLs ($J > 10,000$), the conditions are different tissues ($R = 44$), and mash estimates the effect of each eQTL in each tissue, allowing for cross-tissue sharing and tissue-specificity of eQTLs.

To apply mash, the analyst must first conduct a condition-by-condition analysis to obtain an effect estimate and corresponding standard error for each unit in each condition. These estimates are the inputs for a two-step Empirical Bayes (EB) procedure: (1) learn patterns of sparsity, sharing and correlations among effects from the condition-by-condition results; (2) combine these learned patterns with the condition-by-condition results to produce improved effect estimates and corresponding measures of significance. These steps are summarized here, and in Fig. 1. (See Online Methods for more details.)

In brief, let \mathbf{b} denote the vector of true effects for a single unit across R conditions. We capture correlations and sharing of effects among conditions using a mixture model,

$$p(\mathbf{b}; \boldsymbol{\pi}, \mathbf{U}) = \sum_{k=1}^K \sum_{l=1}^L \pi_{k,l} N_R(\mathbf{b}; \boldsymbol{\theta}, \omega_l \mathbf{U}_k), \quad (1)$$

where $N_R(\cdot; \boldsymbol{\mu}, \Sigma)$ denotes the multivariate normal density in R dimensions with mean $\boldsymbol{\mu}$ and variance-covariance matrix Σ ; each \mathbf{U}_k is a covariance matrix that captures a pattern of effects; each ω_l is a scaling coefficient that corresponds to a different effect size; and the

mixture proportions $\pi_{k,l}$ determine the relative frequency of each covariance-scale combination. The scaling coefficients ω_l take values on a fixed dense grid that spans from “very small” to “very large”, capturing the full range of possible effects.

Step 1 of mash learns patterns of sparsity, sharing and correlations by estimating covariance matrices U and mixture proportions π in two sub-steps:

Step 1a: Generate candidate covariance matrices $U = (U_1, \dots, U_K)$. This list includes both “data-driven” matrices that are estimated from the strongest signals in the condition-by-condition results (see Online Methods), and “canonical” matrices that have simple interpretations. For example, the “canonical” matrices include the identity matrix (representing independent effects across all conditions); a matrix of all ones (representing effects that are equal in all conditions); and R matrices that represent effects that are specific to each condition.

Step 1b: Given U , estimate π by maximum likelihood from the condition-by-condition results. This step can rescue imperfections in Step 1a by assigning low weight to covariance matrices that are not well-supported by the data. This step also adapts to sparse effects; if most effects are zero, or small, this step will put most weight on small effects (*i.e.*, small scaling coefficients, ω).

Step 2 of mash uses Bayes’ theorem to compute the posterior distribution for each effect given the condition-by-condition results and the fitted prior (eq. 1). These posterior distributions yield improved effect estimates—posterior means and standard deviations—that account for sparsity and correlations among effects. We use these estimates to define quantitative measures of effect sharing between any two conditions: “sharing by sign” (effects have the same sign), and “sharing by magnitude” (effects have similar magnitude—here defined to be within a factor of 2, although other thresholds could be used; see Supplementary Note). The posterior distributions also yield a condition-specific measure of significance for each effect, the “local false sign rate”, or $lfsr^{16}$, that is analogous to a false discovery rate, but more stringent because it requires true discoveries to be not only non-zero, but also correctly signed. Finally, mash also computes, for each unit, a Bayes Factor that summarizes the overall significance of the unit—*i.e.*, overall evidence for a non-zero effect in *any* condition.

Simulation studies.

We ran simulations to compare mash with existing methods. We simulated effects for 20,000 units in 44 conditions, with 400 units having non-null effects. Non-null effects were simulated under two different scenarios:

1. “Shared, structured effects”: Effects are largely shared (non-null units have an effect in many conditions), and “structured”—that is, similar in size and direction, with greater similarity among some subsets of conditions. This scenario was based on the fit of the model (eq. 1) to the eQTL data (see Online Methods).

2. “Shared, unstructured effects”: Effects are shared (non-null units have an effect in every condition) but unstructured (independent) across conditions.

We used these simulations to compare mash with several other methods: mash-bmalite, a simplified version of mash that uses only the BMALite models from Flutre *et al.*⁶; ash¹⁶, a univariate analogue of mash that analyzes each condition separately; and metasoft^{9,14}, which implements several multivariate tests corresponding to different models (the “fixed effects”, or FE⁹, model, which assumes equal effects in all conditions, the RE2 model⁹, which allows for normally-distributed variation in effects among conditions, and the BE model, which extends the RE2 model to allow for effects are exactly zero¹⁴).

The simulation results are summarized in Fig. 2 and Supplementary Table 1. We assessed the methods on three important tasks: distinguishing non-null units from null units (Fig. 2a–b); distinguishing non-null effects (*i.e.*, non-null unit-condition pairs) from null effects (Fig. 2c–d); and estimating effect sizes (Fig. 2e–f). (We included metasoft results for the first task only because it does not provide multivariate effect estimates, and because it is unclear how best to combine unit-level significance (*p* value) and condition-specific significance (*m* value) into a significance measure for each unit-condition pair.)

In the “shared, structured” scenario, mash markedly outperformed the other methods in all three tasks (Fig. 2, Supplementary Table 1). This is expected because mash is the only method that attempts to fully exploit the structure in these effects. Further, mash found essentially all (>99%) of the signals that the other methods found, plus additional signals (Supplementary Table 2).

In the “shared, unstructured” scenario, mash and mash-bmalite are the best methods, finding almost exactly the same signals (Supplementary Table 2), with metasoft-RE2 close behind as it was applied to the task it was designed for (testing unit-level significance). The strong similarity in performance between mash and mash-bmalite is explained by mash putting negligible weight on the data-driven covariance matrices, which are unnecessary in this scenario. This illustrates mash’s ability to adapt to the data, increasing power in complex scenarios without being unduly penalized in simpler scenarios.

Note that mash uses two distinct strategies to improve accuracy of effect estimates: (i) it shrinks estimates towards zero, which improves average accuracy because most effects are null; (ii) in the presence of “structured effects”, it shares information across conditions to improve accuracy. For example, if a unit has effects that are similar across a subset of conditions, averaging the effect estimates in those conditions will improve accuracy. Both these strategies contribute to the strong performance of mash in the “shared, structured” scenario.

These simulations are meant to be illustrative rather than comprehensive. Indeed, although numerical comparisons can help choose between methods, we believe that *qualitative* factors are equally important. For example, the fact that mash produces effect size estimates is a key feature that distinguishes it from most existing methods. See Supplementary Note for a detailed discussion on connections and differences between the methods.

GTEx *cis*-eQTL analysis.

To illustrate mash in a substantive application, we analyzed eQTLs across 44 human tissues, using data from the Genotype Tissue Expression (GTEx) project¹⁷. The GTEx project aims to provide insights into the mechanisms of gene regulation by studying gene expression and regulation in multiple tissues from human donors. One fundamental question is which SNPs are eQTLs (*i.e.*, associated with gene expression) in which tissues. Answering this could help distinguish regulatory mechanisms that are tissue-specific or shared, and, by integrating with results from genome-wide association studies (GWAS), could help identify the most relevant tissues for complex disease^{17,18}.

As input, mash requires results from a condition-by-condition analysis (effect estimates and standard errors). We used the effect estimates and standard errors for candidate local (“*cis*”) eQTLs for each gene (GTEx data release 6). These were obtained by performing (univariate) single-SNP eQTL analyses in each tissue using Matrix eQTL¹⁹. Expression levels were corrected for population structure (using genotype principal components²⁰) and other confounding factors (both measured factors such as age and sex, and unmeasured factors estimated using factor analysis²¹), then rank-transformed to the corresponding quantiles of a standard normal distribution. Because these estimates were obtained by single-SNP analysis, the estimated effects for each SNP reflect both the effects of the SNP itself *and SNPs in LD with it*. Thus, our analyses do not distinguish causal eQTLs from those that are in LD with the causal eQTLs; see Discussion and Supplementary Note.

We analyzed the 16,069 genes for which effect estimates were available for all 44 tissues considered; the filtering criteria used by GTEx ensure that these genes show some expression in all 44 tissues.

mash improves model fit.

To assess the improved fit of mash compared with the simpler mash-bmalite, we used cross-validation; we fitted each model to a random subset of units (“training set”), and assessed fit by the log-likelihood on the remaining units (“test set”). We found that mash improved the test set log-likelihood very substantially (by 23,796; Supplementary Fig. 1). Further, mash placed 79% of the weight on the data-driven covariance matrices. These results confirm that our methods for estimating data-driven covariance matrices are sufficiently effective that they better capture most effects than do the canonical matrices used by existing methods.

Identification of data-driven patterns of sharing.

The data-driven covariance matrix with largest mash weight is shown in Fig. 3 (34% weight). This covariance matrix captures several important features: (i) effects are positively correlated among all tissues; (ii) the brain tissues—and, to a lesser extent, testis and pituitary—are particularly strongly correlated with one another, and less correlated with other tissues; (iii) effects in whole blood tend to be less correlated with other tissues. Many other data-driven covariance matrices estimated by mash also have positive correlations among all tissues and/or highlight heterogeneity between brain tissues and other tissues, confirming these as common features of these data (Supplementary Fig. 2). Other components capture

less prevalent patterns, such as effects that are appreciably stronger in one tissue (e.g., Supplementary Fig. 2b).

Patterns of sharing inform effect size estimates.

Having estimated patterns of sharing, mash exploits these patterns to improve effect estimates at each candidate eQTL. Although we cannot directly demonstrate improved accuracy of effect estimates in the GTEx data (for this, see simulations above), individual examples can provide insight into how mash achieves improved accuracy. Figure 4 shows three illustrative examples, which we discuss in turn.

In the first example (Fig. 4, left-hand column), raw effect estimates in each tissue are mostly positive, and strongest in brain tissues. The mash estimates are *all* positive: the few modest negative estimates are outweighed by the strong background information that effects are highly correlated among tissues and mostly positive. This weighing of evidence was done using Bayes' rule. Humans are notoriously bad at weighing background information against specific instances—they tend to underweight background information when presented with specific data²²—so this behavior may or may not be intuitive. The mash effect estimates are also appreciably larger in brain tissues than in other tissues. This is a result of using Bayes' rule to combine the effect estimates at *this* eQTL with the background information on heterogeneity of brain and non-brain effects learned from *all* eQTLs.

In the second example (Fig. 4, middle column), most effect estimates in non-brain tissues are positive (30 out of 34), but modest in size, and only one is nominally significant ($p < 0.05$). However, by combining information among tissues, mash estimated that all effects in non-brain tissues are positive, and mostly “significant” ($lfr < 0.05$). By contrast, the estimated effects in brain tissues are inconsistent (both positive and negative) and so mash is not confident about the sign of effects in brain tissues. This example illustrates that mash can learn to treat subsets of conditions differently; mash learned that effects in brain tissues are occasionally different from effects in other tissues, and therefore did not draw strong inferences in the brain based on the other tissues.

In the third example (Fig. 4, right-hand column), the effect estimates vary in sign, and are modest except for a very strong signal in whole blood. While whole-blood-specific effects were estimated to be rare, mash—again, using Bayes' theorem—recognized that the strong signal at this eQTL outweighs the background information, and estimated a strong effect in blood with insignificant effects in other tissues. This illustrates that mash, although focused on combining information across tissues, can still recognize—and shed light on—tissue-specific patterns when they exist.

In all three examples, the mash estimates are noticeably different from those obtained by mash-bmalite (Fig. 4c) and ash (Supplementary Fig. 3). The decreased flexibility of mash-bmalite compared to mash is particularly evident in the first example where mash-bmalite estimates are almost constant across tissues, hiding the stronger effects in brain-related tissues that are clearly visible in the original data, and maintained by mash.

Increased identification of significant effects.

Consistent with our simulation results, mash identified many more significant effects than either mash-bmalite or ash. To avoid double-counting eQTLs in the same gene that are in LD with one another, we assessed the significance of only the “top” SNP in each gene (the SNP with the largest univariate $|Z|$ -statistic across tissues). Thus, our results are based on 16,069 candidate eQTLs, each with effect estimates in 44 tissues, for a total of $44 \times 16,069 = 707,036$ effects.

The vast majority of top SNPs show a strong signal in at least one tissue (97% have a maximum $|Z|$ score exceeding 4), consistent with most genes containing at least one eQTL in at least one tissue. However, in the tissue-by-tissue analysis ash identified only 13% of these effects as “significant” at $lfr < 0.05$; that is, the univariate analysis was highly confident in the direction of the effect in only 13% of cases. By comparison, mash-bmalite identified 39% as significant at the same threshold, and mash identified 47%. Similar to our simulations, the significant associations identified by mash include the vast majority (95%) of those found by either of the other methods (Supplementary Table 4).

Overall, mash found 76% (12,171 out of 16,069) of the top SNPs to be significant in at least one tissue. We refer to these as the “top eQTLs” in subsequent sections.

Sharing of effects among tissues.

To investigate sharing of the top eQTLs among tissues, we assessed sharing of effects by sign and by magnitude (effects have the same sign and are within a factor of 2 in size). Because of the large differences between brain and non-brain tissues, we also show results separately for these subsets. The results (Table 1, Fig. 5) confirm extensive eQTL sharing among tissues, particularly among brain tissues. Sharing by sign always exceeds 85%, and is as high as 96% among brain tissues. (Furthermore, these numbers may underestimate the sharing by sign of actual causal effects due to the impact of multiple eQTLs in LD; see Supplementary Note and Supplementary Fig. 4.) Sharing by magnitude is necessarily lower because sharing by magnitude implies sharing by sign. On average, 36% of tissues show an effect within a factor of 2 of the strongest effect at each top eQTL. However, within brain tissues this increases to 76%. Thus, not only do eQTLs tend to be shared among brain tissues, but effect sizes tend to be homogeneous. Because our analyses were based on the top eQTLs, our results reflect patterns of sharing only among stronger *cis* eQTLs; weaker eQTLs may show different patterns.

Some tissues share eQTLs more than others. Fig. 6 summarizes eQTL sharing by magnitude between all pairs of tissues (see Supplementary Fig. 5 for sharing by sign). In addition to strong sharing among brain tissues, mash identified increased sharing among other biologically related groups, including arteries (tibial, coronary and aortal), two groups of gut tissues (one group containing esophagus and sigmoid colon, the other containing stomach, terminal ileum of the small intestine and transverse colon), skin (sun-exposed and non-exposed), adipose (subcutaneous and visceral-omentum) and heart (left ventricle and atrial appendage). Figure 6 also reveals heterogeneity in effect sizes in cerebellum versus non-cerebellum tissues, and highlights sharing between pituitary and brain tissues.

Different levels of effect sharing among tissues means that effect estimates in some tissues gain more precision than others from the joint analysis. To quantify this, we computed an “effective sample size” (ESS) for each tissue that reflects typical precision of its effect estimates (Supplementary Fig. 6, Supplementary Note). The ESS values are smallest for tissues that show more “tissue-specific” behaviour (e.g., testis, whole blood), and are largest for coronary artery, reflecting its stronger correlation with other tissues.

Tissue-specific eQTLs.

Despite high average levels of sharing of eQTLs among tissues, mash identified some eQTLs that are relatively “tissue-specific”. Indeed, the distribution of the number of tissues in which an eQTL is shared by magnitude has a mode at 1 (Fig. 5), representing a subset of eQTLs that have a much stronger effect in one tissue than in any other (henceforth, “tissue-specific” for brevity). Breaking down this group by tissue identified testis as the tissue with the most tissue-specific effects (Supplementary Fig. 7). Testis and whole blood stand out as having lower pairwise sharing of eQTLs with other tissues (Fig. 6). Other tissues showing larger than average tissue specificity (Fig. 5, Supplementary Fig. 7) include skeletal muscle, thyroid and transformed cell lines (fibroblasts and LCLs).

One possible explanation for tissue-specific eQTLs is tissue-specific expression; that is, if a gene is strongly expressed only in one tissue, this could cause an eQTL for that gene to show a strong effect only in that tissue. Whether or not a tissue-specific eQTL is due to tissue-specific expression could considerably impact biological interpretation. To assess whether tissue-specific eQTLs could be explained by tissue-specific expression we took genes with tissue-specific eQTLs and examined the distribution of expression in the eQTL-affected tissue relative to expression in other tissues. We found this distribution to be similar to genes without tissue-specific eQTLs (Supplementary Fig. 8). Therefore, most tissue-specific eQTLs identified here do not solely reflect tissue-specific expression.

Discussion

The statistical benefits of joint multivariate analyses compared with univariate analyses are well documented, and increasingly widely appreciated. But we believe this potential nonetheless remains under-exploited. With mash we aim to provide a set of flexible and general tools to facilitate such analyses. In particular, mash is *generic* and *adaptive*. It is generic in that it can take as input any matrix of Z scores (or, preferably, a matrix of effect estimates and a matrix of corresponding standard errors) that test many effects in many conditions. These scores could come from many sources; e.g., from linear regression, generalized linear models or linear mixed models¹³. And mash is adaptive in that it learns correlations among effects from the data, allowing it to maximize power and precision for each setting. Consequently, mash should be widely applicable to many settings involving estimation of multivariate effects.

The mash method uses EB hierarchical modelling, and so is related to other EB methods^{5,6,12}. Indeed, the mash framework essentially includes these methods as special cases (as well as simpler methods such as “fixed effects” and “random effects” meta-analyses^{9,23}). One key feature that distinguishes mash from previous methods is that mash

puts greater emphasis on *quantitative* estimation and assessment of effects. Moving away from binary-based models has at least two advantages. First, allowing for all possible binary configurations can create computational challenges. Second, in practice we have found that effects are often shared broadly among many conditions, and in such cases binary-based methods tend to infer that effects are non-zero in most or all conditions, even when the signal is modest in some conditions. This conclusion may be technically correct—for example, in our GTEx analysis it is not impossible that all eQTLs are somewhat active in all tissues. However, as our analysis has illustrated, a more quantitative focus can reveal variation in effect sizes that may be of considerable biological importance.

One limitation of our eQTL analysis is that, like most eQTL analyses, it does not distinguish between causal associations and those that are due to LD. This issue is particularly important to appreciate when, for example, cross-referencing GWAS associations with eQTL effect estimates; a GWAS-associated SNP may be a “significant” eQTL simply because it is in LD with another causal SNP. For single-tissue eQTL mapping, this problem has been addressed in several ways, including fine-mapping^{24–29} and co-localization^{30–32}. For multi-tissue analysis, only more limited attempts exist to address this problem. For example, eQTLBMA⁵ implements a Bayesian approach to fine-mapping under the simplifying assumption that there is at most one causal SNP per gene^{24,25}. However, this assumption becomes less plausible in analyses of many tissues, and developing more flexible multi-tissue fine-mapping methods seems an important future direction.

Dealing with multiple tests is often described as a “burden”. This likely originates from the fact that controlling family-wise error rate requires more stringent thresholds as the number of tests increases. However, modern analyses prefer to control the FDR³³, which does not depend on the number of tests³⁴. Consequently, the term “burden” is inaccurate and unhelpful. We believe that the results of many tests in many conditions should instead be viewed as an *opportunity*—an opportunity to learn about relationships among underlying effects, and make data-driven decisions to improve inferences. This will inevitably, it seems, involve modelling assumptions, and the challenge is to design flexible models that work well in a range of settings. The methods presented here represent a major step towards this goal.

URLs.

Multivariate adaptive shrinkage (“mash”) software, <https://github.com/stephenslab/mashr>; code and data resources for GTEx analysis, <https://github.com/stephenslab/gtexresults>, <https://doi.org/10.5281/zenodo.1296399>; GTEx project, <http://gtexportal.org>; adaptive shrinkage (“ash”) software, <https://github.com/stephens999/ashr>; Sparse Factor Analysis (SFA) software, <http://stephenslab.uchicago.edu/software.html>; Extreme Deconvolution software, <https://github.com/jobovy/extreme-deconvolution>; METASOFT software, <http://genetics.cs.ucla.edu/meta>; Matrix eQTL software, http://www.bios.unc.edu/research/genomic_software/Matrix_eQTL.

Online Methods

Model and fitting.

Let b_{jr} ($j = 1, \dots, J$; $r = 1, \dots, R$) denote the true value of effect j in condition r . Further, let \hat{b}_{jr} denote the “observed” estimate of this effect, and let \hat{s}_{jr} be the standard error of this estimate, so $z_{jr} = \hat{b}_{jr}/\hat{s}_{jr}$ is the standard Z statistic used to test whether b_{jr} is zero. Let B , \hat{B} , S and Z denote the corresponding $J \times R$ matrices, and let b_j (respectively, \hat{b}_j , z_j) denote the j th row of B (respectively, \hat{B} , Z).

We assume \hat{b}_j is normally distributed about b_j with variance-covariance matrix V_j (defined below), and that b_j follows eq. 1:

$$p(\hat{b}_j \mid b_j, V_j) = N_R(\hat{b}_j; b_j, V_j) \quad (2)$$

$$p(b_j \mid \pi, U) = \sum_{k=1}^K \sum_{l=1}^L \pi_{k,l} N_R(b_j; 0, \omega_l U_k), \quad (3)$$

where $N_R(\cdot; \mu, \Sigma)$ denotes the multivariate normal (MVN) density in R dimensions with mean μ and variance-covariance matrix Σ . Here, each U_k is a covariance matrix that captures a pattern of effects; each ω_l is a scaling coefficient that corresponds to a different effect size; and the mixture proportions $\pi_{k,l}$ determine the relative frequency of each covariance-scale combination. The scaling coefficients ω_l take values on a fixed dense grid that spans “very small” to “very large” so as to capture the full range of effects that could occur.

Combining eqs. 2 and 3 implies that the marginal distribution of \hat{b}_j , integrating out b_j is

$$p(\hat{b}_j \mid \pi, U, V_j) = \sum_{k=1}^K \sum_{l=1}^L \pi_{k,l} N_R(\hat{b}_j; 0, \omega_l U_k + V_j). \quad (4)$$

Each covariance matrix V_j is specified as $V_j = S_j C S_j$, where C is a correlation matrix that accounts for correlations among the measurements in the R conditions, and S_j is the $R \times R$ diagonal matrix with diagonal elements $\hat{s}_{j1}, \dots, \hat{s}_{jR}$. If measurements in the R conditions are independent, one would set $C = I_R$, the $R \times R$ identity matrix, and therefore $V_j = S_j^2$.

However, in our GTEx analysis the measurements are correlated due to sample overlap

(individuals in common) among tissues; we estimate this correlation from the data (see “Estimating the correlation matrix C ” below). The methods place no restriction on the V_j except that they must be symmetric positive definite matrices.

The steps of mash are:

1. Estimate U and π . This involves two sub-steps:
 - a. Define a set of (normalized) data-driven and canonical covariance matrices, \hat{U} .
 - b. Given \hat{U} , estimate π by maximum likelihood. (A key idea is that if some matrices generated in Step 1a do not help capture patterns in the data, they will receive little weight.) Let $\hat{\pi}$ denote this estimate.
2. Compute, for each j , the posterior distribution $p(b_j \mid \hat{b}_j, \hat{\pi}, \hat{U}, V_j)$

These steps are now detailed in turn.

Generate data-driven covariance matrices U_k .

We first identify rows j of matrix \hat{B} that correspond to the “strongest” effects. For example, in the GTEx data we chose rows corresponding to the “top” SNP for each gene, which we defined to be the SNP with the highest value of $Z_j^{\max} = \max_r \hat{b}_{jr} / \hat{s}_{jr}$. (We used the maximum rather than the sum because we wanted to include effects that were strong in a single condition, not just effects that were shared among conditions.) For the simulated data, we ran ash separately for each condition r , computed $lfsr_{jr}$ for each (j, r) , then chose rows j for which $\min_r lfsr_{jr} < 0.05$.

Next, we fit a mixture of MVN distributions to these strongest effects using methods from Bovy *et al.*³⁵. Bovy *et al.* describe an EM algorithm for fitting a model similar to eqs. 2 and 3, with the crucial difference that there are no scaling parameters on the covariances:

$$p(b_j \mid \pi, U) = \sum_{k=1}^K \pi_k N_R(b_j; 0, U_k). \quad (5)$$

The absence of the scaling factors ω_l means that, compared with mash, their model is less well suited to capturing effects that have similar patterns (relative sizes across conditions) but vary in magnitude. By applying it to only the largest effects, we seek to sidestep this issue.

Estimates of U_k from this EM algorithm are sensitive to initialization. Furthermore, we noticed an interesting feature of the EM algorithm: each iteration preserves the rank of the matrices U_k , so that ranks of the estimated matrices are the same as the ranks of the matrices used to initialize the algorithm. We exploited this fact to ensure that some of the estimated

U_k had low rank. This helped stabilize the estimates since rank-penalization is one way to regularize covariance matrix estimation.

To describe the initialization in detail, let \tilde{J} denote the number of “strongest effects” selected above, and let \tilde{Z} denote the column-centered $\tilde{J} \times R$ matrix of Z scores for these “strong effects”. To extract the main patterns in \tilde{Z} , we perform dimension reduction on \tilde{Z} ; specifically, we apply principal components analysis (through a singular value decomposition, or SVD) and sparse factor analysis²¹ (SFA) to \tilde{Z} . SVD yields a set of singular values and singular vectors of \tilde{Z} . Let λ_p , v_p denote the p th singular value and corresponding right singular vector. SFA yields matrix factorization

$$\tilde{Z} = LF + E,$$

where L is a sparse $J \times Q$ matrix of loadings and F is a $Q \times R$ matrix of factors. We use $Q = 5$.

Given this matrix factorization, we use the EM algorithm to fit the mixture model (eq. 5) with $K = 3$, in which the three matrices U_k are initialized as follows:

- $U_1 = \tilde{Z}^T \tilde{Z} / \tilde{J}$, the empirical covariance matrix of \tilde{Z} .
- $U_2 = \sum_{p=1}^P \lambda_p v_p v_p^T / \tilde{J}$, which is a rank- P approximation of the covariance matrix of \tilde{Z} , with $P < Q$. We use $P = 3$.
- $U_3 = F^T L^T L F / \tilde{J}$, which is a rank- Q approximation of the covariance matrix of \tilde{Z} .

The output of the EM algorithm defines U_1 , U_2 and U_3 in the mash model (eq. 1).

In addition to the covariance matrices obtained from this EM algorithm, we define more covariance matrices from the SFA results; specifically, the $Q = 5$ rank-1 matrices $F_q^T L_q^T L_q F_q$, with $q = 1, \dots, Q$. The rationale for adding these rank-1 covariance matrices is that the SFA factors may directly reflect effect patterns in the data, and if so these matrices will be a helpful addition. (If they are not helpful, they will receive little weight when estimating π).

In summary, SFA and the EM algorithm generate 8 data-driven covariance matrices for our analysis of the GTEx and simulated data sets.

Generate canonical covariance matrices U_k .

We use the following “canonical” covariance matrices:

- The identity matrix, I_R . This represents the situation where the effects in different conditions are independent, which may be unlikely in some applications (like the GTEx application here), but seems useful to include if only to show that the data assign low weight to it.

- R rank-1 matrices $e_r e_r^T$, where e_r denotes the unit vector with zeros everywhere except for element r , which is set to 1. These matrices represent effects that occur in only one condition.
- The rank-1 matrix $\mathbf{1}\mathbf{1}^T$, where $\mathbf{1}$ denotes the R -vector of ones—this is the matrix of all ones, and represents effects that are identical among all conditions.

The user can, if desired, add other canonical matrices. For example, if R is moderate then one could add the 2^{R-1} canonical matrices that correspond to shared (equal) effects in each of the 2^{R-1} subsets.

In summary, this procedure produces 46 canonical covariance matrices for our analyses of the GTEx and simulated data sets, with $R = 44$.

Standardize the covariance matrices.

Since eq. 3 uses the same grid of scaling factors for each covariance matrix U_k , we standardize the matrices U_k so that they are similar in scale. Specifically, for each k , we divide all elements U_k by the maximum diagonal element of U_k so that the maximum diagonal element of the rescaled matrix is 1. These rescaled matrices define \hat{U} in Step 1a of mash.

Define grid of scaling factors.

We define a dense grid of scaling factors, ranging from “very small” to “very large”. Stephens¹⁶ suggests a way to select suitable limits ($\omega_{\min}, \omega_{\max}$) for this grid in the univariate case; we apply this method to each condition r , and take the smallest ω_{\min} and the largest ω_{\max} as the grid limits. The internal points of the grid are obtained, as in the univariate case¹⁶, by setting $\omega_l = \omega_{\max}/m^{l-1}$, for $l = 1, \dots, L$, where $m > 1$ is a user-tunable parameter that controls the grid density, and L is chosen to be just large enough so that $\omega_L < \omega_{\min}$. We use $m = \sqrt{2}$.

Estimate π by maximum likelihood.

Given \hat{U} and ω , we estimate the mixture proportions π by maximum likelihood.

To simplify notation, let $\Sigma_{k,l} = \omega_l U_k$, and replace the double index k, l with a single index p ranging from 1 to $P := KL$, so that eqs. 3 & 4 become

$$p(b_j \mid \pi, U) = \sum_{p=1}^P \pi_p N(b_j; 0, \Sigma_p).$$

$$p(\hat{b}_j \mid \pi, U, V_j) = \sum_{p=1}^P \pi_p N(\hat{b}_j; 0, \Sigma_p + V_j).$$

Assuming independence of the rows of \hat{B} , the likelihood for π is

$$L(\pi) = p(\hat{B} \mid \pi, U, V) = \prod_{j=1}^J p(\hat{b}_j \mid \pi, U, V_j) = \prod_{j=1}^J \sum_{p=1}^P \pi_p N_R(\hat{b}_j; 0, \Sigma_p + V_j)$$

If the rows of \hat{B} are not independent, this may be interpreted as a “composite likelihood”³⁶, which generally yields consistent point estimates. (We also empirically confirmed that our GTEx results were robust to LD pruning of SNPs; see Supplementary Fig. 9.)

Maximizing $L(\pi)$ is a convex optimization problem, which we solve using EM³⁷, accelerated using SQUAREM³⁸. If \hat{B} has a large number of rows, we can reduce computational effort by taking a random subset of rows. In the GTEx analysis, we use a random subset of 20,000 rows. (It is important that this is a random subset, and not just the \tilde{J} rows of strong effects.)

Posterior calculations.

If $b \sim N_R(0, U)$ and $\hat{b} \mid b \sim N_R(b, V)$, then

$$\hat{b} \mid b \sim N_R(\tilde{\mu}, \tilde{U}),$$

where

$$\tilde{U} = \tilde{U}(U, V) = (U^{-1} + V^{-1})^{-1} \quad (6)$$

$$\tilde{\mu} = \tilde{\mu}(U, V, \hat{b}) = \tilde{U}(U, V)V^{-1}\hat{b} \quad (7)$$

This result is easily extended to the case where the prior on b is a mixture of MVNs (eq. 3). In this case the posterior distribution is

$$p(b_j \mid \hat{b}_j, V_j, \pi) = \sum_{p=1}^P \tilde{\pi}_{jp} N_R(b_j; \tilde{\mu}_{jp}, \tilde{U}_{jp}),$$

where $\tilde{U}_{jp} = \tilde{U}(\Sigma_p, V_j)$ (eq. 6) and $\tilde{\mu}_{jp} = \tilde{\mu}(\Sigma_p, V_j, \hat{b}_j)$ (eq. 7), and

$$\tilde{\pi}_{jp} = \frac{\pi_p N_R(\hat{b}_j; 0, \Sigma_p + V_j)}{\sum_{p'=1}^P \pi_{p'} N_R(\hat{b}_j; 0, \Sigma_{p'} + V_j)}.$$

The posterior mean and variance are

$$\hat{b}_j = E(b_j \mid \hat{b}_j, V_j, \pi) = \sum_{p=1}^P \tilde{\pi}_{jp} \tilde{\mu}_{jp} \quad (8)$$

$$\text{Var}(b_{jr} \mid \hat{b}_j, V_j, \pi) = \sum_{p=1}^P \tilde{\pi}_{jp} (\tilde{U}_{jp,rr} + \tilde{\mu}_{jp,r}^2) - \left(\sum_{p'=1}^P \tilde{\pi}_{jp'} \tilde{\mu}_{jp',r} \right)^2.$$

The local false sign rate (*lfsr*) is defined by Stephens¹⁶ as:

$$lfsr_{jr} = \min\{\Pr(b_{jr} \geq 0 \mid D), \Pr(b_{jr} \leq 0 \mid D)\},$$

where D denotes the available data. Intuitively, $lfsr_{jr}$ is the probability that we would incorrectly predict the sign of the effect if we were to use our best guess of the sign (positive or negative). Thus, a small *lfsr* indicates high confidence in the sign of an effect. The *lfsr* is more conservative than the local false discovery rate (*lfdr*)³⁹ because requiring high confidence in the sign is more stringent than requiring high confidence that the effect be non-zero. More importantly, the *lfsr* is more robust to modelling assumptions than the *lfdr*¹⁶, a particularly important issue in multivariate analyses where modelling assumptions play a larger role.

Bayes Factors testing global null.

The Bayes Factor for $\mathbf{b} \neq 0$ against the null model ($\mathbf{b} = 0$) is

$$BF_j = p(\hat{b}_j \mid \pi, U, V_j) / p(\hat{b}_j \mid \mathbf{b} = 0, V_j).$$

The numerator is given by eq. 4, and the denominator is given by eq. 2, with $\mathbf{b} = 0$.

The EZ model, and applying mash to Z scores.

The model (eq. 3) assumes the b_j are independent of their standard errors, V_j . We refer to this as the “Exchangeable

Effects” (EE) model²³. An alternative is to allow the effects to scale with standard error, so that effects with larger standard error tend to be larger,

$$S_j^{-1} b_j \mid \pi, U, V_j \sim g,$$

where g is shorthand for the mixture of MVNs (eq. 3). We refer to this as the “Exchangeable Z” (EZ) model because the left of this equation is the vector of Z scores for effect j .

As explained in Stephens¹⁶, the EZ model can be fit by applying the same code as the EE model to the Z statistics instead of the effect estimates, \hat{B} , in which the standard errors of the Z statistics are set to 1; in other words, set $\hat{b}_{jr} = z_{jr}$ and $\hat{s}_{jr} = 1$. (The correlation matrix, C , remains the same.) One advantage of the EZ model is that it can be fit using only the Z scores; it does not require access to both the effect estimates and their standard errors. The lfr can also be computed using only the Z scores. However, the posterior mean estimates that arise from this model are estimates of $S_j^{-1}b_j$; transforming these to estimates of effect sizes b_j requires knowledge of S_j .

We analyzed the GTEx data using both EE and EZ models. Results were qualitatively similar in terms of patterns of sharing, but the EZ model performed better in cross-validation tests of model fit (see below), so we used results from the EZ model.

Estimating the correlation matrix C .

C is an $R \times R$ correlation matrix that accounts for correlations among measurements in the R conditions. To estimate this matrix, we exploit the fact that C is the correlation matrix of the Z scores under the null ($b_j = 0$). Specifically, we estimate C from the empirical correlation matrix of the Z scores for the effects j that are most consistent with the null,

$$C = \frac{1}{|\mathcal{N}|} \sum_{j \in \mathcal{N}} z_j z_j^T,$$

where $\mathcal{N} = \left\{ j : \max_r |z_{jr}| < 2 \right\}$. For the GTEx data, we found that measurements in different tissues were not highly correlated; all elements of estimated C had magnitude less than 0.2, and 95% had magnitude less than 0.1. However, in cross-validation tests (below), this estimated C produced better model fit than ignoring correlations (*i.e.*, $C = I_R$).

Cross-validation of model fit.

To compare model fitting procedures, we randomly divide the rows of the data matrix \hat{B} into two subsets—a “training set” \hat{B}_{train} , and a “test set” \hat{B}_{test} . We then apply mash to \hat{B}_{train} , yielding estimates $(\hat{\pi}, \hat{U})$, and assess the “fit” of these estimates by computing the log-likelihood in the test data, $\log p(\hat{B}_{test} \mid \hat{\pi}, \hat{U}, V)$, which is given by eq. 4. This cross-validation strategy can also be used to compare different approaches to estimating \hat{U} , and our current strategy was developed and refined using this framework.

When applying cross-validation to the GTEx data, we create test and training sets by randomly selecting half the genes, rather than half the rows (gene-SNP pairs). Specifically, we use genes on even-numbered chromosomes as the training set, and genes on odd-numbered chromosomes as the test set. This ensures that rows in the test set were independent of rows in the training set.

Visualizing U_k .

In the simulations, and in our application to the GTEx data, $R = 44$, so each U_k is a 44×44 covariance matrix, and each component of the mixture (eq. 1) is a distribution in 44 dimensions. Visualizing such a distribution is challenging. However, we can get some insight from the first eigenvector of U_k , denoted by v_k , which captures the principal direction of the effects. If U_k is dominated by this principal direction, then we can think of effects from that component as being of the form λv_k for some scalar λ . For example, if the elements of the vector v_k are approximately equal, then component k captures effects that are approximately equal in all conditions; alternatively, if v_k has one large element, with other elements close to zero, then component k corresponds to an effect that is strong in only one condition. See Fig. 3 and Supplementary Fig. 2 for illustration.

RRMSE (accuracy of estimates in simulation studies).

In simulations, we used the “relative root mean squared error” (RRMSE) to assess error in effect estimates (e.g., Fig. 2). Let b_{jr}^* denote the estimate of effect j in condition r (for mash,

$b_{jr}^* = \hat{b}_{jr}$). Then

$$\text{RRMSE} = \frac{\sqrt{\frac{1}{JR} \sum_{j=1}^J \sum_{r=1}^R (b_{jr} - b_{jr}^*)^2}}{\sqrt{\frac{1}{JR} \sum_{j=1}^J \sum_{r=1}^R (b_{jr} - \hat{b}_{jr})^2}}.$$

ROC curves.

For the ROC curves in Fig. 2c–d, the True Positive Rate (TPR) and False Positive Rate (FPR) are computed at any given threshold t as

$$\text{TPR} = \frac{|CS \cap S|}{|T|}$$

$$\text{FPR} = \frac{|N \cap S|}{|M|},$$

where S is the set of significant results at threshold t , CS the set of correctly signed results, T the set of true (non-zero) effects and N the set of null effects:

$$S = \{j, r : |f_{sr}| > t\}$$

$$CS = \{j, r : E(b_{jr} | D) \times b_{jr} > 0\}$$

$$N = \{j, r : b_{jr} = 0\}$$

$$T = \{j, r : b_{jr} \neq 0\}.$$

Thus, to be considered a true positive, we require that an effect be both “significant” and correctly signed.

For the ROC curves in Fig. 2a–b, the TPR and FPR are computed by treating units j as “discoveries”. For example, suppose a method produces a p value p_j for testing unit j . Then, at any threshold t , the TPR and FPR are

$$TPR = |\tilde{S}|/|\tilde{T}|$$

$$FPR = |\tilde{S}|/|\tilde{N}|,$$

where \tilde{S} is the set of significant units at threshold t , \tilde{T} the set of true (non-null) units, and \tilde{N} the set of null units:

$$\tilde{S} = \{j : p_j \leq t\}$$

$$\tilde{N} = \{j : b_{jr} = 0, \forall r\}$$

$$\tilde{T} = \{j : \exists r \text{ s.t. } b_{jr} \neq 0\}.$$

Effective sample size.

We define the “effective sample size” (ESS) for tissue r to be

$$n_r^{eff} = n_r^{orig} \times \text{median}_j \left(\hat{s}_{jr}^2 / \hat{s}_{jr}^2 \right)$$

where n_r^{orig} is the number of available samples in tissue r (Supplementary Table 3), and \hat{s}_{jr} and \tilde{s}_{jr} are the standard error and posterior standard deviation, respectively, for effect j in tissue r .

Normalized effects.

For estimating sharing of effects by magnitude (e.g., Supplementary Fig. 7), we define a “normalized” effect \tilde{b} in each condition as the ratio of its effect in that condition to the largest effect across all conditions:

$$\tilde{b}_{jr} = b_{jr} / b_{jr_0},$$

Where

$$r_0 := \operatorname{argmax}_r |b_{jr}|$$

For example, in the application to multi-tissue eQTL analysis, a normalized effect $\tilde{b}_{jr} = 0.5$ means that the effect of eQTL j in tissue r is half that of its effect in the strongest tissue.

Sharing by sign and magnitude.

We define two effects to be “shared in sign” if they have the same sign; we define them to be “shared in magnitude” if they *also* have an effect within a factor of 2 of one another. (In some applications it may make sense to drop the requirement that effects are of the same sign when assessing sharing by magnitude, but not here.)

In Fig. 5 and Table 1, we assess sharing for each eQTL by comparing its estimated effect in each tissue against the strongest estimated effect across all tissues. Thus, for example, if an eQTL has a very strong positive effect in one tissue and weak negative effects in all other tissues, then the number of tissues “shared by sign” for this eQTL would be 1.

In Fig. 6 (and Supplementary Fig. 5), we investigate sharing between each pair of tissues. When computing sharing between tissues r and s , we consider only QTLs that are “significant” ($lfsr < 0.05$) in at least one of r and s . This is to avoid sharing estimates being driven by estimates of null or nearly-null effects.

More formally, define

$$QTL_r := \left\{ j : lfsr_{jr} < 0.05 \right\}$$

$$SS_{rs} := \left\{ j : \operatorname{sign}(\hat{b}_{jr}) = \operatorname{sign}(\hat{b}_{js}) \right\}$$

$$SM_{rs} := \left\{ j : \frac{1}{2} \leq \hat{b}_{jr} / \hat{b}_{js} \leq 2 \right\}.$$

Then the sharing by sign between r and s is given by

$$\frac{|SS_{rs} \cap (QTL_r \cup QTL_s)|}{|QTL_r \cup QTL_s|},$$

and sharing by magnitude between r and s is given by

$$\frac{|SM_{rs} \cap (QTL_r \cup QTL_s)|}{|QTL_r \cup QTL_s|}.$$

These estimates of sharing are based on the posterior estimates \hat{b}_j to simplify their calculation. However, we obtained similar estimates of sharing when we accounted for uncertainty in \hat{b}_j by sampling from the posterior distributions of the effects.

ash:

We applied ash¹⁶ separately to each tissue using the same 20,000 randomly-selected gene-SNP pairs as for mash. We then computed the posterior means and *Ifsr* for the top SNPs.

mash-bmalite:

This is a simplification of mash without the data-driven covariance matrices, and with particular choices for the canonical covariance matrices and scaling factors. It can be thought of as a version of BMALite⁵ that outputs effect size estimates and *Ifsr* values. (The original software does not compute these quantities.)

Specifically, the U for mash-bmalite includes the 44 singleton configurations ($U_k = e_k e_k^T$), and matrices corresponding to the models in Flutre *et al.*⁵, with heterogeneity parameters $H = \{0, 0.25, 0.5, 1\}$. (When $H=0$, effects are equal in all conditions; when $H=1$, effects are independent among conditions.) We used $\omega = \{0.1, 0.4, 1.6, 6.4, 25.6\}$, similar to the coarse grid used in Flutre *et al.*⁵, and chosen to capture the full range of the GTEx (real and simulated) Z statistics.

Statistical tests.

For the ash, mash and mash-bmalite methods, an association was considered significant if the *Ifsr*¹⁶ was less than 0.05. For all metasoft approaches included in the simulation experiments (FE, RE2, BE), p values were computed as described in ^{9,14}. The metasoft p value threshold was varied to assess performance in the simulation experiments (Fig. 2).

“Shared, structured effects” simulations.

We simulated each b_j from the mixture model (eq. 3), with mixture parameters based on fit to the GTEx data. This is the simulation procedure in more detail:

1. We took the 8 “data-driven” covariance matrices U_1, \dots, U_8 learned from the GTEx data, standardized as described above.

2. We simulated 400 “non-null units”: independently for $j = 1, \dots, 400$, we (a) chose a component k uniformly at random from $1, \dots, 8$, (b) simulated a scaling factor ω as the absolute value of an $N(0,1)$ random variable, and (c) simulated $b_j \sim N_{44}(0, \omega U_k)$.
3. For the 19,600 “null units”, we set $b_j = 0$.
4. For all 20,000 units, we simulated $\hat{b}_j \sim N_{44}(b_j, V_j)$, where V_j was the diagonal matrix with diagonal elements 0.1^2 (so that standard errors were 0.1, similar to the estimates from the GTEx data).

“Shared, unstructured effects” simulations.

We simulated as for “shared, structured effects”, except that the 400 non-null units in Step 2c were simulated as $b_j \sim N_{44}(0, 0.1^2 I)$.

“Independent effects” simulations.

We also simulated data from effects that were entirely independent across conditions:

1. Independently for each $r = 1, \dots, 44$, we chose a random set of 400 $j \in \{1, \dots, 2 \times 10^4\}$ to be the “true effects”.
2. For the “true effects”, we simulated $b_{jr} \sim N(0, \Sigma^2)$, where Σ^2 was chosen with equal probability from the set $\{0.1, 0.5, 0.75, 1\}$. In this way, the true effects captured both small and large effects for each condition. All other effects were set to zero.
3. We simulated $\hat{b}_j \sim N_{44}(b_j, V_j)$ as in the other simulations.

Analysis of simulated data.

Each simulated dataset was analyzed using mash, following the procedure described above. Note that we re-estimated (U, π) from the data without making use of the true values used to simulate the data. We estimated effects by their posterior means (eq. 8), and assessed significance by the *Ifsr*. The ash and mash-bmalite methods were applied to the simulated data in the same way that they were applied to the GTEx data (see above).

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

This work was supported by NIH grants MH090951 and HG02585 to M.S., and by a grant from the Gordon and Betty Moore Foundation (GBMF #4559) to M.S. S.M.U. was supported by NIH grant T32HD007009. Expression (GTEx) Project was supported by the Common Fund of the Office of the Director of the National Institutes of Health. Additional funds were provided by the NCI, NHGRI, NHLBI, NIDA, NIMH, and NINDS. Donors were enrolled at Biospecimen Source Sites funded by NCI/SAIC-Frederick, Inc. (SAIC-F) subcontracts to the National Disease Research Interchange (10XS170), Roswell Park Cancer Institute (10XS171), and Science Care, Inc. (10XS172). The Laboratory, Data Analysis, and Coordinating Center (LDACC) was funded through a contract

(HHSN268201000029C) to The Broad Institute, Inc. Biorepository operations were funded through an SAIC-F subcontract to Van Andel Institute (10ST1035). Additional data repository and project management were provided by SAIC-F (HHSN261200800001E). The Brain Bank was supported by a supplement to University of Miami grants DA006227 & DA033684 and to contract N01MH000028. Statistical Methods development grants were made to the University of Geneva (MH090941 & MH101814), the University of Chicago (MH090951, MH090937, MH101820, MH101825), the University of North Carolina-Chapel Hill (MH090936 & MH101819), Harvard University (MH090948), Stanford University (MH101782), Washington University St. Louis (MH101810), and the University of Pennsylvania (MH101822). The data used for the analyses described in this manuscript were obtained from the GTEx Portal on October 17, 2015.

References

1. Blischak JD, Tailleux L, Mitrano A, Barreiro LB & Gilad Y Mycobacterial infection induces a specific human innate immune response. *Scientific Reports* 5, 16882 (2015). [PubMed: 26586179]
2. Ferguson JP, Cho JH & Zhao H A new approach for the joint analysis of multiple ChIP-Seq libraries with application to histone modification. *Statistical Applications in Genetics and Molecular Biology* 11 (2012).
3. Pickrell J, Berisa T, Séguérel L, Tung JY & Hinds D Detection and interpretation of shared genetic influences on 40 human traits. *Nature Genetics* 48, 709–717. [PubMed: 27182965]
4. Dimas AS et al. Common regulatory variation impacts gene expression in a cell type-dependent manner. *Science* 325, 1246–1250 (2009). [PubMed: 19644074]
5. Flutre T, Wen X, Pritchard J & Stephens M A statistical framework for joint eQTL analysis in multiple tissues. *PLoS Genetics* 9, e1003486 (2013). [PubMed: 23671422]
6. Li G, Shabalin AA, Rusyn I, Wright FA & Nobel AB An Empirical Bayes approach for multiple tissue eQTL Analysis. *Biostatistics* 19, 391–406 (2017).
7. Petretto E et al. New insights into the genetic control of gene expression using a Bayesian multi-tissue approach. *PLoS Computational Biology* 6, e1000737 (2010). [PubMed: 20386736]
8. Wen X & Stephens M Using linear predictors to impute allele frequencies from summary of pooled genotype data. *Annals of Applied Statistics* 4, 1158–1182 (2010). [PubMed: 21479081]
9. Han B & Eskin E Random-effects model aimed at discovering associations in meta-analysis of genome-wide association studies. *American Journal of Human Genetics* 88, 586–598 (2011). [PubMed: 21565292]
10. Stephens M Unified framework for association analysis with multiple related phenotypes. *PLoS ONE* 8, e65245 (2013). [PubMed: 23861737]
11. Sul JH, Han B, Ye C, Choi T & Eskin E Effectively identifying eQTLs from multiple tissues by combining mixed model and meta-analytic approaches. *PLoS Genetics* 9, e1003491 (2013). [PubMed: 23785294]
12. Wei Y, Tenzen T & Ji H Joint analysis of differential gene expression in multiple studies using correlation motifs. *Biostatistics* 16, 31–46 (2015). [PubMed: 25143368]
13. Zhou X & Stephens M Efficient multivariate linear mixed model algorithms for genome-wide association studies. *Nature Methods* 11, 407–409 (2014). [PubMed: 24531419]
14. Han B & Eskin E Interpreting meta-analyses of genome-wide association studies. *PLoS Genetics* 8, e1002555 (2012). [PubMed: 22396665]
15. Lebrech JJ, Stijnen T & van Houwelingen HC Dealing with heterogeneity between cohorts in genomewide SNP association studies. *Statistical Applications in Genetics and Molecular Biology* 9 (2010).
16. Stephens M False discovery rates: a new deal. *Biostatistics* 18: 275–294 (2017). [PubMed: 27756721]
17. GTEx Consortium. The Genotype-Tissue Expression (GTEx) pilot analysis: multitissue gene regulation in humans. *Science* 348, 648–660 (2015). [PubMed: 25954001]
18. Nicolae DL, Gamazon E, Zhang W, Duan S, Dolan ME & Cox NJ Trait-associated SNPs are more likely to be eQTLs: annotation to enhance discovery from GWAS. *PLoS Genetics* 6, e1000888 (2010). [PubMed: 20369019]
19. Shabalin AA Matrix eQTL: ultra fast eQTL analysis via large matrix operations. *Bioinformatics* 28, 1353–1358 (2012). [PubMed: 22492648]

20. Price AL, Patterson NJ, Plenge RM, Weinblatt ME, Shadick NA & Reich D Principal components analysis corrects for stratification in genome-wide association studies. *Nature Genetics* 38, 904–909 (2006). [PubMed: 16862161]
21. Engelhardt BE & Stephens M Analysis of population structure: a unifying framework and novel methods based on sparse factor analysis. *PLoS Genetics* 6, e1001117 (2010). [PubMed: 20862358]
22. Tversky A & Kahneman D Judgment under uncertainty: heuristics and biases. *Science* 185, 1124–1131 (1974). [PubMed: 17835457]
23. Wen X & Stephens M Bayesian methods for genetic association analysis with heterogeneous subgroups: from meta-analyses to gene-environment interactions. *Annals of Applied Statistics* 8, 176–203 (2014). [PubMed: 26413181]
24. Servin B & Stephens M Imputation-based analysis of association studies: candidate regions and quantitative traits. *PLoS Genetics* 3, e114 (2007). [PubMed: 17676998]
25. Veyrieras J-B, Kudaravalli S, Kim SY, Dermitzakis ET, Gilad Y, Stephens M & Pritchard JK High-resolution mapping of expression-QTLs yields insight into human gene regulation. *PLoS Genetics* 4, e1000214 (2008). [PubMed: 18846210]
26. Hormozdiari F, Kostem E, Kang EY, Pasaniuc B & Eskin E Identifying causal variants at loci with multiple signals of association. *Genetics* 198, 497–508 (2014). [PubMed: 25104515]
27. Kichaev G, Yang W-Y, Lindstrom S, Hormozdiari F, Eskin E, Price AL, Kraft P & Pasaniuc B Integrating functional data to prioritize causal variants in statistical fine-mapping studies. *PLoS Genetics* 10, e1004722 (2014). [PubMed: 25357204]
28. Chen W, Larrabee BR, Ovsyannikova IG, Kennedy RB, Haralambieva IH, Poland GA & Schaid DJ Fine mapping causal variants with an approximate Bayesian method using marginal test statistics. *Genetics* 200, 719–736 (2015). [PubMed: 25948564]
29. Moyerbrailean GA, Kalita CA, Harvey CT, Wen X, Luca F & Pique-Regi R Which genetic variants in DNase-seq footprints are more likely to alter binding? *PLOS Genetics* 12, e1005875 (2016). [PubMed: 26901046]
30. Giambartolomei C, Vukcevic D, Schadt EE, Franke L, Hingorani AD, Wallace C & Plagnol V Bayesian test for colocalisation between pairs of genetic association studies using summary statistics. *PLoS Genetics* 10, e1004383 (2014). [PubMed: 24830394]
31. Fortune MD, Guo H, Burren O, Schofield E, Walker NM, Ban M, Sawcer SJ, Bowes J, Worthington J, Barton A, Eyre S, Todd JA & Wallace C Statistical colocalization of genetic risk variants for related autoimmune diseases in the context of common controls. *Nature Genetics* 47, 839–846 (2015). [PubMed: 26053495]
32. Nica AC, Montgomery SB, Dimas AS, Stranger BE, Beazley C, Barroso I & Dermitzakis ET Candidate causal regulatory effects by integration of expression QTLs with complex trait genetic associations. *PLoS Genetics* 6, e1000895 (2010). [PubMed: 20369022]
33. Benjamini Y & Hochberg Y Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society, Series B* 57, 289–300 (1995).
34. Storey JD The positive false discovery rate: a Bayesian interpretation and the q-value. *Annals of Statistics* 31, 2013–2035 (2003).

Methods-only references

35. Bovy J, Hogg DW & Roweis ST Extreme Deconvolution: inferring complete distribution functions from noisy, heterogeneous and incomplete observations. *Annals of Applied Statistics* 5, 1657–1677 (2011).
36. Larribe F & Fearnhead P Composite likelihood methods in statistical genetics. *Statistica Sinica* 21, 43–69 (2011).
37. Dempster AP, Laird NM & Rubin DB Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B* 39, 1–38 (1977).
38. Varadhan R & Roland C Simple and globally convergent methods for accelerating the convergence of any EM algorithm. *Scandinavian Journal of Statistics* 35, 335–353 (2008).
39. Efron B Microarrays, empirical Bayes and the two-groups model. *Statistical Science* 23, 1–22 (2008).

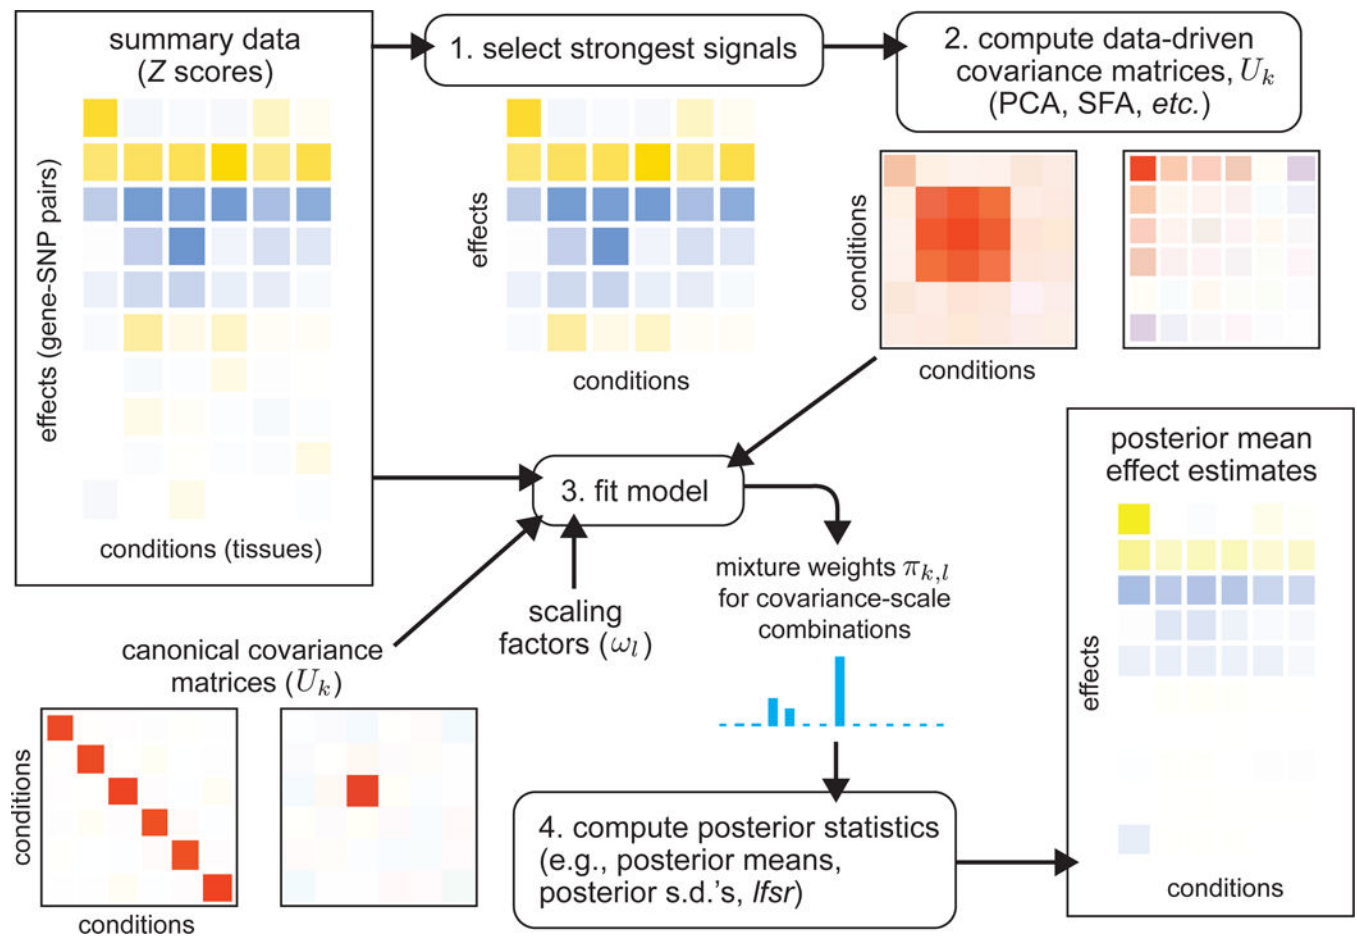


Figure 1. Overview of fitting procedure in mash, which estimates the multivariate distribution of effects present in the data.

The data consist of a matrix of summary data (e.g., Z scores) for a large number of units (e.g., gene-SNP pairs) in multiple conditions (e.g., tissues), and, optionally, their standard errors (not shown). Color indicates the sign (positive, negative) of an effect (blue, yellow) or covariance (blue, red), with shading intensity indicating size. After selecting rows containing the strongest signals (1)—in this example, the top 6 rows—we apply covariance estimation techniques to estimate candidate “data-driven” covariance matrices U_k (2). To these, we add “canonical” covariance matrices U_k , including the identity matrix, and matrices representing condition-specific effects. Each covariance matrix represents a pattern of effects that may occur in the data. We scale each covariance matrix by a grid of scaling factors, ω_l , varying from “very small” to “very large”, which allows for *a priori* effect sizes to range from very small to very large. Using the entire data set, we compute maximum-likelihood estimates of the weights (relative frequencies) $\pi_{k,l}$ for each (U_k, ω_l) combination (3), thereby learning how commonly each pattern-effect size combination occurs in the data. Finally, we compute posterior statistics using the fitted model (4); the posterior mean estimates shown in the bottom-right illustrate that effect estimates are “shrunk” adaptively using the fitted mash model.

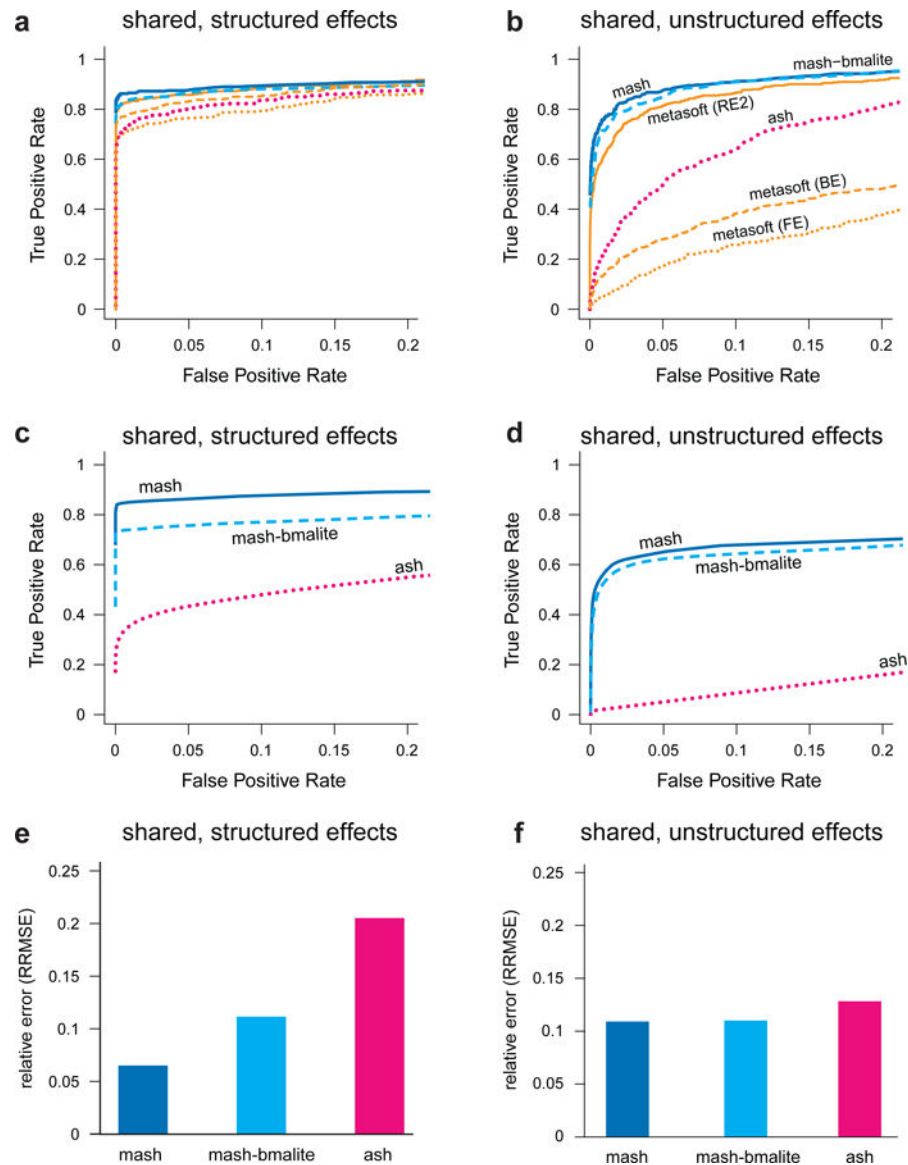


Figure 2. Comparison of methods on simulated data.

Results are shown for two simulation scenarios: “shared, structured effects”, in which the non-zero effects are shared among conditions in complex, structured ways similar to patterns of eQTL sharing in the GTEx data; and “shared, unstructured effects”, in which the non-zero effects are shared among conditions but independent. Each simulation result involves $n = 20,000$ independent units observed at $R = 44$ conditions, with 400 non-null units. Panels **a–b** show ROC curves for detecting significant units ($n = 20,000$ discoveries), based on unit-specific measures of significance (as in traditional meta-analyses). Panels **c–d** show ROC curves for detecting significant effects ($n \times R = 44 \times 20,000 = 880,000$ discoveries), which requires effect-specific measures of significance. In **c–d**, we also require the estimated sign (+/–) of each significant effect to be correct to be considered a “true positive”. Panels **e** and **f** summarize the error in the estimated effects relative to the error from a simple condition-by-condition analysis (Relative Root Mean Squared Error, or RRMSE for short). Our new

method (mash) outperformed other methods, particularly in the “shared, structured effects” scenario.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

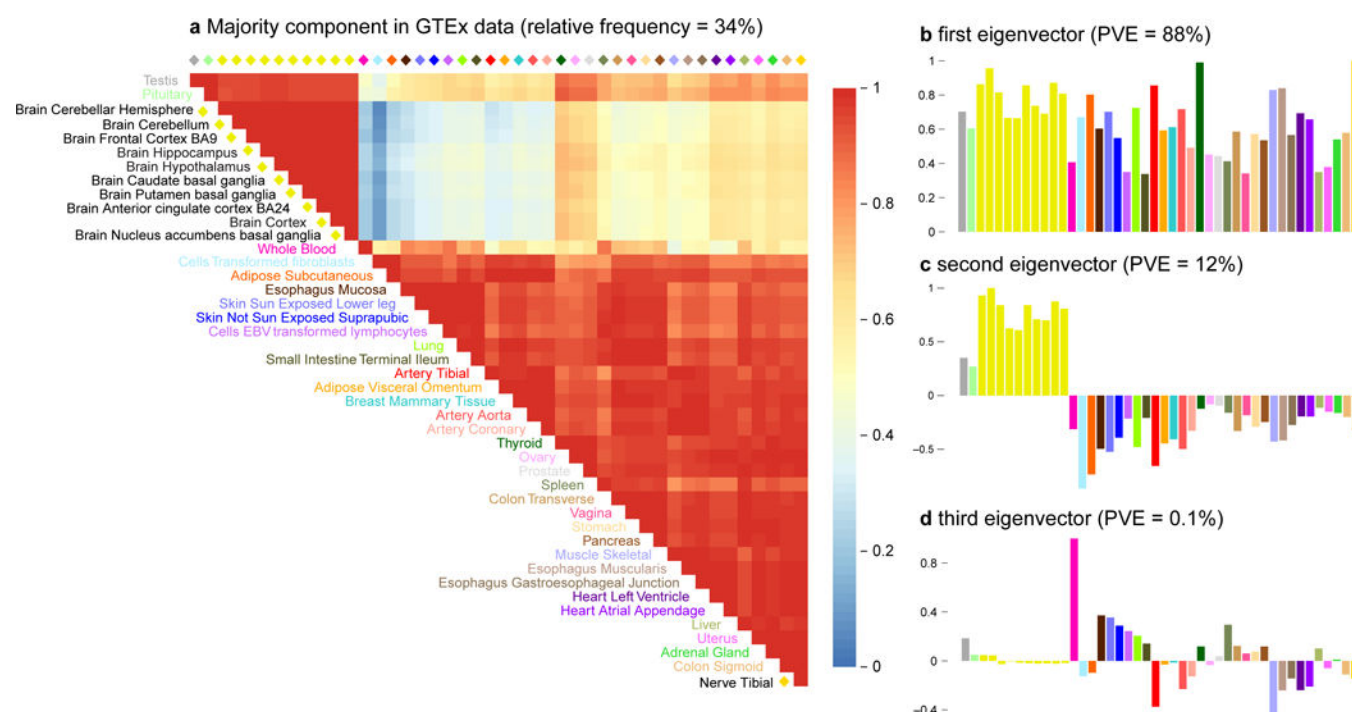


Figure 3. Summary of primary patterns identified by mash in GTEx data.

Shown are the heatmap of the correlation matrix (**a**) and bar plots of the first three eigenvectors (**b**, **c**, **d**) of the covariance matrix U_k corresponding to the dominant mixture component identified by mash ($n = 16,069$ independent gene-SNP pairs). This component accounts for 34% of all weight in the GTEx data. Tissues are color-coded as indicated by the tissue labels in the heatmap. The first eigenvector (**b**) reflects broad sharing among all tissues, with all effects in the same direction; the second eigenvector (**c**) captures differences between brain (and, to a lesser extent, testis and pituitary) and other tissues; the third eigenvector (**d**) primarily captures effects that are stronger in whole blood.

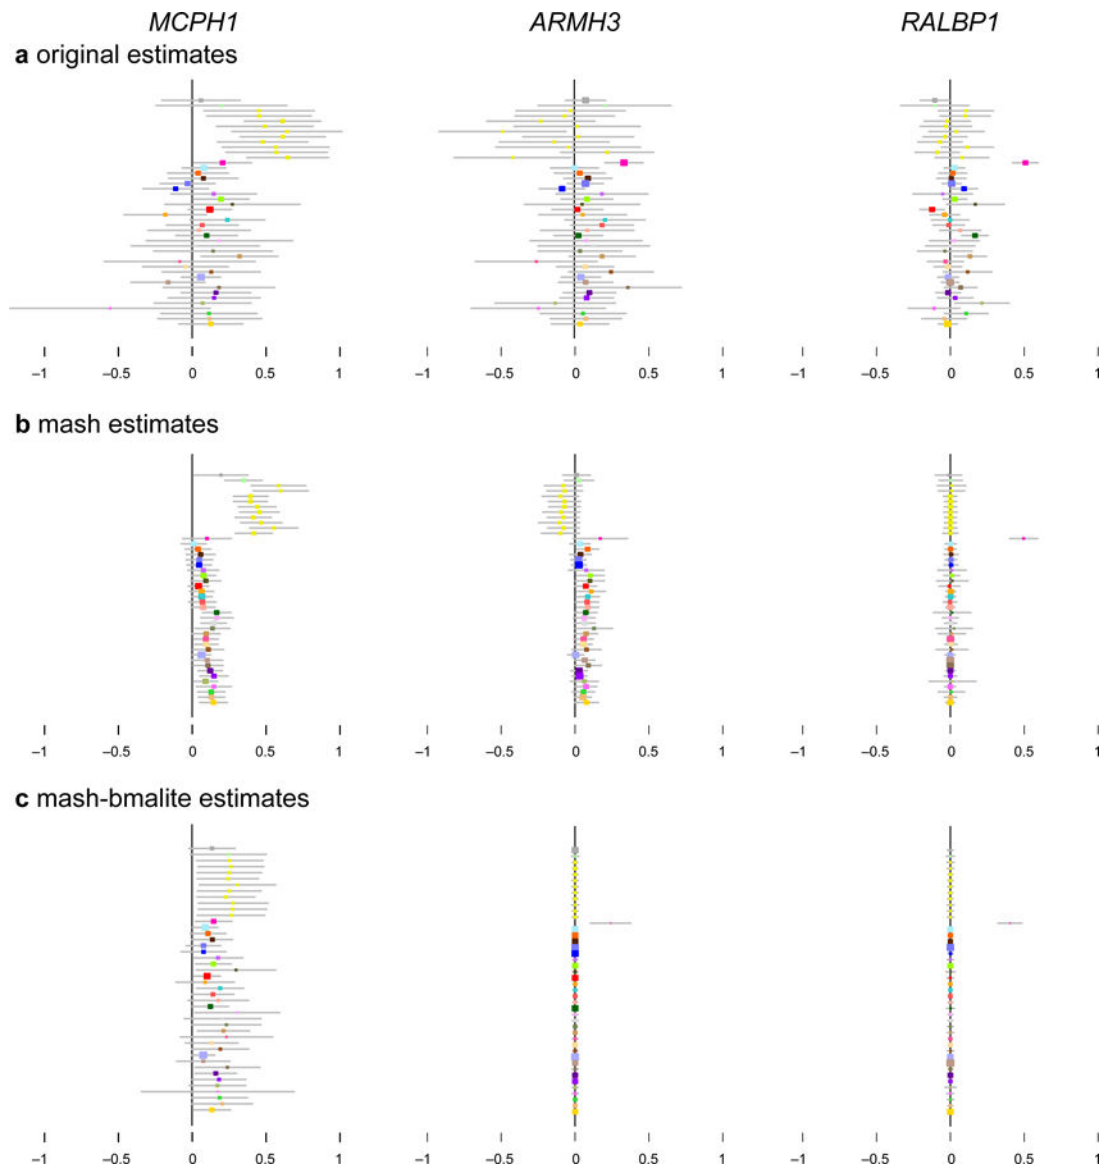


Figure 4. Examples illustrating that mash uses learned patterns of sharing to inform effect estimates in the GTEx data.

In panel **a**, each colored dot shows the original ("raw") effect estimate for a single tissue (color-coded as in Fig. 3), with grey bars indicating ± 2 standard errors. These are the data provided to mash. Panel **b** shows the corresponding mash estimates. In each case, mash combines information across all tissues, using the background information (patterns of sharing) learned from data on all eQTLs to produce more precise estimates. Panel **c** shows, for contrast, the corresponding estimates from mash-bmalite, which, due to its more restricted model, fails to capture features clearly apparent in the original data, such as strong brain effects in *MCPH1*. In **b** and **c**, colored dots are posterior means, and error bars depict ± 2 posterior standard deviations. For all estimates, $n = 83$ –430 individuals, depending on the tissue (Supplementary Table 3).

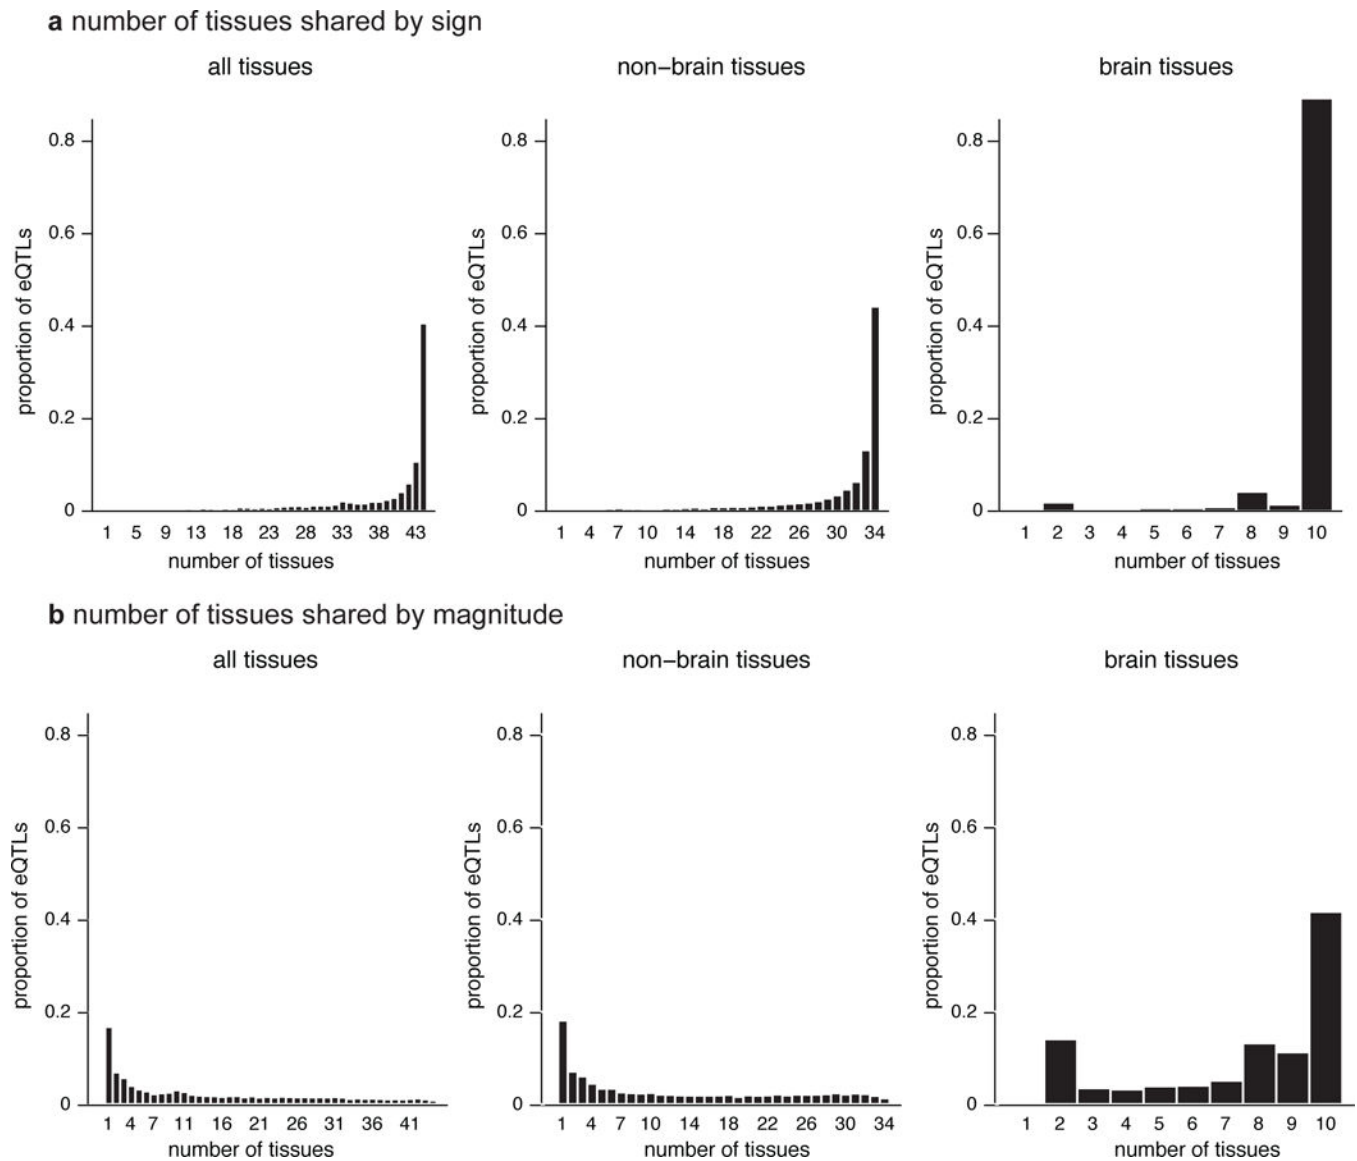


Figure 5. Number of tissues shared by sign and magnitude.

Histograms show estimated number of tissues in which top eQTLs are “shared,” considering all tissues ($n = 12,171$ gene-SNP pairs with a significant eQTL in at least one tissue), non-brain tissues ($n = 12,117$), and brain tissues only ($n = 8,474$), and using two different sharing definitions, by sign (**a**) and by magnitude (**b**). Sharing by sign means that the eQTLs have the same sign in the estimated effect; sharing by magnitude means that they also have similar effect sizes (within a factor of 2).

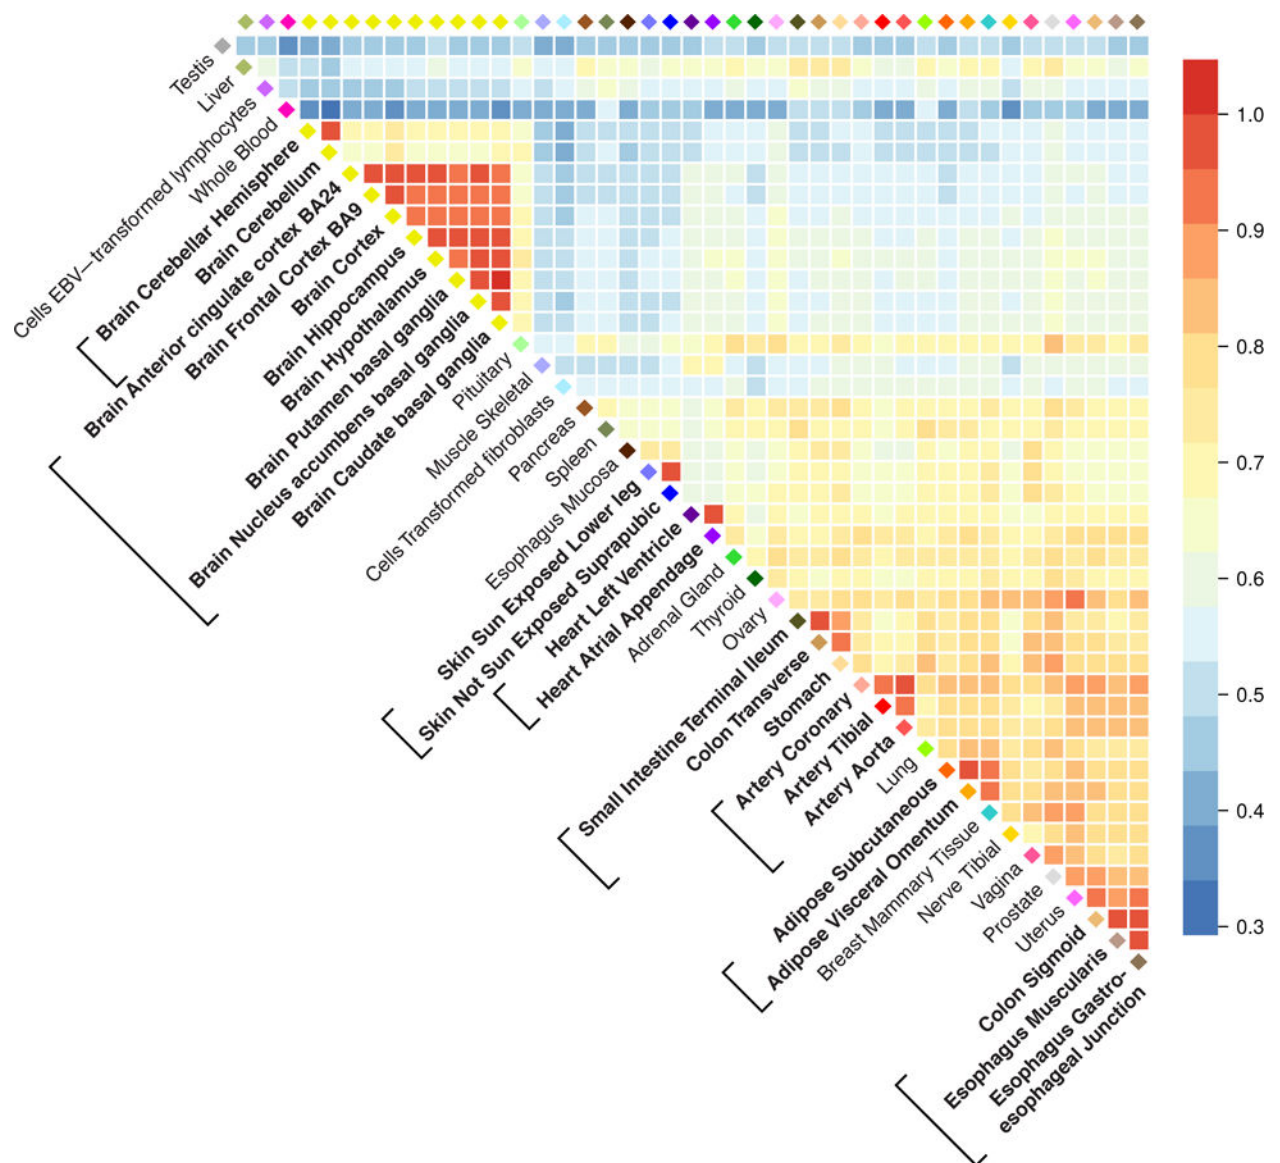


Figure 6. Pairwise sharing by magnitude of eQTLs among tissues.

For each pair of tissues, we considered the top eQTLs that were significant ($P < 0.05$) in at least one of the two tissues, and plotted the proportion of these that are “shared in magnitude”—that is, have effect estimates that are the same sign and within a factor of 2 in size of one another ($n = 5,605$ – $9,811$ gene-SNP pairs, depending on pair of tissues compared). Brackets surrounding tissue labels highlight groups of biologically related tissues mentioned in the text as showing particularly high levels of sharing.

Table 1
Summary of sharing among top eQTLs.

| Numbers give the proportion of effects meeting a given sharing criterion: “shared by sign” requires that the effect have the same sign as the strongest effect across tissues; “shared by magnitude” requires that the effect also be within a factor of 2 of the strongest effect. Numbers in parentheses were obtained from secondary mash analyses conducted separately in brain and non-brain tissues.

sharing criterion	proportion of effects		
	all tissues	non-brain	brain
shared by sign	0.85	0.85 (0.88)	0.96 (0.98)
shared by magnitude	0.36	0.40 (0.44)	0.76 (0.85)