PCA outperforms popular hidden variable inference methods for QTL mapping

Heather J. Zhou[1], Lei Li[2], Yumei Li[3], Wei Li[3], Jingyi Jessica Li[1,4,5,6,*]

[1]Department of Statistics, University of California, Los Angeles, CA 90095
[2]Institute of Systems and Physical Biology, Shenzhen Bay Laboratory, Shenzhen, China 518055
[3]Division of Computational Biomedicine, Department of Biological Chemistry, School of Medicine, University of California, Irvine, Irvine, CA 92697
[4]Department of Human Genetics, University of California, Los Angeles, CA 90095
[5]Department of Computational Medicine, University of California, Los Angeles, CA 90095
[6]Department of Biostatistics, University of California, Los Angeles, CA 90095
[*]To whom correspondence should be addressed. Email: lijy03@g.ucla.edu

Abstract

Estimating and accounting for hidden variables is widely practiced as an important step in quantitative trait locus (QTL) analysis for improving the power of QTL identification. Here we benchmark popular hidden variable inference methods including surrogate variable analysis (SVA), probabilistic estimation of expression residuals (PEER), and hidden covariates with prior (HCP) against principal component analysis (PCA; a well-established dimension reduction and factor discovery method) through comprehensive simulation studies and show that PCA is faster and better-performing. In particular, we show that under realistic simulation scenarios, PEER—the most popular hidden variable inference method for QTL mapping to date—does not have the alleged advantage that its performance does not deteriorate as the number of inferred covariates increases. Additionally, our real data analysis shows that PEER produces identical results as PCA at best, is at least three orders of magnitude slower, and can fail to capture important variance components of the molecular phenotype data in certain cases. Based on these results and the fact that PCA is much easier to interpret and use, we contend that PCA should be the preferred hidden variable inference method for QTL mapping until evidence suggests otherwise. To help researchers use PCA in their QTL analysis, we provide an R package `PCAForQTL` along with a detailed tutorial.

1 Introduction

Genome-wide association studies (GWASs) have identified thousands of genetic variants associated with human traits or diseases [1–4]. However, the majority of GWAS variants are located in non-coding regions of the genome, making it challenging to interpret the GWAS associations [5, 6]. In response to this, (molecular) quantitative trait locus (QTL) analysis has

emerged as an important field in human genetics, interrogating the relationship between genetic variants and intermediate, molecular traits and potentially explaining the GWAS associations [7, 8].

Based on the type of molecular phenotype studied, QTL analyses can be categorized into gene expression QTL (eQTL) analyses [9, 10], alternative splicing QTL (sQTL) analyses [10], three prime untranslated region alternative polyadenylation QTL ($3'$aQTL) analyses [11], and so on [7, 8]. Among these categories, eQTL analyses, which investigate the association between genetic variants and gene expression levels, are the most common. To date, most (single-tissue) QTL studies are carried out using regression-based methods such as Matrix eQTL [12] and FastQTL [13].

In QTL analysis, a major challenge is that measurements of gene expression levels and other molecular phenotypes can be affected by a number of technical or biological variables other than the genetic variants (e.g., batch, sex, and age). If these variables are known, then they can be directly included in the QTL pipeline as covariates. However, many of these variables may be unknown or unmeasured. Therefore, it has become standard practice to *first* infer the hidden variables and *then* include the inferred variables as covariates (or account for them in other ways; Table 1) [9–11, 14–23]. This type of approach has been shown to both improve the power of QTL identification in simulation settings [24] and empirically increase the number of discoveries in QTL studies [9, 10, 16, 21–23].

Surrogate variable analysis (SVA) [25, 26] is one of the first popular hidden variable inference methods for large-scale genomic analysis. Although initially proposed as a hidden variable inference method for both QTL mapping and differential expression (DE) analysis, currently SVA is primarily used in DE and similar analyses as opposed to QTL mapping [27–30]. We believe this is partly because the SVA package [31] is difficult to apply in QTL settings in that it requires the user to input at least one variable of interest and using too many variables of interest causes the package to fail (Figure 1; Section S4); while there are usually at most a few variables of interest in a DE study, there are often millions of single nucleotide polymorphisms (SNPs; variables of interest) in a QTL study. Historically, there have been two versions of the SVA method: two-step SVA [25] and iteratively reweighted SVA (IRW-SVA) [26]; the latter supersedes the former. Therefore, we focus on IRW-SVA in this work.

Probabilistic estimation of expression residuals (PEER) [24, 32] is arguably the most popular hidden variable inference method for QTL mapping to date. It is used in the Genotype-Tissue Expression (GTEx) project [9, 10] and many other high-impact studies [11, 14–21]. The PEER method has two main perceived advantages: (1) it can take known covariates into account when estimating the hidden covariates, and (2) its performance does not deteriorate as the number of inferred covariates increases (i.e., it does not "overfit"). One drawback of PEER, though, is that there is no consensus in the literature on *how* it should be used. For example, when there are known covariates available, PEER can be run with or without the known covariates—Stegle et al. [32] do not give an explicit recommendation as to which approach should be used, and both approaches are used in practice (e.g., [9, 10] vs. [11, 16]). Further, PEER outputs both inferred covariates and residuals of the inputted molecular phenotypes (Figure 1), so the user needs to

decide which set of outputs to use (Section S4; we refer to the approach using the inferred covariates as the "factor approach" and the approach using the residuals as the "residual approach"). Such "flexibility" of PEER could be considered a benefit, but we believe it not only leads to confusion for practitioners but also reduces the transparency and reproducibility of published QTL research.

Hidden covariates with prior (HCP) [33] is another popular hidden variable inference method for QTL mapping. Though less popular than PEER, it is also used in some high-impact studies [22, 23]. To determine which method is the best and whether PEER indeed has the perceived advantages, we thoroughly evaluate SVA, PEER, and HCP for the first time in the literature. Given that principal component analysis (PCA) [34, 35] underlies the methodology behind each of these methods (Section 2.4), we also include PCA in our evaluation. Through simulation studies (Section 2.1) and real data analysis (Sections 2.2, 2.3 and 2.5), we show that PCA is better than or as good as all of the other methods in various regards (Figure 1).

2 Results

2.1 Comprehensive simulation studies show that PCA is faster and better-performing

We compare the runtime and performance of 15 methods, including Ideal (assuming the hidden covariates are known), Unadjusted (not estimating or accounting for the hidden covariates), and 13 variants of PCA, SVA, PEER, and HCP (Table 1), based on two simulation studies. In the first simulation study (Simulation Design 1; Section S2), we follow the data simulation in Stegle et al. [24], the original PEER publication, while addressing its data analysis and overall design limitations (Section S1). In the second simulation study (Simulation Design 2; Section S3), we further address the data simulation limitations of Stegle et al. [24] (Section S1) by simulating the data in a more realistic and comprehensive way, roughly following Wang et al. [36]—the SuSiE publication—but introducing the existence of hidden covariates. A summary of the main differences between the two simulation designs is provided in Table S1. The key difference is that in Simulation Design 1, the gene expression levels are primarily driven by trans-regulatory effects rather than cis-regulatory effects or covariate effects (Table S2), inconsistent with the common belief that trans-regulatory effects are generally weaker than cis-regulatory effects. In contrast, in Simulation Design 2, we focus on cis-QTL detection and carefully control the genotype effects and covariate effects in 176 experiments with two replicates per experiment (Section S3).

The details of the 15 methods are described in Section S4, and the evaluation metrics are described in Section 3.1. For convenience, we refer to the simulated molecular phenotypes as gene expression levels throughout our simulation studies; however, they can be interpreted as any type of molecular phenotype after data preprocessing and transformation, e.g., alternative splicing phenotypes and alternative polyadenylation phenotypes (Table S3).

The results from our simulation studies are summarized in Figures 2, 3, S1, S3, and S4. We find that PCA and HCP are orders of magnitude faster than SVA, which in turn is orders of magnitude faster than PEER, and that PCA outperforms SVA, PEER, and HCP in terms of area under the precision-recall curve (AUPRC; Figures 2 and 3). On a dataset-by-dataset basis, PCA outperforms

the other methods in terms of AUPRC in 11% to 88% of the simulated data sets and underperforms them in close to 0% of the simulated data sets in Simulation Design 2 (Figure S3(d)). In addition, PCA has the highest average concordance score(s), a metric for the concordance between the true hidden covariates and the inferred covariates (Section 3.1; Figures S1 and S4), which explains why PCA performs the best in terms of AUPRC.

To contrast the results in Stegle et al. [24], we also compare the powers of the different methods in Simulation Design 1 (Figure S1). We find that PCA is more powerful than SVA, PEER, and HCP. Notably, SVA and PEER have very low power in identifying trans-QTL relations—an especially unfavorable result for SVA and PEER, considering that the gene expression levels are primarily driven by trans-regulatory effects in Simulation Design 1 (Table S2).

Incidentally, Figures 2 and S3 also provide us with the following insights into the different ways of using PEER (Section S4). First, running PEER *with* the known covariates has no advantage over running PEER *without* the known covariates in terms of AUPRC, given the choice of K (the number of inferred covariates) and the choice between the factor approach and the residual approach. In fact, running PEER *with* the known covariates significantly increases the runtime of PEER in real data (Figure S6). Second, contrary to claims in Stegle et al. [24, 32], the performance of PEER *does* deteriorate as the number of PEER factors increases. The only exception is when the residual approach is used in Simulation Design 1 (Figure 2). But given that Simulation Design 2 is more realistic than Simulation Design 1 and that the factor approach is more popular than the residual approach [9–11, 17–20], the take-home message should be that in general, the performance of PEER is worse when we use a large K rather than the true K. Third, whether the factor approach or the residual approach performs better depends on the choice of K. When we use the true K, the factor approach performs better, but when we use a large K, the residual approach performs better. All in all, PCA outperforms all different ways of using PEER in both of our simulation studies (Figure 2).

2.2 PEER factors sometimes fail to capture important variance components of the molecular phenotype data

For our real data analysis, we examine the most recent GTEx eQTL and sQTL data [10] (Sections 2.3 and 2.5) and the 3′aQTL data prepared by Li et al. [11] from GTEx RNA-seq reads [9] (Section 2.2). While the exact data analysis pipelines are different (Table S3), these studies all choose PEER as their hidden variable inference method.

Unlike PCs, which are always uncorrelated (Section S5.1), PEER factors are not guaranteed to be uncorrelated. Here we show through the above-mentioned 3′aQTL data that PEER factors can be highly correlated with each other (to the extent that many or all of them are practically identical) and thus fail to capture important variance components of the molecular phenotype data.

Given a post-imputation alternative polyadenylation phenotype matrix (each entry is between zero and one, representing a proportion), Li et al. [11] run PEER without further data transformation using the number of PEER factors chosen by GTEx [9] (Table S3). To assess the impact of data transformation on the PEER factors, we also run PEER after transforming the data in three ways:

(1) center and scale (to unit variance) each feature, (2) apply inverse normal transform (INT) [37] to each feature ("INT within feature"), and (3) apply INT to each sample ("INT within sample"). Among these methods, GTEx [9, 10] uses "INT within feature" for its eQTL data and "INT within sample" for its sQTL data (Table S3). To quantify how many "distinct" or "nonrepetitive" PEER factors there are, given a set of PEER factors, we group them into clusters such that in each cluster, the correlation between any two PEER factors is above a pre-defined threshold (0.99, 0.9, or 0.8) in absolute value; this is done via hierarchical clustering [38] with complete linkage and the distance defined as one minus the absolute value of the correlation. Therefore, the number of PEER factor clusters can be interpreted as the number of distinct or nonrepetitive PEER factors.

Our results show that in many cases, the number of distinct PEER factors is considerably smaller than the number of PEER factors requested (Figure 4), and when this issue is severe (e.g., "No transformation" and "INT within sample"), the PEER factors fail to capture important variance components of the molecular phenotype data (Figure S5). Since the number of discoveries increases substantially with the number of PEER factors in GTEx's eQTL analyses [9, 10], where the PEER factors are essentially identical to PCs (Section 2.3), it is possible that replacing the nearly-all-identical PEER factors with appropriate numbers of PCs in Li et al. [11]'s 3′aQTL analysis can lead to more discoveries. This is a potential direction for a future study.

2.3 PEER factors are almost identical to PCs but take three orders of magnitude longer to compute in GTEx eQTL and sQTL data

We report the surprising finding that in both GTEx eQTL and sQTL data [10], the PEER factors obtained by GTEx and used in its QTL analyses are almost identical to PCs. Specifically, given a fully processed molecular phenotype matrix, there is almost always a near-perfect one-to-one correspondence between the PEER factors and the top PCs (Figure 5). This means that after the variational Bayesian inference in PEER initializes with PCs [24], it does not update the PCs much beyond scaling them (see Section 2.4 for a theoretical explanation). Therefore, it is no surprise that replacing the PEER factors with PCs in GTEx's FastQTL pipeline [10, 13] produces almost identical QTL results (Figures S6 and S7) because in linear regressions (the basis of both Matrix eQTL [12] and FastQTL [13]), scaling and/or shifting the predictors does not change the p-values of t-tests for non-intercept terms (neither does scaling and/or shifting the response, for that matter).

However, PEER is at least three orders of magnitude slower than PCA (Figure S6). For a given expression matrix, running PEER without the known covariates (GTEx's approach) takes up to about 32 hours, while running PCA (with centering and scaling; our approach) takes no more than a minute.

2.4 PCA, SVA, PEER, and HCP are closely related statistical methods

We report that PCA, SVA, PEER, and HCP are closely related statistical methods despite their apparent dissimilarities. In particular, the methodology behind SVA, PEER, and HCP can all be traced back to PCA (Figure 6).

PCA [34, 35] is traditionally derived by optimizing some objective functions (either maximum variance or minimum reconstruction error; Section S5.1), but more recently, it is shown [39] that PCA can be derived as a limiting case of probabilistic principal component analysis (PPCA) [39], which in turn is a special case of factor analysis [35, 40]—a dimension reduction method commonly used in psychology and the social sciences that is based on a frequentist probabilistic model.

PEER [24, 32] is based on a Bayesian probabilistic model and can be considered a Bayesian version of factor analysis (with the not-very-useful ability to explicitly model the known covariates; Section 2.1). Therefore, PEER and PCA are closely related statistical methods, which explains why PEER factors are almost identical to PCs in GTEx eQTL and sQTL data [10] (Section 2.3).

SVA [25, 26] is purely an algorithm and is not defined based on a probabilistic model or objective function. Though complicated, the SVA algorithm is heavily based on PCA. In a nutshell, SVA iterates between two steps: (1) reweight the features of the molecular phenotype matrix, and (2) perform PCA on the resulting matrix (with centering but without scaling) [26].

Lastly, HCP [33] is defined by minimizing a loss function that is very similar to the minimum-reconstruction-error loss function of PCA (Section S5.2). In short, SVA, PEER, and HCP can all be considered extensions or more complex versions of PCA (though we show that the complexity is a burden rather than a benefit; Figure 1).

2.5 PCA provides insight into the choice of K

Choosing K, the number of inferred covariates in the context of hidden variable inference or the number of dimensions or clusters in more general contexts, is always a difficult task. Nonetheless, based on the proportion of variance explained (PVE) by each PC (Section S5.1), PCA offers convenient ways of choosing K such as the elbow method and the Buja and Eyuboglu (BE) algorithm [41] (more details below). Since SVA is heavily based on PCA (Section 2.4), it is able to adapt and make use of the BE algorithm. In contrast, PEER and HCP do not offer easy ways of choosing K; for lack of a better method, users of PEER and HCP often choose K by maximizing the number of discoveries [9, 10, 16, 21–23]. Not only is this approach of choosing K extremely computationally expensive and theoretically questionable, here we also show from the perspective of PCA that it may yield inappropriate choices of K.

Recall from Section 2.3 that PEER factors are almost identical to PCs in GTEx eQTL data [10] (the number of PEER factors is chosen by maximizing the number of discovered cis-eGenes for each pre-defined sample size bin; Table S3). Therefore, for each tissue type, we compare the number of PEER factors selected by GTEx to (1) the number of PCs chosen via an automatic elbow detection method (Algorithm S2) and (2) the number of PCs chosen via the BE algorithm (Algorithm S3; the default parameters are used). The BE algorithm is a permutation-based approach for choosing K in PCA. Intuitively, it retains PCs that explain more variance in the data than by random chance and discards those that do not. Hence, based on the statistical interpretation of the BE algorithm and the scree plots (examples shown in Figure 7), we believe that the number of PCs chosen via BE

should be considered an upper bound of the reasonable number of PCs to choose in GTEx eQTL data [10].

Our results show that the number of PEER factors selected by GTEx is almost always greater than the number of PCs chosen via *BE*, which in turn is almost always greater than the number of PCs chosen via *elbow* (Figure 7). In particular, the number of PEER factors selected by GTEx far exceeds the number of PCs chosen via *BE* for many tissue types with sample size above 350, suggesting that the number of PEER factors selected by GTEx may be too large. This hypothesis is further supported by the fact that we can reduce the number of inferred covariates to between 20% and 40% of the number of PEER factors selected by GTEx without significantly reducing the number of discovered cis-eGenes (Figure 7).

3 Methods

3.1 Evaluation metrics

Given a simulated data set, we evaluate each of the 15 methods in Table 1 mainly in three ways (when applicable): runtime, AUPRC, and adjusted R^2 measures (including adjusted R^2, reverse adjusted R^2, and concordance score).

First, we record the runtime of the hidden variable inference step (Section S4; not applicable for Ideal and Unadjusted).

Second, we calculate the area under the precision-recall curve (AUPRC) of the QTL result. We use AUPRC rather than the area under the receiver operating characteristic curve (AUROC) because AUPRC is more appropriate for data sets with imbalanced classes (there are far more negatives than positives in our simulated data sets and in QTL settings in general). Since AUPRC measures the trade-off between the true positive rate (i.e., power) and the false discovery rate (i.e., one minus precision), it is a more comprehensive metric than power. However, to contrast the results in Stegle et al. [24], we also compare the powers of the different methods in Simulation Design 1.

Third, for each simulated data set, each method except Ideal and Unadjusted gets an adjusted R^2 score (short as "adjusted R^2"), a reverse adjusted R^2 score (short as "reverse adjusted R^2"), and a concordance score. The adjusted R^2 score summarizes how well the true hidden covariates can be captured by the inferred covariates; the reverse adjusted R^2 score summarizes how well the inferred covariates can be captured by the true hidden covariates (a low score indicates that the inferred covariates are invalid or "meaningless"); lastly, the concordance score is the average of the previous two scores and thus measures the concordance between the true hidden covariates and the inferred covariates. Specifically, given m true hidden covariates and n inferred covariates, first, we calculate m adjusted R^2's (regressing each true hidden covariate against the inferred covariates) and n reverse adjusted R^2's (regressing each inferred covariate against the true hidden covariates); then, we average the m adjusted R^2's to obtain the adjusted R^2 score and average the n reverse adjusted R^2's to obtain the reverse adjusted R^2 score; finally, we define the concordance score as the average of the adjusted R^2 score and the reverse adjusted R^2 score.

3.2 Selection of representative methods for detailed comparison

Here we describe how we select a few representative methods from the 15 methods for detailed comparison in Simulation Design 2 (Table 1). From Figures 2(d) and S3, we see that the two PCA methods perform almost identically, so for simplicity, we select PCA_direct_screeK. The two SVA methods perform almost identically as well, so we select SVA_BE. For PEER, whether the known covariates are inputted when PEER is run has little effect on the AUPRC. Further, we observe that when we use the true K, the factor approach outperforms the residual approach, but when we use a large K, the residual approach outperforms the factor approach. Therefore, we select PEER_withCov_trueK_factors and PEER_withCov_largeK_residuals as the representative PEER methods. In addition, Ideal, Unadjusted, and HCP_trueK are selected.

4 Discussion

Hidden variable inference is widely practiced as an important step in QTL mapping for improving the power of QTL identification. Popular hidden variable inference methods include SVA, PEER, and HCP. In this work, we show that PCA not only underlies the statistical methodology behind the popular methods (Section 2.4), but is also better than or as good as them in various regards (Figure 1).

On the simulation front, we compare the runtime and performance of PCA, SVA, PEER, and HCP via two simulation studies (Section 2.1). In the first simulation study, we follow the data simulation in Stegle et al. [24], the original PEER publication, while addressing its data analysis and overall design limitations. In the second simulation study, we further address the data simulation limitations of Stegle et al. [24] by simulating the data in a more realistic and comprehensive way. Both simulation studies unanimously show that PCA is faster and better-performing. Further, they show that running PEER *with* the known covariates has no advantage over running PEER *without* the known covariates—in fact, running PEER *with* the known covariates makes PEER significantly slower (Figure S6)—and that contrary to claims in Stegle et al. [24, 32], the performance of PEER *does* deteriorate as the number of PEER factors increases (Section 2.1). One caveat of our simulation studies, though, is that the genotype and covariates all have linear effects on the gene expression levels (consistent with Stegle et al. [24] and Wang et al. [36]). But since PCA, SVA, PEER, and HCP are all linear methods or assume linearity (Section 2.4), and so does linear regression, we do not believe our conclusions would change qualitatively if we simulate the data in a nonlinear fashion.

On the real data front, we examine the most recent GTEx eQTL and sQTL data [10] and the $3'$aQTL data prepared by Li et al. [11] from GTEx RNA-seq reads [9]. While the exact data analysis pipelines are different (Table S3), these studies all choose PEER as their hidden variable inference method (due to lack of data availability, we do not examine more real data sets). Our analysis shows that PEER, the most popular hidden variable inference method for QTL mapping currently, produces identical results as PCA at best (Section 2.3), is at least three orders of magnitude slower than PCA (Figure S6), and can be full of pitfalls. Specifically, we show that in certain cases, PEER factors can be highly correlated with each other and thus fail to capture

important variance components of the molecular phenotype data, leading to potential loss of power in QTL identification (Section 2.2). Further, we show from the perspective of PCA that choosing the number of PEER factors by maximizing the number of discoveries (a common approach used by practitioners) may yield inappropriate choices of K, leading to model overfit and potential loss of power and precision (Section 2.5).

Between the two PCA approaches, PCA_direct (running PCA on the fully processed molecular phenotype matrix *directly* and filtering out the known covariates that are captured well by the top PCs afterwards) and PCA_resid (running PCA after regressing out the effects of the known covariates from the molecular phenotype matrix) (Table 1; Section S4), we recommend PCA_direct because the two approaches perform similarly in our simulation studies and PCA_direct is simpler. In addition, PCA_direct can better hedge against the possibility that the known covariates are not actually important confounders because in PCA_direct, the known covariates do not affect the calculation of the PCs. We also advise the users to make sure to center and scale their data when running PCA unless they are experts and have a good reason not to.

To help researchers use PCA in their QTL analysis, we provide an R package `PCAForQTL`, which implements highly interpretable methods for choosing the number of PCs (Algorithms S2 and S3), a graphing function, and more, along with a detailed tutorial. Both resources are freely available at `https://github.com/heatherjzhou/PCAForQTL`. We believe that using PCA rather than SVA, PEER, or HCP will substantially improve and simplify hidden variable inference in QTL mapping as well as increase the transparency and reproducibility of QTL research.

Figures

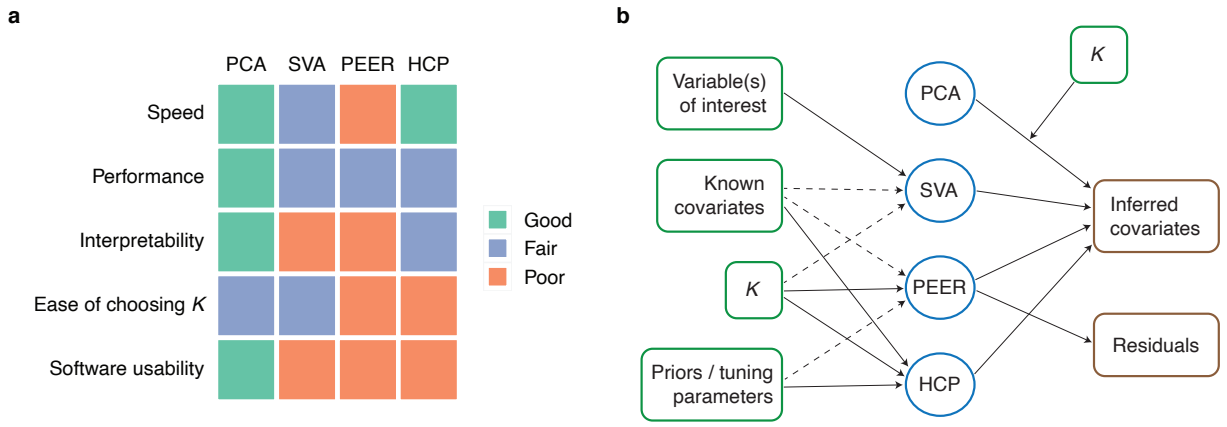

Figure 1: Overall comparison of PCA, SVA, PEER, and HCP and summary of their inputs and outputs. In this work, we use K to denote the number of inferred covariates, which are called PCs, SVs, PEER factors, and HCPs in PCA, SVA, PEER, and HCP, respectively. **a** PCA is better than or as good as all of the other methods in various regards. For speed and performance comparison, see Section 2.1 (and to a lesser extent, Sections 2.2 and 2.3). For interpretability and ease of choosing K, see Sections 2.4 and 2.5, respectively. In terms of software usability, SVA is difficult to apply in QTL settings (Section S4), PEER is difficult to install, and HCP is poorly documented. In addition, PEER suffers from the disadvantage that there is no consensus in the literature on how it should be used (Section S4). **b** Inputs (green boxes) and outputs (brown boxes) of the four methods. The fully processed molecular phenotype matrix (after the effects of the known covariates are regressed out in the case of PCA_resid; Table 1) is a required input for all four methods and is thus omitted in the diagram. Dashed arrows indicate optional inputs. PEER outputs both inferred covariates and residuals of the inputted molecular phenotype matrix [32].

	Inference method	**Method**	**Response, covariates**	**Method abbr. (if selected)**
	(A)	(B)	(C)	(D)
1		**Ideal**	Y, X_1 + X_2	Ideal
2		**Unadjusted**	Y, X_1	Unadjusted
3	PCA_direct	**PCA**_direct_screeK	Y, X_1 (filtered) + top PCs	PCA
4	PCA_resid	**PCA**_resid_screeK	Y, X_1 + top PCs	
5	SVA_trueK	**SVA**_trueK	Y, X_1 + SVs	
6	SVA_BE	**SVA**_BE	Y, X_1 + SVs	SVA
7	PEER_noCov_trueK	**PEER**_noCov_trueK_factors	Y, X_1 (filtered) + PEER factors	
8		**PEER**_noCov_trueK_residuals	Y_{resid}, `NULL`	
9	PEER_noCov_largeK	**PEER**_noCov_largeK_factors	Y, X_1 (filtered) + PEER factors	
10		**PEER**_noCov_largeK_residuals	Y_{resid}, `NULL`	
11	PEER_withCov_trueK	**PEER**_withCov_trueK_factors	Y, X_1 + PEER factors	PEER, true K, factors
12		**PEER**_withCov_trueK_residuals	Y_{resid}, `NULL`	
13	PEER_withCov_largeK	**PEER**_withCov_largeK_factors	Y, X_1 + PEER factors	
14		**PEER**_withCov_largeK_residuals	Y_{resid}, `NULL`	PEER, large K, residuals
15	HCP_trueK	**HCP**_trueK	Y, X_1 + HCPs	HCP

Table 1: Summary of the 15 methods we compare based on simulation studies, including Ideal, Unadjusted, and 13 variants of PCA, SVA, PEER, and HCP (Section S4). Out of the 15 methods, we select a few representative methods (Section 3.2) for detailed comparison in Simulation Design 2, the abbreviations of which are shown in (D). Y denotes the gene expression matrix, Y_{resid} denotes the residual matrix outputted by PEER, X_1 denotes the known covariate matrix, and X_2 denotes the hidden covariate matrix. In Line 3, PCA is run on Y directly; in Line 4, PCA is run after the effects of X_1 are regressed out from Y (Section S4). The addition signs in (C) denote column concatenation. "filtered" means that we filter out the known covariates that are captured well by the inferred covariates (unadjusted $R^2 \geq 0.9$); this filtering is only needed when the hidden variable inference method in (A) does not explicitly take the known covariates into account.

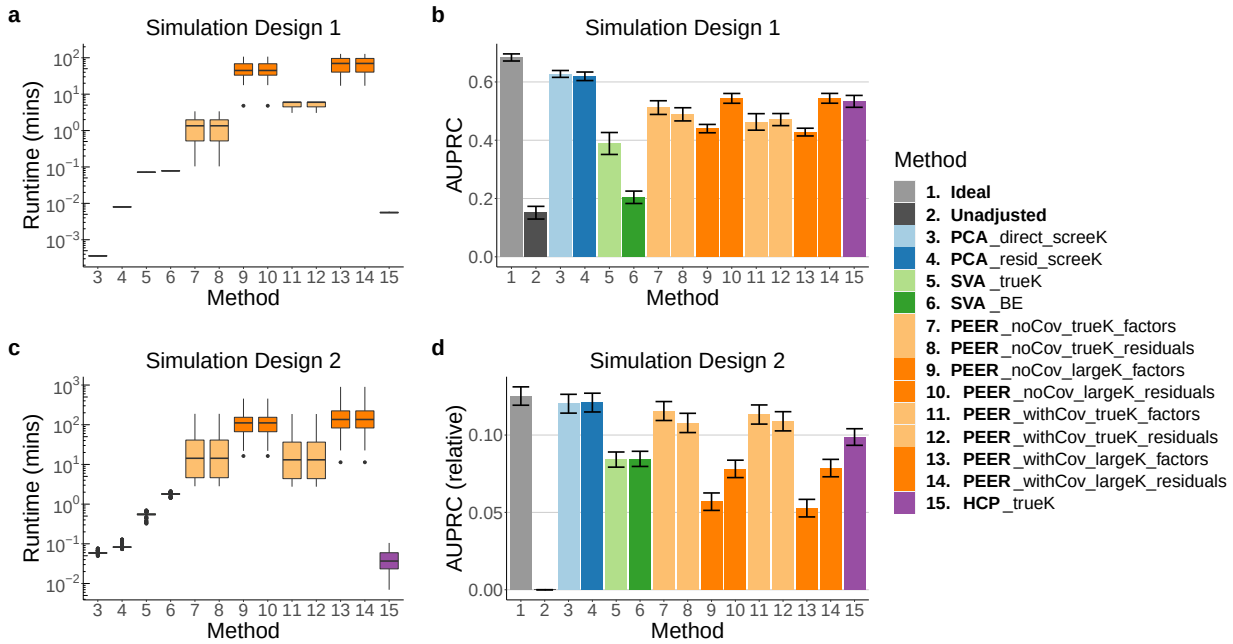

Figure 2: Runtime and AUPRC comparison of all 15 methods (Table 1) in Simulation Design 1 and Simulation Design 2. **a, c** PCA and HCP each takes within a few seconds, SVA takes up to a few minutes, and PEER takes up to about $1,000$ minutes, equivalent to about 17 hours. In particular, PEER takes longer to run when K is large (dark orange vs. light orange boxes). **b, d** PCA outperforms SVA, PEER, and HCP in terms of AUPRC. The height of each bar represents the average across simulated data sets. For ease of visualization, in **d**, the y-axis displays $\left(\text{AUPRC} - \text{AUPRC}_{\text{Unadjusted}}\right)/\text{AUPRC}_{\text{Unadjusted}}$. In this work, error bars indicate standard errors unless otherwise specified (whiskers in box plots are not considered error bars).

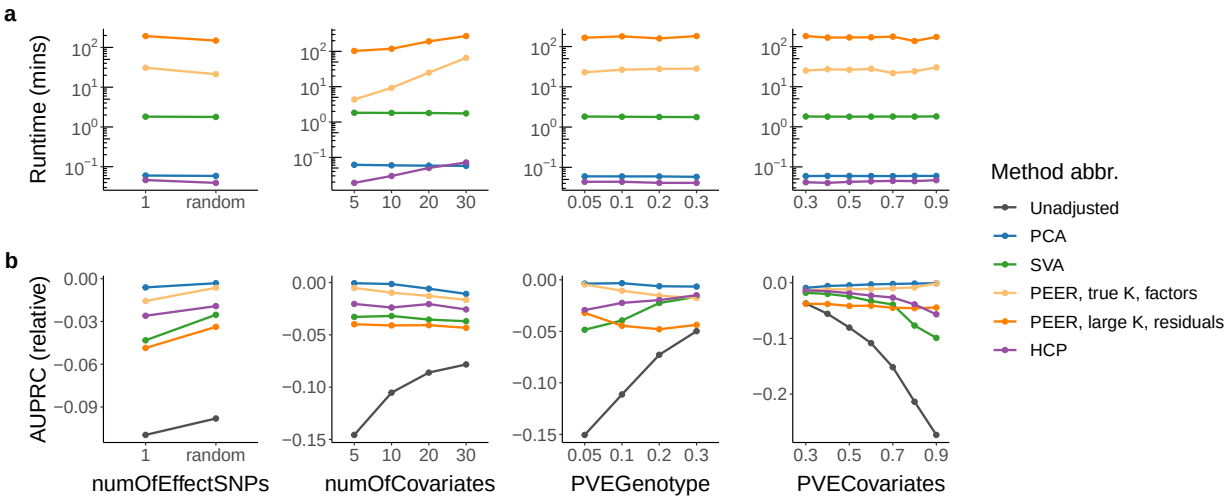

Figure 3: Detailed runtime and AUPRC comparison of the selected representative methods (Table 1) in Simulation Design 2. Each point represents the average across simulated data sets. The x-axes are: number of effect SNPs per gene (`numOfEffectSNPs`), number of covariates (`numOfCovariates`; including known and hidden covariates), proportion of variance explained by genotype (`PVEGenotype`), and proportion of variance explained by covariates (`PVECovariates`) (Section S3). **a** PCA and HCP are orders of magnitude faster than SVA, which in turn is orders of magnitude faster than PEER. **b** PCA outperforms SVA, PEER, and HCP in terms of AUPRC across different simulation settings. For ease of visualization, the y-axis displays $(\text{AUPRC} - \text{AUPRC}_{\text{Ideal}})/\text{AUPRC}_{\text{Ideal}}$. Consistent with our expectation, the difference between Unadjusted and Ideal is the largest (and thus accounting for hidden covariates is the most important) when `numOfCovariates` is small, when `PVEGenotype` is small, and when `PVECovariates` is large.

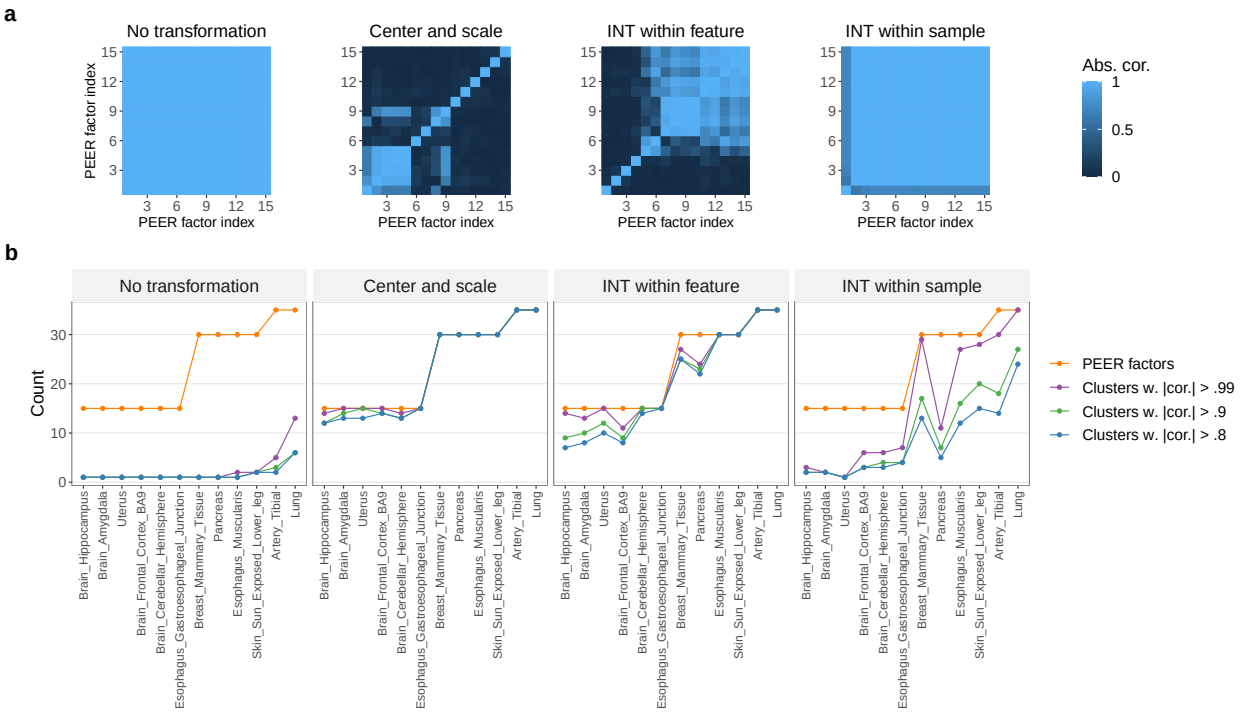

Figure 4: In the 3′aQTL data prepared by Li et al. [11] from GTEx RNA-seq reads [9], PEER factors can be highly correlated with each other to the extent that many or all of them are practically identical. **a** Correlation heatmaps of PEER factors for Brain_Hippocampus. For ease of visualization, the PEER factors are reordered based on results from hierarchical clustering (Section 2.2). **b** The x-axis shows 12 randomly selected tissue types with increasing sample sizes. The y-axis shows the number of PEER factors requested (orange line) or the number of PEER factor clusters. In each cluster, the correlation between any two PEER factors is above 0.99, 0.9, or 0.8 in absolute value. Therefore, the number of PEER factor clusters can be interpreted as the number of distinct or nonrepetitive PEER factors. We find that in many cases, the number of distinct PEER factors is considerably smaller than the number of PEER factors requested, and when this issue is severe (e.g., "No transformation" and "INT within sample"), the PEER factors fail to capture important variance components of the molecular phenotype data (Figure S5).

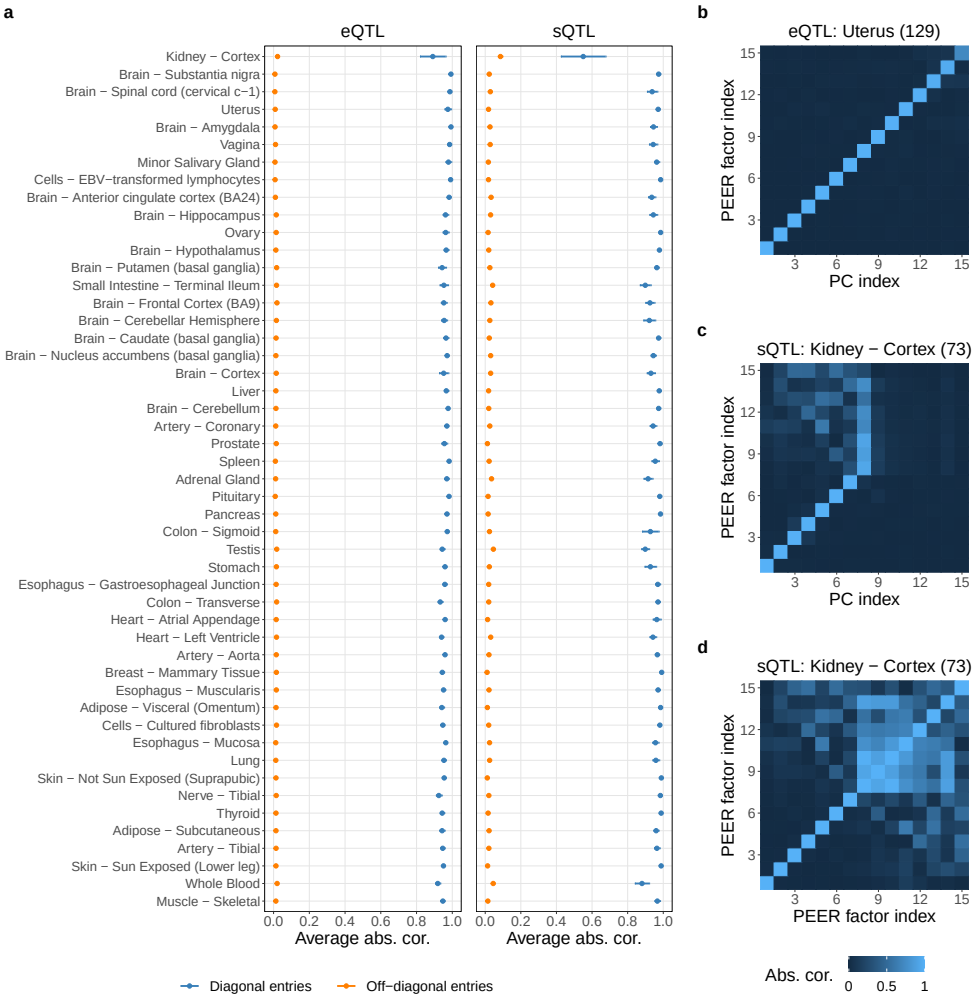

Figure 5: PEER factors are almost identical to PCs in GTEx eQTL and sQTL data [10]. **a** The y-axis shows all 49 tissue types with GTEx QTL analyses ordered by sample size (from small to large). Given a fully processed molecular phenotype matrix, we summarize the correlation matrix (in absolute value) between the PEER factors and the top PCs into two numbers: the average of the diagonal entries and the average of the off-diagonal entries. With the exception of Kidney - Cortex sQTL data, the diagonal entries have averages close to one, and the off-diagonal entries have averages close to zero (both have minimal standard errors). **b** A typical correlation heatmap showing near-perfect one-to-one correspondence between the PEER factors and the top PCs. **c** In Kidney - Cortex sQTL data, the PEER factors and the top PCs do not have a perfect one-to-one correspondence. The reason is because the PEER factors are highly correlated with each other (**d**), while PCs are always uncorrelated (Section S5.1). The numbers in the parentheses represent sample sizes. To produce this figure, we reorder the PEER factors based on the PCs (Algorithm S1), although in almost all cases, this reordering does not change the original ordering of the PEER factors because PEER initializes with PCs [24].

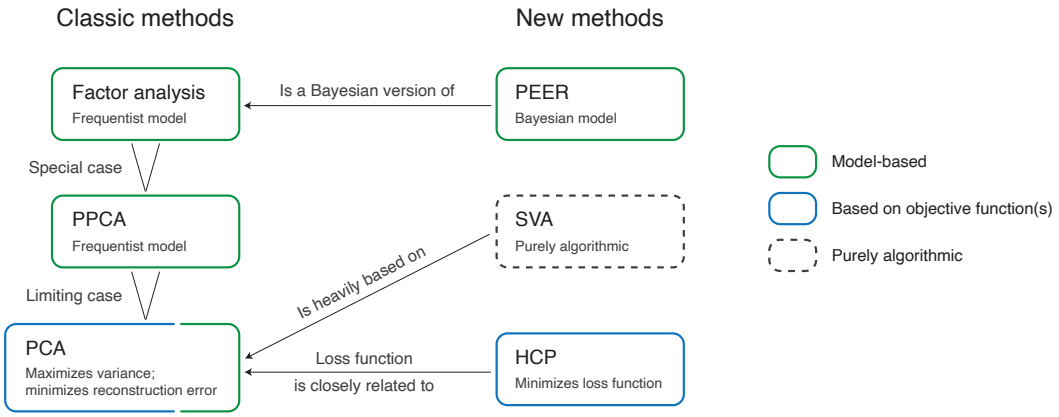

Figure 6: PCA, SVA, PEER, and HCP are closely related statistical methods despite their apparent dissimilarities. In particular, the methodology behind SVA, PEER, and HCP can all be traced back to PCA. PCA [34, 35] is traditionally derived by optimizing some objective functions (either maximum variance or minimum reconstruction error; Section S5.1), but more recently, it is shown [39] that PCA can be derived as a limiting case of probabilistic principal component analysis (PPCA) [39], which in turn is a special case of factor analysis [35, 40]. PEER [24, 32] is based on a Bayesian probabilistic model and can be considered a Bayesian version of factor analysis. SVA [25, 26] is purely an algorithm and is not defined based on a probabilistic model or objective function; though complicated, the SVA algorithm is heavily based on PCA. Lastly, HCP [33] is defined by minimizing a loss function that is very similar to the minimum-reconstruction-error loss function of PCA (Section S5.2).

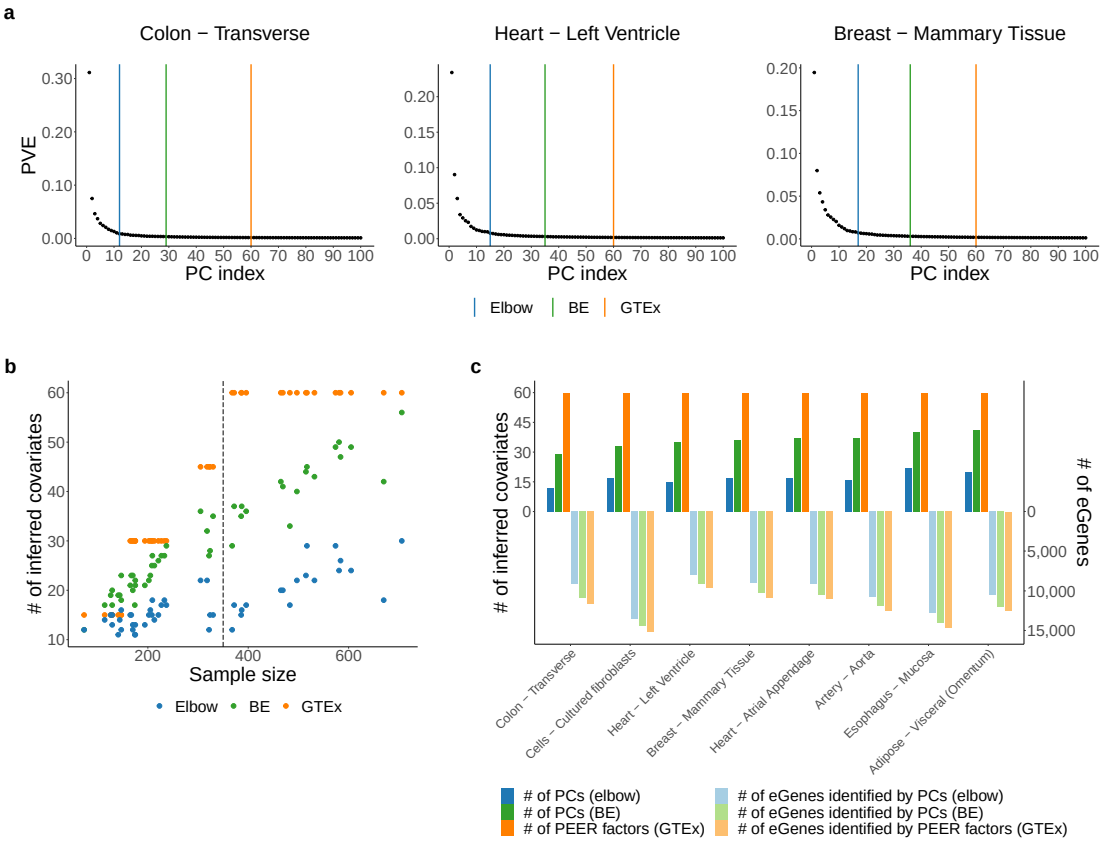

Figure 7: PCA provides insight into the choice of K. Recall from Section 2.3 that PEER factors are almost identical to PCs in GTEx eQTL data [10]. Therefore, for each tissue type, we compare the number of PEER factors selected by GTEx to (1) the number of PCs chosen via an automatic elbow detection method (Algorithm S2) and (2) the number of PCs chosen via the BE algorithm (Algorithm S3; the default parameters are used). **a** Example scree plots. **b** This scatter plot contains 49 dots of each color, corresponding to the 49 tissue types with GTEx eQTL analyses. The number of PEER factors selected by GTEx far exceeds the number of PCs chosen via *BE* for many tissue types with sample size above 350 (dashed line), suggesting that the number of PEER factors selected by GTEx may be too large. **c** For the eight tissue types with the largest absolute differences between the number of PEER factors chosen by GTEx and the number of PCs chosen via *BE* (all eight tissue types have sample size above 350), we replace the PEER factors with smaller numbers of PCs in GTEx's FastQTL pipeline [10, 13] and find that we can reduce the number of inferred covariates to between 20% ($12/60 = 20\%$, Colon - Transverse) and 40% ($22/60 \approx 36.67\%$, Esophagus - Mucosa) of the number of PEER factors selected by GTEx without significantly reducing the number of discovered cis-eGenes.

Availability of data and materials

The R package `PCAForQTL` and a tutorial on using PCA for hidden variable inference in QTL mapping are available at `https://github.com/heatherjzhou/PCAForQTL`. The code used to

generate the results in this work is available at http://doi.org/10.5281/zenodo.6342065. In addition, this work makes use of the following data and software:

- GTEx V8 public data [10], including fully processed gene expression matrices, fully processed splicing phenotype matrices, known covariates, PEER factors, and QTL results, are downloaded from https://gtexportal.org/home/datasets.
- GTEx V8 protected data [10], specifically, the whole genome sequencing (WGS) phased genotype data, are downloaded from the AnVIL repository with an approved dbGaP application (see https://gtexportal.org/home/protectedDataAccess).
- 3′aQTL data prepared by Li et al. [11] from GTEx RNA-seq reads [9] are available from the authors by request.
- SVA R package Version 3.40.0 (https://bioconductor.org/packages/sva/).
- PEER R package (https://bioconda.github.io/recipes/r-peer/README.html).
- HCP R package Version 1.6 (https://rdrr.io/github/mvaniterson/Rhcpp/).
- FastQTL (https://github.com/francois-a/fastqtl).

Competing interests

The authors declare that they have no competing interests.

Funding

This work is supported by the following grants: NSF DGE-1829071 and NIH/NHLBI T32HL139450 (to H.J.Z.); NIH/NCI R01CA193466 and R01CA228140 (to W.L.); NIH/NIGMS R01GM120507 and R35GM140888, NSF DBI-1846216 and DMS-2113754, Johnson & Johnson WiSTEM2D Award, Sloan Research Fellowship, and UCLA David Geffen School of Medicine W.M. Keck Foundation Junior Faculty Award (to J.J.L.).

Authors' contributions

H.J.Z and J.J.L. conceived the project. H.J.Z. performed the experiments and data analyses and wrote the manuscript. L.L. provided the 3'aQTL data. L.L., Y.L., and W.L. served as advisors. J.J.L. supervised the project. All authors participated in discussions and approved the final manuscript.

Acknowledgments

The authors would like to thank former and current members of Junction of Statistics and Biology at UCLA for their valuable insight and suggestions, including Elvis Han Cui, Kexin Li, Dr. Xinzhou Ge, and Dr. Ruochen Jiang.

References

[1] Jacqueline MacArthur, Emily Bowler, Maria Cerezo, Laurent Gil, Peggy Hall, Emma Hastings, Heather Junkins, Aoife McMahon, Annalisa Milano, Joannella Morales, Zoe May Pendlington, Danielle Welter, Tony Burdett, Lucia Hindorff, Paul Flicek, Fiona Cunningham, and Helen Parkinson. The new NHGRI-EBI Catalog of published genome-wide association studies (GWAS Catalog). *Nucleic Acids Research*, 45(D1):D896–D901, 2017.

[2] Annalisa Buniello, Jacqueline A. L. MacArthur, Maria Cerezo, Laura W. Harris, James Hayhurst, Cinzia Malangone, Aoife McMahon, Joannella Morales, Edward Mountjoy, and Elliot Sollis. The NHGRI-EBI GWAS Catalog of published genome-wide association studies, targeted arrays and summary statistics 2019. *Nucleic Acids Research*, 47(D1):D1005–D1012, 2019.

[3] Hongyu Zhao. Roles of statistical modeling in characterizing the genetic basis of human diseases and traits. *Quantitative Biology*, 9(4):371–377, 2021.

[4] Loic Yengo, Julia Sidorenko, Kathryn E. Kemper, Zhili Zheng, Andrew R. Wood, Michael N. Weedon, Timothy M. Frayling, Joel Hirschhorn, Jian Yang, and Peter M. Visscher. Meta-analysis of genome-wide association studies for height and body mass index in ∼700 000 individuals of European ancestry. *Human Molecular Genetics*, 27(20):3641–3649, 2018.

[5] Matthew T. Maurano, Richard Humbert, Eric Rynes, Robert E. Thurman, Eric Haugen, Hao Wang, Alex P. Reynolds, Richard Sandstrom, Hongzhu Qu, Jennifer Brody, Anthony Shafer, Fidencio Neri, Kristen Lee, Tanya Kutyavin, Sandra Stehling-Sun, Audra K. Johnson, Theresa K. Canfield, Erika Giste, Morgan Diegel, Daniel Bates, R. Scott Hansen, Shane Neph, Peter J. Sabo, Shelly Heimfeld, Antony Raubitschek, Steven Ziegler, Chris Cotsapas, Nona Sotoodehnia, Ian Glass, Shamil R. Sunyaev, Rajinder Kaul, and John A. Stamatoyannopoulos. Systematic localization of common disease-associated variation in regulatory DNA. *Science*, 337(6099):1190–1195, 2012.

[6] Edward Mountjoy, Ellen M. Schmidt, Miguel Carmona, Jeremy Schwartzentruber, Gareth Peat, Alfredo Miranda, Luca Fumis, James Hayhurst, Annalisa Buniello, Mohd Anisul Karim, Daniel Wright, Andrew Hercules, Eliseo Papa, Eric B. Fauman, Jeffrey C. Barrett, John A. Todd, David Ochoa, Ian Dunham, and Maya Ghoussaini. An open approach to systematically prioritize causal variants and genes at all published human GWAS trait-associated loci. *Nature Genetics*, 53(11):1527–1533, 2021.

[7] Eddie Cano-Gamez and Gosia Trynka. From GWAS to function: Using functional genomics to identify the mechanisms underlying complex diseases. *Frontiers in Genetics*, 11:424, 2020.

[8] Youqiong Ye, Zhao Zhang, Yaoming Liu, Lixia Diao, and Leng Han. A multi-omics perspective of quantitative trait loci in precision medicine. *Trends in Genetics*, 36(5):318–336, 2020.

[9] GTEx Consortium. Genetic effects on gene expression across human tissues. *Nature*, 550 (7675):204–213, 2017.

[10] GTEx Consortium. The GTEx Consortium atlas of genetic regulatory effects across human tissues. *Science*, 369(6509):1318–1330, 2020.

[11] Lei Li, Kai-Lieh Huang, Yipeng Gao, Ya Cui, Gao Wang, Nathan D. Elrod, Yumei Li, Yiling Elaine Chen, Ping Ji, Fanglue Peng, William K. Russell, Eric J. Wagner, and Wei Li. An atlas of alternative polyadenylation quantitative trait loci contributing to complex trait and disease heritability. *Nature Genetics*, 53(7):994–1005, 2021.

[12] Andrey A. Shabalin. Matrix eQTL: Ultra fast eQTL analysis via large matrix operations. *Bioinformatics*, 28(10):1353–1358, 2012.

[13] Halit Ongen, Alfonso Buil, Andrew Anand Brown, Emmanouil T. Dermitzakis, and Olivier Delaneau. Fast and efficient QTL mapper for thousands of molecular phenotypes. *Bioinformatics*, 32(10):1479–1485, 2016.

[14] The Geuvadis Consortium, Tuuli Lappalainen, Michael Sammeth, Marc R. Friedländer, Peter A. C. 't Hoen, Jean Monlong, Manuel A. Rivas, Mar Gonzàlez-Porta, Natalja Kurbatova, Thasso Griebel, Pedro G. Ferreira, Matthias Barann, Thomas Wieland, Liliana Greger, Maarten van Iterson, Jonas Almlöf, Paolo Ribeca, Irina Pulyakhina, Daniela Esser, Thomas Giger, Andrew Tikhonov, Marc Sultan, Gabrielle Bertier, Daniel G. MacArthur, Monkol Lek, Esther Lizano, Henk P. J. Buermans, Ismael Padioleau, Thomas Schwarzmayr, Olof Karlberg, Halit Ongen, Helena Kilpinen, Sergi Beltran, Marta Gut, Katja Kahlem, Vyacheslav Amstislavskiy, Oliver Stegle, Matti Pirinen, Stephen B. Montgomery, Peter Donnelly, Mark I. McCarthy, Paul Flicek, Tim M. Strom, Hans Lehrach, Stefan Schreiber, Ralf Sudbrak, Ángel Carracedo, Stylianos E. Antonarakis, Robert Häsler, Ann-Christine Syvänen, Gert-Jan van Ommen, Alvis Brazma, Thomas Meitinger, Philip Rosenstiel, Roderic Guigó, Ivo G. Gut, Xavier Estivill, and Emmanouil T. Dermitzakis. Transcriptome and genome sequencing uncovers functional variation in humans. *Nature*, 501(7468):506–511, 2013.

[15] Xianjun Dong, Zhixiang Liao, David Gritsch, Yavor Hadzhiev, Yunfei Bai, Joseph J. Locascio, Boris Guennewig, Ganqiang Liu, Cornelis Blauwendraat, Tao Wang, Charles H. Adler, John C. Hedreen, Richard L. M. Faull, Matthew P. Frosch, Peter T. Nelson, Patrizia Rizzu, Antony A. Cooper, Peter Heutink, Thomas G. Beach, John S. Mattick, Ferenc Müller, and Clemens R. Scherzer. Enhancers active in dopamine neurons are a primary link between genetic variation and neuropsychiatric disease. *Nature Neuroscience*, 21(10):1482–1492, 2018.

[16] Christopher E. Gillies, Rosemary Putler, Rajasree Menon, Edgar Otto, Kalyn Yasutake, Viji Nair, Paul Hoover, David Lieb, Shuqiang Li, Sean Eddy, Damian Fermin, Michelle T. McNulty, Nir Hacohen, Krzysztof Kiryluk, Matthias Kretzler, Xiaoquan Wen, Matthew G. Sampson, John Sedor, Katherine Dell, Marleen Schachere, Kevin Lemley, Lauren Whitted, Tarak Srivastava, Connie Haney, Christine Sethna, Kalliopi Grammatikopoulos, Gerald Appel, Michael Toledo, Laurence Greenbaum, Chia-shi Wang, Brian Lee, Sharon Adler, Cynthia Nast, Janine LaPage, Ambarish Athavale, Alicia Neu, Sara Boynton, Fernando Fervenza, Marie Hogan, John C. Lieske, Vladimir Chernitskiy, Frederick Kaskel, Neelja Kumar, Patricia Flynn, Jeffrey Kopp, Eveleyn Castro-Rubio, Jodi Blake, Howard Trachtman, Olga Zhdanova, Frank Modersitzki, Suzanne Vento, Richard Lafayette, Kshama Mehta,

Crystal Gadegbeku, Duncan Johnstone, Daniel Cattran, Michelle Hladunewich, Heather Reich, Paul Ling, Martin Romano, Alessia Fornoni, Laura Barisoni, Carlos Bidot, Matthias Kretzler, Debbie Gipson, Amanda Williams, Renee Pitter, Patrick Nachman, Keisha Gibson, Sandra Grubbs, Anne Froment, Lawrence Holzman, Kevin Meyers, Krishna Kallem, Fumei Cerecino, Kamal Sambandam, Elizabeth Brown, Natalie Johnson, Ashley Jefferson, Sangeeta Hingorani, Kathleen Tuttle, Laura Curtin, S. Dismuke, Ann Cooper, Barry Freedman, Jen Jar Lin, Stefanie Gray, Matthias Kretzler, Larua Barisoni, Crystal Gadegbeku, Brenda Gillespie, Debbie Gipson, Lawrence Holzman, Laura Mariani, Matthew G. Sampson, Peter Song, Johnathan Troost, Jarcy Zee, Emily Herreshoff, Colleen Kincaid, Chrysta Lienczewski, Tina Mainieri, Amanda Williams, Kevin Abbott, Cindy Roy, Tiina Urv, and John Brooks. An eQTL landscape of kidney tissue in human nephrotic syndrome. *The American Journal of Human Genetics*, 103(2):232–244, 2018.

[17] Satria P. Sajuthi, Peter DeFord, Yingchun Li, Nathan D. Jackson, Michael T. Montgomery, Jamie L. Everman, Cydney L. Rios, Elmar Pruesse, James D. Nolin, Elizabeth G. Plender, Michael E. Wechsler, Angel C. Y. Mak, Celeste Eng, Sandra Salazar, Vivian Medina, Eric M. Wohlford, Scott Huntsman, Deborah A. Nickerson, Soren Germer, Michael C. Zody, Gonçalo Abecasis, Hyun Min Kang, Kenneth M. Rice, Rajesh Kumar, Sam Oh, Jose Rodriguez-Santana, Esteban G. Burchard, and Max A. Seibold. Type 2 and interferon inflammation regulate SARS-CoV-2 entry factor expression in the airway epithelium. *Nature Communications*, 11(1):5139, 2020.

[18] Luz D. Orozco, Hsu-Hsin Chen, Christian Cox, Kenneth J. Katschke, Rommel Arceo, Carmina Espiritu, Patrick Caplazi, Sarajane Saturnio Nghiem, Ying-Jiun Chen, Zora Modrusan, Amy Dressen, Leonard D. Goldstein, Christine Clarke, Tushar Bhangale, Brian Yaspan, Marion Jeanne, Michael J. Townsend, Menno van Lookeren Campagne, and Jason A. Hackney. Integration of eQTL and a single-cell atlas in the human eye identifies causal genes for age-related macular degeneration. *Cell Reports*, 30(4):1246–1259.e6, 2020.

[19] Jing Gong, Shufang Mei, Chunjie Liu, Yu Xiang, Youqiong Ye, Zhao Zhang, Jing Feng, Renyan Liu, Lixia Diao, An-Yuan Guo, Xiaoping Miao, and Leng Han. PancanQTL: Systematic identification of cis-eQTLs and trans-eQTLs in 33 cancer types. *Nucleic Acids Research*, 46(D1):D971–D976, 2018.

[20] Anubha Mahajan, Daniel Taliun, Matthias Thurner, Neil R. Robertson, Jason M. Torres, N. William Rayner, Anthony J. Payne, Valgerdur Steinthorsdottir, Robert A. Scott, Niels Grarup, James P. Cook, Ellen M. Schmidt, Matthias Wuttke, Chloé Sarnowski, Reedik Mägi, Jana Nano, Christian Gieger, Stella Trompet, Cécile Lecoeur, Michael H. Preuss, Bram Peter Prins, Xiuqing Guo, Lawrence F. Bielak, Jennifer E. Below, Donald W. Bowden, John Campbell Chambers, Young Jin Kim, Maggie C. Y. Ng, Lauren E. Petty, Xueling Sim, Weihua Zhang, Amanda J. Bennett, Jette Bork-Jensen, Chad M. Brummett, Mickaël Canouil, Kai-Uwe Ec kardt, Krista Fischer, Sharon L. R. Kardia, Florian Kronenberg, Kristi Läll, Ching-Ti Liu, Adam E. Locke, Jian'an Luan, Ioanna Ntalla, Vibe Nylander, Sebastian Schönherr, Claudia Schurmann, Loïc Yengo, Erwin P. Bottinger, Ivan Brandslund, Cramer Christensen, George Dedoussis, Jose C. Florez, Ian Ford, Oscar H. Franco, Timothy M. Frayling, Vilmantas Giedraitis, Sophie Hackinger, Andrew T.

Hattersley, Christian Herder, M. Arfan Ikram, Martin Ingelsson, Marit E. Jørgensen, Torben Jørgensen, Jennifer Kriebel, Johanna Kuusisto, Symen Ligthart, Cecilia M. Lindgren, Allan Linneberg, Valeriya Lyssenko, Vasiliki Mamakou, Thomas Meitinger, Karen L. Mohlke, Andrew D. Morris, Girish Nadkarni, James S. Pankow, Annette Peters, Naveed Sattar, Alena Stančáková, Konstantin Strauch, Kent D. Taylor, Barbara Thorand, Gudmar Thorleifsson, Unnur Thorsteinsdottir, Jaakko Tuomilehto, Daniel R. Witte, Josée Dupuis, Patricia A. Peyser, Eleftheria Zeggini, Ruth J. F. Loos, Philippe Froguel, Erik Ingelsson, Lars Lind, Leif Groop, Markku Laakso, Francis S. Collins, J. Wouter Jukema, Colin N. A. Palmer, Harald Grallert, Andres Metspalu, Abbas Dehghan, Anna Köttgen, Goncalo R. Abecasis, James B. Meigs, Jerome I. Rotter, Jonathan Marchini, Oluf Pedersen, Torben Hansen, Claudia Langenberg, Nicholas J. Wareham, Kari Stefansson, Anna L. Gloyn, Andrew P. Morris, Michael Boehnke, and Mark I. McCarthy. Fine-mapping type 2 diabetes loci to single-variant resolution using high-density imputation and islet-specific epigenome maps. *Nature Genetics*, 50(11):1505–1513, 2018.

[21] Stephanie Feupe Fotsing, Jonathan Margoliash, Catherine Wang, Shubham Saini, Richard Yanicky, Sharona Shleizer-Burko, Alon Goren, and Melissa Gymrek. The impact of short tandem repeat variation on gene expression. *Nature Genetics*, 51(11):1652–1659, 2019.

[22] Rebecca L. Walker, Gokul Ramaswami, Christopher Hartl, Nicholas Mancuso, Michael J. Gandal, Luis de la Torre-Ubieta, Bogdan Pasaniuc, Jason L. Stein, and Daniel H. Geschwind. Genetic control of expression and splicing in developing human brain informs disease mechanisms. *Cell*, 179(3):750–771, 2019.

[23] Alexis Battle, Sara Mostafavi, Xiaowei Zhu, James B. Potash, Myrna M. Weissman, Courtney McCormick, Christian D. Haudenschild, Kenneth B. Beckman, Jianxin Shi, Rui Mei, Alexander E. Urban, Stephen B. Montgomery, Douglas F. Levinson, and Daphne Koller. Characterizing the genetic basis of transcriptome diversity through RNA-sequencing of 922 individuals. *Genome Research*, 24(1):14–24, 2014.

[24] Oliver Stegle, Leopold Parts, Richard Durbin, and John Winn. A Bayesian framework to account for complex non-genetic factors in gene expression levels greatly increases power in eQTL studies. *PLoS Computational Biology*, 6(5):e1000770, 2010.

[25] Jeffrey T. Leek and John D. Storey. Capturing heterogeneity in gene expression studies by surrogate variable analysis. *PLoS Genetics*, 3(9):e161, 2007.

[26] Jefferey T. Leek and John D. Storey. A general framework for multiple testing dependence. *Proceedings of the National Academy of Sciences*, 105(48):18718–18723, 2008.

[27] James C. Cronk, Anthony J. Filiano, Antoine Louveau, Ioana Marin, Rachel Marsh, Emily Ji, Dylan H. Goldman, Igor Smirnov, Nicholas Geraci, Scott Acton, Christopher C. Overall, and Jonathan Kipnis. Peripherally derived macrophages can engraft the brain independent of irradiation and maintain an identity distinct from microglia. *Journal of Experimental Medicine*, 215(6):1627–1647, 2018.

[28] Jeffrey W. Tyner, Cristina E. Tognon, Daniel Bottomly, Beth Wilmot, Stephen E. Kurtz, Samantha L. Savage, Nicola Long, Anna Reister Schultz, Elie Traer, Melissa Abel, Anupriya

Agarwal, Aurora Blucher, Uma Borate, Jade Bryant, Russell Burke, Amy Carlos, Richie Carpenter, Joseph Carroll, Bill H. Chang, Cody Coblentz, Amanda d'Almeida, Rachel Cook, Alexey Danilov, Kim-Hien T. Dao, Michie Degnin, Deirdre Devine, James Dibb, David K. Edwards, Christopher A. Eide, Isabel English, Jason Glover, Rachel Henson, Hibery Ho, Abdusebur Jemal, Kara Johnson, Ryan Johnson, Brian Junio, Andy Kaempf, Jessica Leonard, Chenwei Lin, Selina Qiuying Liu, Pierrette Lo, Marc M. Loriaux, Samuel Luty, Tara Macey, Jason MacManiman, Jacqueline Martinez, Motomi Mori, Dylan Nelson, Ceilidh Nichols, Jill Peters, Justin Ramsdill, Angela Rofelty, Robert Schuff, Robert Searles, Erik Segerdell, Rebecca L. Smith, Stephen E. Spurgeon, Tyler Sweeney, Aashis Thapa, Corinne Visser, Jake Wagner, Kevin Watanabe-Smith, Kristen Werth, Joelle Wolf, Libbey White, Amy Yates, Haijiao Zhang, Christopher R. Cogle, Robert H. Collins, Denise C. Connolly, Michael W. Deininger, Leylah Drusbosky, Christopher S. Hourigan, Craig T. Jordan, Patricia Kropf, Tara L. Lin, Micaela E. Martinez, Bruno C. Medeiros, Rachel R. Pallapati, Daniel A. Pollyea, Ronan T. Swords, Justin M. Watts, Scott J. Weir, David L. Wiest, Ryan M. Winters, Shannon K. McWeeney, and Brian J. Druker. Functional genomic landscape of acute myeloid leukaemia. *Nature*, 562(7728):526–531, 2018.

[29] Lindsay F. Rizzardi, Peter F. Hickey, Varenka Rodriguez DiBlasi, Rakel Tryggvadóttir, Colin M. Callahan, Adrian Idrizi, Kasper D. Hansen, and Andrew P. Feinberg. Neuronal brain-region-specific DNA methylation and chromatin accessibility are associated with neuropsychiatric trait heritability. *Nature Neuroscience*, 22(2):307–316, 2019.

[30] GTEx Consortium, Taru Tukiainen, Alexandra-Chloé Villani, Angela Yen, Manuel A. Rivas, Jamie L. Marshall, Rahul Satija, Matt Aguirre, Laura Gauthier, Mark Fleharty, Andrew Kirby, Beryl B. Cummings, Stephane E. Castel, Konrad J. Karczewski, François Aguet, Andrea Byrnes, Tuuli Lappalainen, Aviv Regev, Kristin G. Ardlie, Nir Hacohen, and Daniel G. MacArthur. Landscape of X chromosome inactivation across human tissues. *Nature*, 550 (7675):244–248, 2017.

[31] Jeffrey T. Leek, W. Evan Johnson, Hilary S. Parker, Andrew E. Jaffe, and John D. Storey. The sva package for removing batch effects and other unwanted variation in high-throughput experiments. *Bioinformatics*, 28(6):882–883, 2012.

[32] Oliver Stegle, Leopold Parts, Matias Piipari, John Winn, and Richard Durbin. Using probabilistic estimation of expression residuals (PEER) to obtain increased power and interpretability of gene expression analyses. *Nature Protocols*, 7(3):500–507, 2012.

[33] Sara Mostafavi, Alexis Battle, Xiaowei Zhu, Alexander E. Urban, Douglas Levinson, Stephen B. Montgomery, and Daphne Koller. Normalizing RNA-sequencing data by modeling hidden covariates with prior knowledge. *PLoS ONE*, 8(7):e68141, 2013.

[34] Ian T. Jolliffe. *Principal Component Analysis*. Springer, New York, second edition, 2002.

[35] Richard A. Johnson and Dean W. Wichern. *Applied Multivariate Statistical Analysis*. Pearson Prentice Hall, Upper Saddle River, NJ, sixth edition, 2007.

[36] Gao Wang, Abhishek Sarkar, Peter Carbonetto, and Matthew Stephens. A simple new approach to variable selection in regression, with application to genetic fine mapping. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82(5):1273–1300, 2020.

[37] T. Mark Beasley, Stephen Erickson, and David B. Allison. Rank-based inverse normal transformations are increasingly used, but are they merited? *Behavior Genetics*, 39(5):580–595, 2009.

[38] Stephen C. Johnson. Hierarchical clustering schemes. *Psychometrika*, 32(3):241–254, 1967.

[39] Michael E. Tipping and Christopher M. Bishop. Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B*, 61(3):611–622, 1999.

[40] K. V. Mardia, J. T. Kent, and J. M. Bibby. *Multivariate Analysis*. Academic Press, London, 1979.

[41] Andreas Buja and Nermin Eyuboglu. Remarks on parallel analysis. *Multivariate Behavioral Research*, 27(4):509–540, 1992.

[42] Otto Bretscher. *Linear Algebra With Applications*. Pearson Prentice Hall, Upper Saddle River, NJ, fourth edition, 2009.

[43] Ian T. Jolliffe and Jorge Cadima. Principal component analysis: A review and recent developments. *Philosophical Transactions of the Royal Society A*, 374(2065), 2016.

[44] Harold Hotelling. Analysis of a complex of statistical variables into principal components. *Journal of Educational Psychology*, 24(6):417–441, 1933.

[45] Karl Pearson. LIII. On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):559–572, 1901.

Supplementary materials for

PCA outperforms popular hidden variable inference methods for QTL mapping

Heather J. Zhou, Lei Li, Yumei Li, Wei Li, Jingyi Jessica Li

S1 Limitations of the original PEER simulation study

The simulation study in the original PEER publication [24] is limited. We categorize its limitations into three categories: (1) data analysis limitations, (2) overall design limitations, and (3) data simulation limitations.

The data analysis limitations include:

(a) The study only compares PEER against the other methods in terms of power, not in terms of false positive rate or false discovery rate (see Section 3.1 for our evaluation metrics).
(b) The study does not use PCA or SVA properly (we do; Section S4).
(c) The study does not evaluate the different ways of using PEER (we do; Section S4).
(d) The study uses ad hoc priors for PEER that are different from the default priors (we use the default priors; Section S4).

The overall design limitations include:

(a) The study only simulates one replicate of one experiment. That is, the entire simulation study is based on one simulated data set (we simulate 10 replicates in our first simulation study; Section S2).

The data simulation limitations include (see Table S1 for our solutions):

(a) The data dimensions are minimal, with $q = 100$ SNPs in the entire genome.
(b) The SNP genotypes are simulated independently and identically with a target minor allele frequency (MAF) of 0.4, so there is no linkage disequilibrium (LD) and a higher average MAF than in real data (the average MAF in GTEx data [10], after SNPs with MAF under 0.01 are filtered out, is about 0.15; Section S3.1).
(c) The gene expression levels are primarily driven by trans-regulatory effects rather than cis-regulatory effects or covariate effects (Table S2), inconsistent with the common belief that trans-regulatory effects are generally weaker than cis-regulatory effects.

In addition, the simulation study in the original PEER publication [24] is imperfect in that the description of the data simulation and analysis is vague and inconsistent, and there is no reproducible code. In contrast, we describe our data simulation and analysis in detail (Sections S2 to S4) and provide the code we use to generate the results (see Availability of data and materials).

S2 Simulation Design 1

S2.1 Data simulation

In Simulation Design 1, we perform 10 replicates of the same experiment, where in each replicate, we follow the data simulation in Stegle et al. [24] as closely as possible.

In each replicate, we simulate a data set with $n = 80$ individuals, $p = 400$ genes, $q = 100$ SNPs in the entire genome, $K_1 = 3$ known covariates, and $K_2 = 7$ hidden covariates. Let $i,\ j,\ l$, and k be the indices of individuals, genes, SNPs, and covariates, respectively. That is, $i = 1, \cdots, n$; $j = 1, \cdots, p$; $l = 1, \cdots, q$; and $k = 1, \cdots, (K_1 + K_2)$. The data simulation consists of three steps.

In the first step, we simulate $Y_{\text{beforeDSE}}$, the gene expression matrix before downstream effect, based on

$$\underset{n\times p}{Y_{\text{beforeDSE}}} = \underset{n\times q}{S} \left(\underset{q\times p}{I_1} \odot \underset{q\times p}{B_1} \right) + \underset{n\times(K_1+K_2)}{X} \underset{(K_1+K_2)\times p}{B_2} + \underset{n\times p}{E} \ , \tag{S1}$$

where $\odot$ denotes element-wise multiplication. Specifically, in the genotype component, we have

- S: genotype matrix. Each entry is drawn independently from $Binom\,(2, \text{prob} = 0.4)$. That is, the target MAF is 0.4. In this work, all random sampling is independent unless otherwise specified.
- I_1: effect indicator matrix. Each entry is drawn from $Ber\,(0.01)$.
- B_1: effect size matrix. Each entry is drawn from $N(0, \text{var} = 4)$.

In the covariate component, we have

- X: covariate matrix. Each entry is drawn from $N(0, \text{var} = 0.6)$. The first K_1 columns are designated as the known covariates (X_1, $n \times K_1$), and the last K_2 columns are designated as the hidden covariates (X_2, $n \times K_2$).
- B_2: effect size matrix. First, we draw $\sigma_k^2 \sim 0.8\,(\Gamma\,(\text{shape} = 2.5, \text{rate} = 0.6))^2$, the covariate-specific effect size variance. Then, we draw $(B_2)_{kj} \sim N\left(0, \text{var} = \sigma_k^2\right)$.

Lastly, in the noise component, we have

- E: noise matrix. First, we draw $\tau_j \sim \Gamma\,(\text{shape} = 3, \text{rate} = 1)$, the gene-specific noise precision. Then, we draw $(E)_{ij} \sim N\left(0, \text{var} = 1/\tau_j\right)$.

In the second step, we simulate Y_{DSE}, the gene expression matrix due to the downstream effect of genes, based on

$$\underset{n\times p}{Y_{\text{DSE}}} = \underset{n\times p}{Y_{\text{beforeDSE}}} \left(\underset{p\times p}{I_3} \odot \underset{p\times p}{B_3} \right) \ , \tag{S2}$$

where we have

- I_3: effect indicator matrix. To simulate I_3, we start with a zero matrix. Then, we randomly choose three rows corresponding to genes with at least one cis-QTL (Section S2.2). For each of these three rows, we randomly assign 30 entries corresponding to genes other than the current gene in consideration (avoiding self-loops) to be one.
- B_3: effect size matrix. Each entry is drawn from $N(8, \text{var} = 0.8)$ for "strong downstream effects" [24].

As we see in Section S2.2, the downstream effect of genes induces trans-QTL relations.

In the third and last step, we define Y, the final, observed gene expression matrix, as

$$\underset{n\times p}{Y} = \underset{n\times p}{Y_{\text{beforeDSE}}} + \underset{n\times p}{Y_{\text{DSE}}} \,. \tag{S3}$$

S2.2 Definition of truth

In a simulated data set, the cis-QTL relations are encoded in the $q \times p$ binary matrix I_1. The lj-th entry being one means that SNP l and gene j form a cis-QTL pair (i.e., SNP l is a cis-QTL for gene j).

The trans-QTL relations are encoded in J, also a $q \times p$ binary matrix. J is defined based on I_1 and I_3. Specifically, SNP l and gene j form a trans-QTL pair if and only if SNP l is a cis-QTL for gene j' *and* gene j' has downstream effect on gene j, $j' \neq j$.

The overall truth is encoded in $\mathbb{1}\big((I_1 + J) \geq 1\big)$, again a $q \times p$ binary matrix. We use this matrix as the truth when calculating AUPRCs. The lj-th entry being one means that SNP l and gene j form a cis-QTL or trans-QTL pair (or both).

	Simulation Design 1	**Simulation Design 2**
Data simulation	Follows Stegle et al. [24]	Loosely based on Wang et al. [36]
# of experiments	1	176
# of replicates per experiment	10	2
# of simulated data sets	10	352
Genotype data	Simulated (no LD, high MAF)	Real genotype data from GTEx [10]
Cis-QTL relations present	✓	✓
Trans-QTL relations present	✓	✗
Source(s) of expression variation	Primarily trans-regulatory effects	Carefully controlled genotype effects and covariate effects
# of individuals	$n = 80$	$n = 838$
# of genes	$p = 400$	$p = 1,000$
# of SNPs	$q = 100$ SNPs in the entire genome	$q = 1,000$ local common SNPs per gene
# of known covariates	$K_1 = 3$	$K_1 = 2,\ 3,\ 5,$ or 8 depending on the experiment
# of hidden covariates	$K_2 = 7$	$K_2 = 3,\ 7,\ 15,$ or 22 depending on the experiment

Table S1: Summary of the main differences between Simulation Design 1 and Simulation Design 2. Highlighted in blue are the major data simulation limitations (Section S1) of Simulation Design 1, all of which we address in Simulation Design 2.

Replicate	$\mathrm{Var}(Y_{\mathrm{beforeDSE}})$	$\mathrm{Var}(Y_{\mathrm{DSE}})$	$\mathrm{Var}(Y)$	$\mathrm{Var}(Y_{\mathrm{beforeDSE}}) / \mathrm{Var}(Y)$	$\mathrm{Var}(Y_{\mathrm{DSE}}) / \mathrm{Var}(Y)$
1	124.34	1757.07	1889.57	6.58%	92.99%
2	140.92	1505.17	1677.64	8.40%	89.72%
3	213.56	929.96	1169.18	18.27%	79.54%
4	71.85	761.01	855.39	8.40%	88.97%
5	123.07	2434.45	2574.51	4.78%	94.56%
6	74.94	1029.29	1092.65	6.86%	94.20%
7	148.61	2490.72	2628.93	5.65%	94.74%
8	79.36	796.55	868.05	9.14%	91.76%
9	54.62	1340.10	1390.72	3.93%	96.36%
10	65.64	831.89	895.90	7.33%	92.86%
Average				7.93%	**91.57%**

Table S2: In Simulation Design 1, which follows the data simulation in Stegle et al. [24] as closely as possible, the gene expression levels are primarily driven by trans-regulatory effects rather than cis-regulatory effects or covariate effects. $\mathrm{Var}(Y_{\mathrm{beforeDSE}})$ is defined as the variance of the $n \times p$ entries of $Y_{\mathrm{beforeDSE}}$, and the other variances in the table are defined similarly. We find that $\mathrm{Var}(Y_{\mathrm{DSE}}) / \mathrm{Var}(Y)$ is above 90% on average.

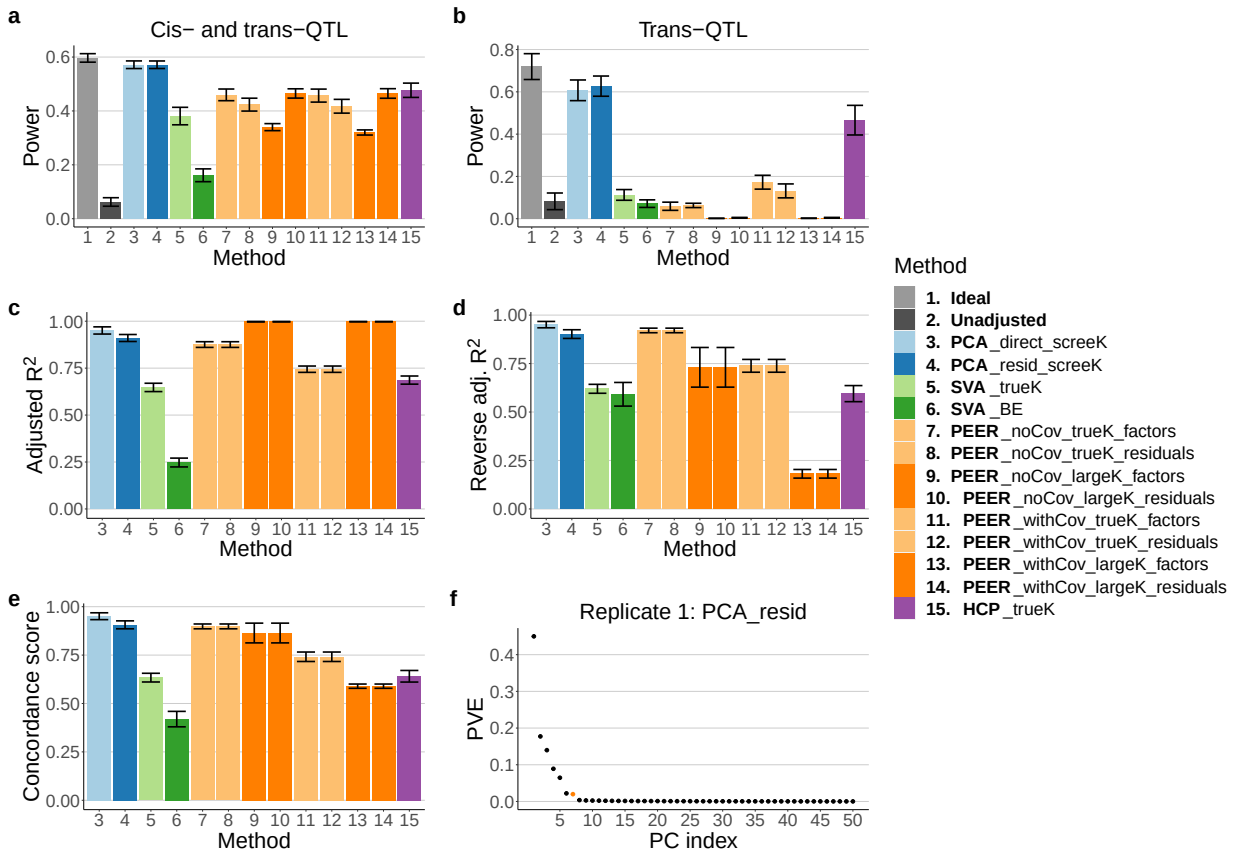

Figure S1: Comparison of all 15 methods (Table 1) in terms of power and adjusted R^2 measures in Simulation Design 1 (the height of each bar represents the average across simulated data sets) and an example scree plot. **a, b** PCA is more powerful than SVA, PEER, and HCP both when we consider all QTL relations (**a**) and when we focus on trans-QTL relations (**b**). Binary decisions are made based on p-values using the Benjamini-Hochberg (BH) procedure and a target false discovery rate of 0.05. **c, d, e** PCA performs the best in terms of concordance score. PEER with a large K (dark orange bars) performs well in terms of adjusted R^2 but less well in terms of reverse adjusted R^2. **f** An example scree plot that unambiguously suggests the true number of hidden covariates, seven in this case, as the reasonable number of PCs to choose (the y-axis represents the proportion of variance explained).

S3 Simulation Design 2

S3.1 Data simulation

In Simulation Design 2, we use real genotype data from GTEx [10], focus on cis-QTL detection, and carefully control the genotype effects and covariate effects in 176 experiments with two replicates per experiment. This simulation design takes inspiration from and is loosely based on Wang et al. [36].

In each experiment-replicate combination, we simulate a data set with $n = 838$ individuals, $p = 1,000$ genes, $q = 1,000$ local common SNPs per gene, K_1 known covariates, and K_2 hidden covariates (the values of K_1 and K_2 depend on the experiment; see below). Again, let i, j, l, and k be the indices of individuals, genes, SNPs, and covariates, respectively. That is, $i = 1, \cdots, n$; $j = 1, \cdots, p$; $l = 1, \cdots, q$; and $k = 1, \cdots, (K_1 + K_2)$.

We begin by obtaining *SArray*, the $n \times q \times p$ genotype array that remains constant throughout Simulation Design 2. $SArray[\,,\,,j]$, an $n \times q$ matrix, is the genotype matrix for the q local common SNPs for gene j. We obtain *SArray* with the following steps:

(a) Download the whole genome sequencing (WGS) phased genotype data for $n = 838$ individuals from GTEx V8 [10].
(b) Randomly select $p = 1,000$ genes from the more than 20,000 genes on chromosomes 1 to 22.
(c) For each gene, obtain the genotype data for the $q = 1,000$ SNPs with MAF≥ 0.01 that are the closest to the gene's transcription start site (TSS); we find that these SNPs are almost always within 1 Mb of the TSS. The average MAF of *SArray*, calculated as $\overline{SArray}/2$, the average of all entries of *SArray* divided by 2, is $0.1474385 \approx 0.15$.

Each experiment is characterized by four attributes:

(a) Number of effect SNPs per gene (`numOfEffectSNPs`): 1 or `random`.
(b) Number of covariates (`numOfCovariates`): 5, 10, 20, or 30.
    - Number of known covariates (K_1): 2, 3, 5, 8, respectively.
    - Number of hidden covariates (K_2): 3, 7, 15, 22, respectively.
(c) Proportion of variance explained by genotype (`PVEGenotype`): 0.05, 0.1, 0.2, or 0.3.
(d) Proportion of variance explained by covariates (`PVECovariates`): minimum 0.3, maximum $1 - 0.05 -$ `PVEGenotype`, in increments of 0.1. For example, when `PVEGenotype` $= 0.05$, `PVECovariates` takes seven possible values: 0.3, 0.4, 0.5, $\cdots$, 0.9.

Therefore, we have a total of $2 \times 4 \times (7+6+5+4) = 8 \times 22 = 176$ experiments covering typical scenarios in QTL studies [36]. Following Wang et al. [36], we use the term "effect SNPs" to refer to SNPs that have a nonzero cis effect on a given gene.

Given `numOfEffectSNPs`, `numOfCovariates`, `PVEGenotype`, and `PVECovariates`, we simulate each data set based on

$$\underset{n\times p}{Y} = \underset{n\times q\times p}{SArray} \otimes \left(\underset{q\times p}{I} \odot \underset{q\times p}{B_1} \right) + \underset{n\times(K_1+K_2)}{X} \underset{(K_1+K_2)\times p}{B_2} + \underset{n\times p}{E}\,, \tag{S4}$$

where Y is the gene expression matrix, and $\otimes$ is defined as

$$\underset{n\times p}{C} = \underset{n\times q\times p}{A} \otimes \underset{q\times p}{B} \quad \Leftrightarrow \quad \underset{n\times 1}{C[\,,j]} = \underset{n\times q}{A[\,,\,,j]} \times \underset{q\times 1}{B[\,,j]}\,,\ j = 1, \cdots, p\,. \tag{S5}$$

Specifically, in the genotype component, we have

- *SArray*: genotype array. $SArray[\,,\,,j]$, an $n \times q$ matrix, is the genotype matrix for the q local common SNPs for gene j (see above).
- I: effect indicator matrix.
  - If `numOfEffectSNPs` $= 1$, then for each column, we randomly assign one entry to be one while keeping the other entries zero.
  - If `numOfEffectSNPs` $=$ `random`, then each entry of I is drawn from $Ber(1/q)$. This means that for each gene, the number of effect SNPs is drawn from $Binom(q, \text{prob} = 1/q)$. This binomial distribution approximates the empirical distribution of the number of independent cis-eQTLs per gene in GTEx data [10] well (Figure S2).
- B_1: effect size matrix. Each entry is drawn from $N(0,1)$.

In the covariate component, we have

- X: covariate matrix. Each entry is drawn from $N(0,1)$. As in Simulation Design 1, the first K_1 columns are designated as the known covariates (X_1, $n \times K_1$), and the last K_2 columns are designated as the hidden covariates (X_2, $n \times K_2$).
- B_2: effect size matrix. Each entry is drawn from $N(0,1)$ and scaled (see below).

Lastly, in the noise component, we have

- E: noise matrix. Each entry is drawn from $N(0,1)$ and scaled (see below).

Alternatively, (S4) can be written as

$$\underset{n\times 1}{(Y)_j} = \underset{n\times q}{S_j}\ \underset{q\times 1}{(IB_1)_j} + \underset{n\times(K_1+K_2)}{X}\ \underset{(K_1+K_2)\times 1}{(B_2)_j} + \underset{n\times 1}{(E)_j}\,,\ j = 1, \cdots, p\,, \tag{S6}$$

where $(Y)_j$, $(IB_1)_j$, $(B_2)_j$, and $(E)_j$ denote the jth column of Y, $I \odot B_1$, B_2, and E, respectively, and S_j denotes $SArray[\,,\,,j]$.

The scaling for B_2 and E is to ensure that `PVEGenotype` and `PVECovariates` are as desired. Specifically, for gene j, if $\text{Var}\left(S_j\,(IB_1)_j\right) \neq 0$, then we scale $(B_2)_j$ so that

$$\frac{\text{Var}\left(X\,(B_2)_j\right)}{\text{Var}\left(S_j\,(IB_1)_j\right)} = \frac{\texttt{PVECovariates}}{\texttt{PVEGenotype}} \tag{S7}$$

and separately scale $(E)_j$ so that

$$\frac{\text{Var}\left((E)_j\right)}{\text{Var}\left(S_j\,(IB_1)_j\right)} = \frac{1 - \texttt{PVEGenotype} - \texttt{PVECovariates}}{\texttt{PVEGenotype}}\,. \tag{S8}$$

If $\mathrm{Var}\left(S_j\,(IB_1)_j\right) = 0$ (which is the case when $(IB_1)_j$ is a zero vector, i.e., when gene j has zero effect SNPs), then we only scale $(E)_j$ so that

$$\frac{\mathrm{Var}\left((E)_j\right)}{\mathrm{Var}\left(X\,(B_2)_j\right)} = \frac{1 - \texttt{PVECovariates}}{\texttt{PVECovariates}}. \tag{S9}$$

S3.2 Definition of truth

In a simulated data set, I is a $q \times p$ binary matrix. The lj-th entry being one means that the lth local common SNP for gene j is an effect SNP for gene j. However, due to LD, the expression level of a gene may be strongly associated with SNPs other than its effect SNPs.

Therefore, we define I_{cor}, also a $q \times p$ binary matrix, based on *SArray* and I and use it as the truth when calculating AUPRCs. The lj-th entry of I_{cor} is defined as one if and only if the lth local common SNP for gene j is highly correlated with *any* of gene j's effect SNPs (correlation ≥ 0.9 in absolute value).

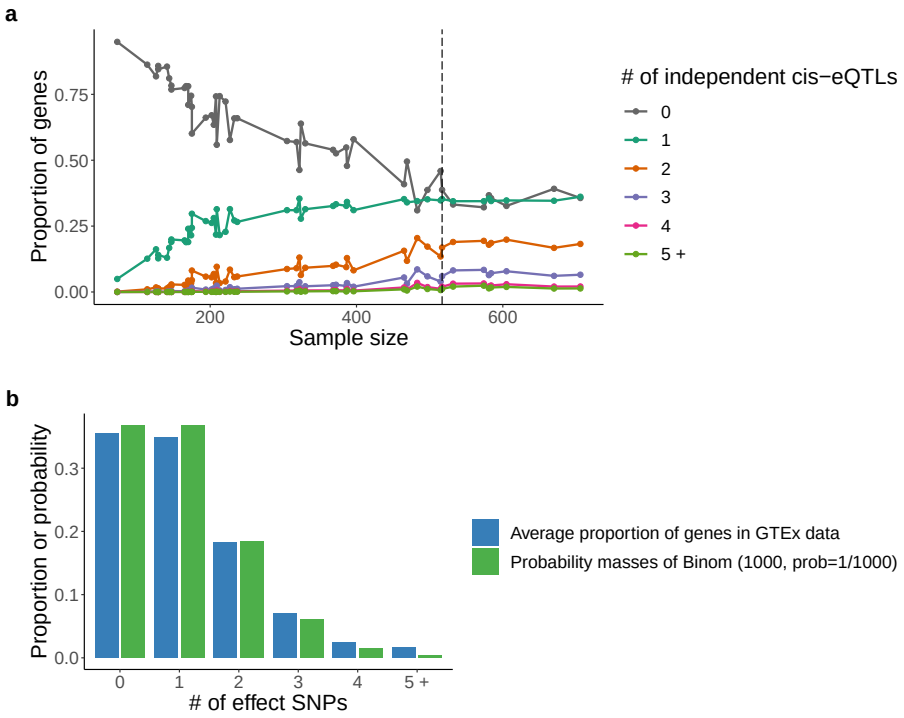

Figure S2: In Simulation Design 2, we find that $Binom(1000, \text{prob} = 1/1000)$ approximates the empirical distribution of the number of independent cis-eQTLs per gene in GTEx data [10] well. **a** Given a tissue type, which corresponds to a sample size, we plot the proportion of genes with 0, 1, 2, 3, 4, or 5 or more independent cis-eQTLs (the proportions add up to one; data from GTEx [10]). We find that the proportions stabilize once the sample size reaches about 517 (dashed line). **b** For the eight tissue types with sample size ≥ 517, we take the average proportion of genes with 0 independent cis-eQTLs, 1 independent cis-eQTL, etc. and plot them in the blue bars. The green bars represent the probability mass function of $Binom(1000, \text{prob} = 1/1000)$ (with the tail probabilities combined together).

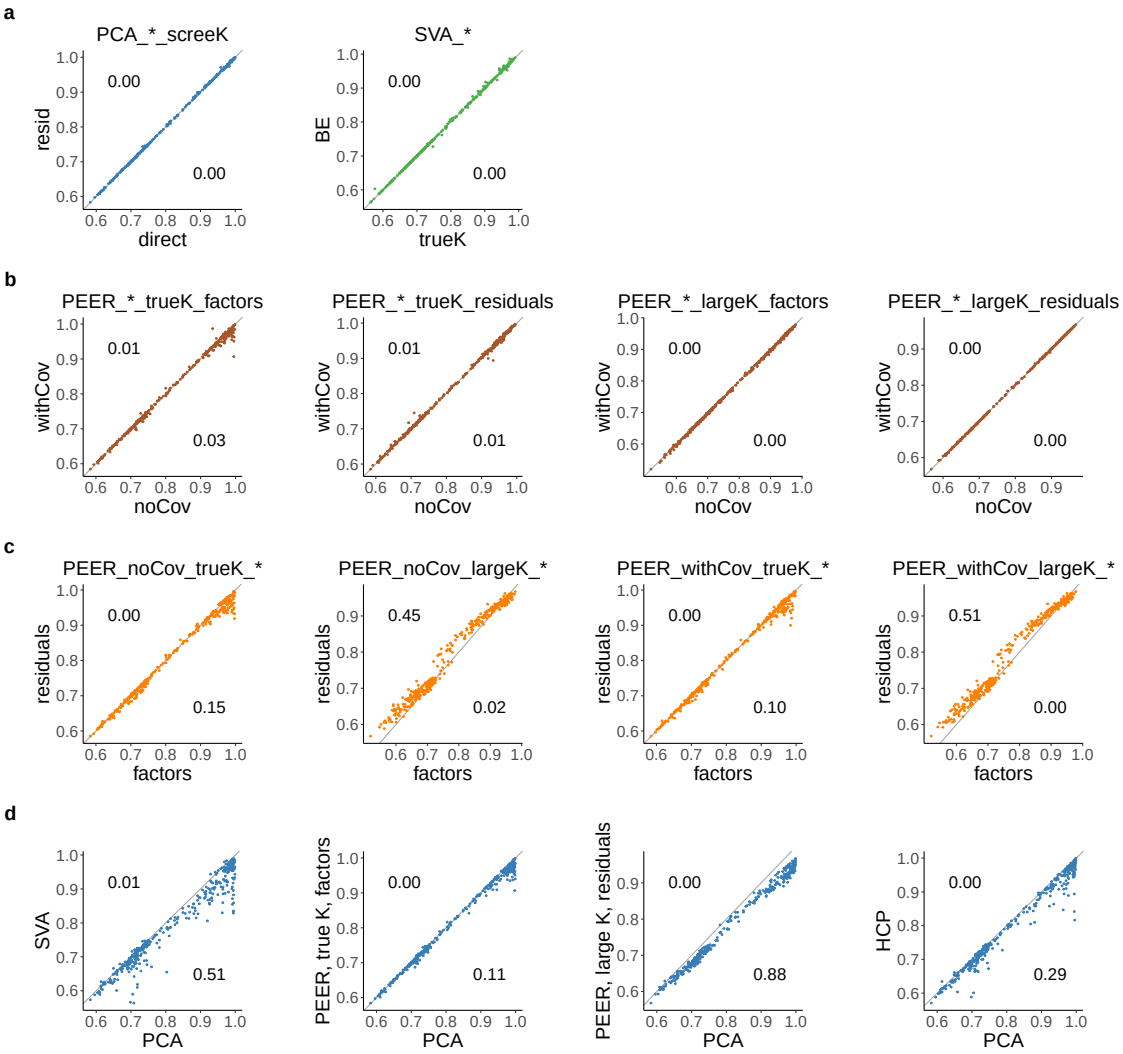

Figure S3: This figure shows how we select a few representative methods from the 15 methods for detailed comparison in Simulation Design 2 (**a, b, c**) and a dataset-by-dataset comparison of the selected representative methods (**d**). The x-axis and y-axis both represent AUPRCs of different methods. Each scatter plot contains 352 points, each of which corresponds to a simulated data set in Simulation Design 2. The number on the upper-left corner of each scatter plot represents the proportion of points that satisfy $y > 1.02\,x$, and the number on the lower-right corner represents the proportion of points that satisfy $x > 1.02\,y$, where x and y denote the coordinates of each point. **a** The two PCA methods perform almost identically, so for simplicity, we select PCA_direct_screeK. The two SVA methods perform almost identically as well, so we select SVA_BE. **b** Whether the known covariates are inputted when PEER is run has little effect on the AUPRC. **c** When we use the true K, the factor approach outperforms the residual approach, but when we use a large K, the residual approach outperforms the factor approach. Therefore, we select PEER_withCov_trueK_factors and PEER_withCov_largeK_residuals as the representative PEER methods. **d** Among the selected representative methods, PCA outperforms SVA, PEER, and HCP in terms of AUPRC in 11% to 88% of the simulated data sets and underperforms them in close to 0% of the simulated data sets.

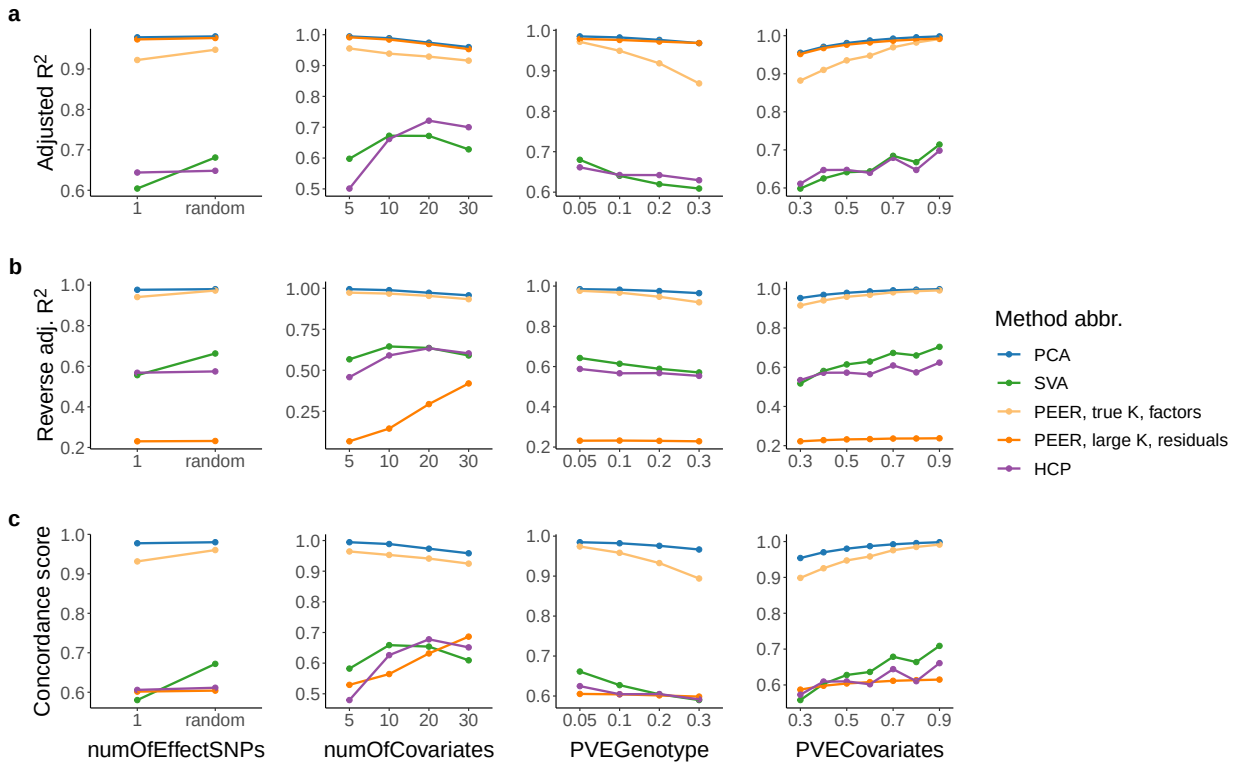

Figure S4: Detailed adjusted R^2, reverse adjusted R^2, and concordance score comparison of the selected representative methods (Table 1) in Simulation Design 2. Each point represents the average across simulated data sets. PCA performs the best in all three regards. PEER with a large K (dark orange line) performs well in terms of adjusted R^2 but falls short in terms of reverse adjusted R^2.

S4 Compared methods

We compare the runtime and performance of 15 methods based on simulation studies, including Ideal, Unadjusted, and 13 variants of PCA, SVA, PEER, and HCP (Table 1). The details of Simulation Design 1 and Simulation Design 2 are described in Sections S2 and S3, respectively. Recall that in each simulated data set, Y denotes the gene expression matrix ($n \times p$, sample by gene), X_1 denotes the known covariate matrix ($n \times K_1$, sample by covariate), and X_2 denotes the hidden covariate matrix ($n \times K_2$, sample by covariate). The genotype information is stored in S in Simulation Design 1 and $SArray$ in Simulation Design 2. In this work, we use K to denote the number of inferred covariates, which are called PCs, SVs, PEER factors, and HCPs in PCA, SVA, PEER, and HCP, respectively.

Given a simulated data set, each of the 15 methods consists of two steps: hidden variable inference step (not applicable for Ideal and Unadjusted) and QTL step. In the hidden variable inference step, we run PCA, SVA, PEER, or HCP to obtain the inferred covariates (and the expression residuals in the case of PEER; Figure 1). In the QTL step, given a gene-SNP pair, we

run a linear regression with the gene expression vector (or the residual vector from PEER) as the *response* and the genotype vector and covariates as *predictors*, where the choice of the response and covariates depends on the method (Table 1); thus we obtain the p-value for the null hypothesis that the coefficient corresponding to the genotype vector is zero given the covariates. In Simulation Design 1, we investigate the association between each gene's expression level and each SNP in the *entire genome* for a simultaneous detection of cis-QTL and trans-QTL relations. In Simulation Design 2, we investigate the association between each gene's expression level and each of the gene's *local common SNPs* for a cis-QTL analysis.

For Ideal, we assume that X_2 is known. Therefore, we use X_1 and X_2 as covariates in the QTL step. For Unadjusted, we use X_1 as the covariates.

We devise two ways to use PCA to account for the hidden covariates. For PCA_direct_screeK, we run PCA on Y directly. For PCA_resid_screeK, we first residualize Y against X_1 and then run PCA on the residual matrix. In this work, PCA is run *with* centering and scaling unless otherwise specified; given A, an $n \times p_1$ matrix, and B, an $n \times p_2$ matrix, both observation by feature, to *residualize A against B* means to take each column of A, regress it against B, and replace the original column of A with the residuals from the linear regression. For both methods, since the scree plots always suggest the true "number of hidden covariates" ($K_1 + K_2$ for PCA_direct_screeK, K_2 for PCA_resid_screeK) as the reasonable number of PCs to choose within plus or minus one (usually exactly; Figure S1), we set the number of PCs to be the true "number of hidden covariates". For PCA_direct_screeK, we filter out the known covariates that are captured well by the top PCs (unadjusted $R^2 \geq 0.9$) and use the remaining known covariates along with the top PCs as covariates in the QTL step. For PCA_resid_screeK, no filtering is needed.

Here we describe the two hidden variable inference methods for SVA: SVA_trueK and SVA_BE. Since the SVA package [31] requires the user to input at least one variable of interest (Figure 1) and using too many variables of interest causes the package to fail, when running SVA, we input the top PC of the genotype matrix (S in Simulation Design 1, collapsed version of $SArray$ in Simulation Design 2) as the variable of interest. We also input X_1 as the known covariates because the package documentation indicates that the known covariates should be provided if available. The SVA package allows the user to specify K. Alternatively, it can automatically choose K using a variation of the Buja and Eyuboglu (BE) algorithm [41]. Therefore, in SVA_trueK, we set $K = K_2$, and in SVA_BE, we let the package choose K automatically. In both cases, we use X_1 and the surrogate variables (SVs) as covariates in the QTL step.

There are several different ways to use PEER [32] but no consensus in the literature on which one is the best. In the hidden variable inference step, PEER can be run with or without the known covariates when there are known covariates available (Stegle et al. [32] do not give an explicit recommendation as to which approach should be used, and both approaches are used in practice [9–11, 16]), and K has to be specified by the user (Stegle et al. [24, 32] claim that the performance of PEER does not deteriorate as K increases). In the QTL step, one can include the PEER factors as covariates (we call this the "factor approach") or use the expression residuals outputted by PEER as the response (and not use any known or inferred covariates; we call this the "residual approach"). For completeness, we compare $2^3 = 8$ ways of using PEER (the default priors are always used):

PEER is run with or without the known covariates; PEER is run using the true "number of hidden covariates" ($K_1 + K_2$ when PEER is run without the known covariates, K_2 when PEER is run with the known covariates) or using a large K (K=50); and either the factor approach or the residual approach is used in the QTL step.

The HCP package requires the user to specify K and three tuning parameters: λ_1, λ_2, and λ_3 (Section S5.2). The package documentation suggests choosing K and the tuning parameters via a grid search. However, no specific recommendations are given regarding the choice of the score function. In practice, users of HCP often choose K and the tuning parameters by maximizing the number of discoveries [22, 23]. For our simulation studies, such an approach would be computationally prohibitive. Therefore, for simplicity, we set $K = K_2$ and $\lambda_1 = \lambda_2 = \lambda_3 = 1$; the latter is because we do not want to give more weight to the penalty terms than the main term in the objective function (Section S5.2).

Reference	**GTEx data version**	**QTL analysis**	**Data transformation**	**Known covariates inputted**	**# of PEER factors**	**Factor or residual approach**
(A)	(B)	(C)	(D)	(E)	(F)	(G)
GTEx Consortium [9]	V6p	eQTL (cis and trans)	INT within feature	No	Maximizes cis-eGenes	Factor
GTEx Consortium [10]	V8	eQTL (cis and trans)	INT within feature	No	Maximizes cis-eGenes	Factor
		sQTL (cis and trans)	INT within sample	No	15	Factor
Li et al. [11]	V7	3′aQTL (cis)	No transformation	Yes	Follows GTEx [9]	Factor

Table S3: Summary of QTL analyses performed by GTEx [9, 10] and Li et al. [11]. "INT" in (D) stands for "inverse normal transform" [37]. (E), (F), and (G) summarize how PEER is used (Section S4). GTEx [9, 10] chooses the number of PEER factors for its eQTL analyses (including cis and trans) by maximizing the number of discovered cis-eGenes for each pre-defined sample size bin. The number of PEER factors selected is 15 for $n < 150$, 30 for $n \in [150, 250)$, and 35 for $n \geq 250$ for GTEx V6p eQTL analyses [9] and 15 for $n < 150$, 30 for $n \in [150, 250)$, 45 for $n \in [250, 350)$, and 60 for $n \geq 350$ for GTEx V8 eQTL analyses [10], where n denotes the sample size. Li et al. [11] use the numbers of PEER factors chosen by GTEx [9].

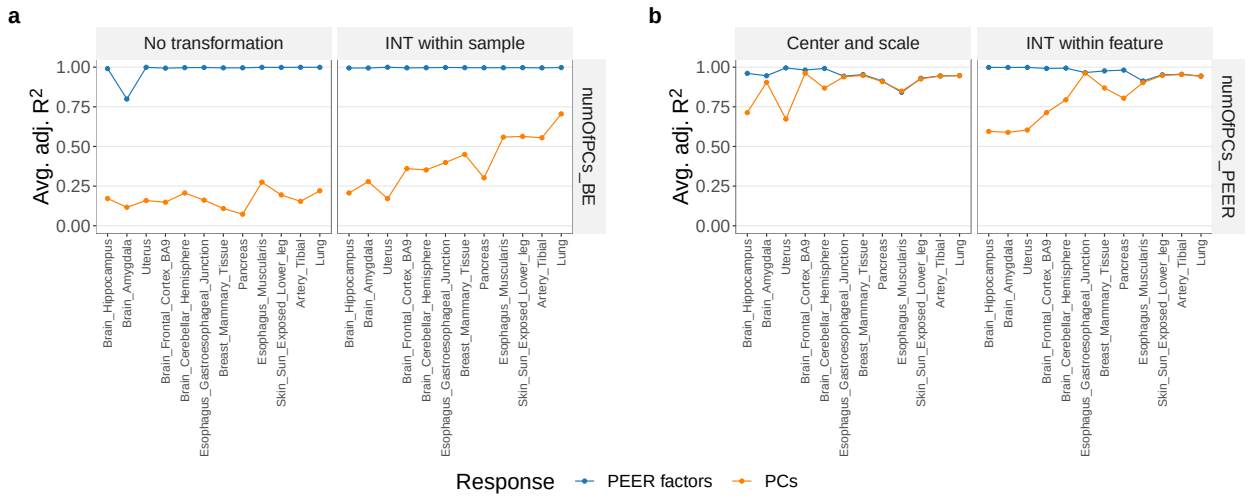

Figure S5: In the 3′aQTL data prepared by Li et al. [11] from GTEx RNA-seq reads [9], PEER factors fail to capture important variance components of the molecular phenotype data when the data transformation method is "No transformation" or "INT within sample" (**a**; the numbers of PCs are chosen via BE (Algorithm 3)); on the other hand, PEER factors span roughly the same linear subspace as the top PCs when the data transformation method is "Center and scale" or "INT within feature" (**b**; the numbers of PCs are equal to the numbers of PEER factors). Given m PEER factors and n PCs from the same post-transformation molecular phenotype matrix ($m \geq n$ in **a**, $m = n$ in **b**), we calculate m adjusted R^2's by regressing each PEER factor against the PCs and plot the average in blue. Similarly, we calculate n adjusted R^2's by regressing each PC against the PEER factors and plot the average in orange.

Algorithm S1: Reordering of PEER factors based on PCs (Figure 5).

Inputs:

- K PEER factors.
- K PCs.

Output: K PEER factors (reordered).

for $k \leftarrow 1$ **to** K **do**
  Select the PEER factor that is the most highly correlated with the kth PC from the PEER factors that have not been selected yet.
end
return the PEER factors in the order that they are selected.

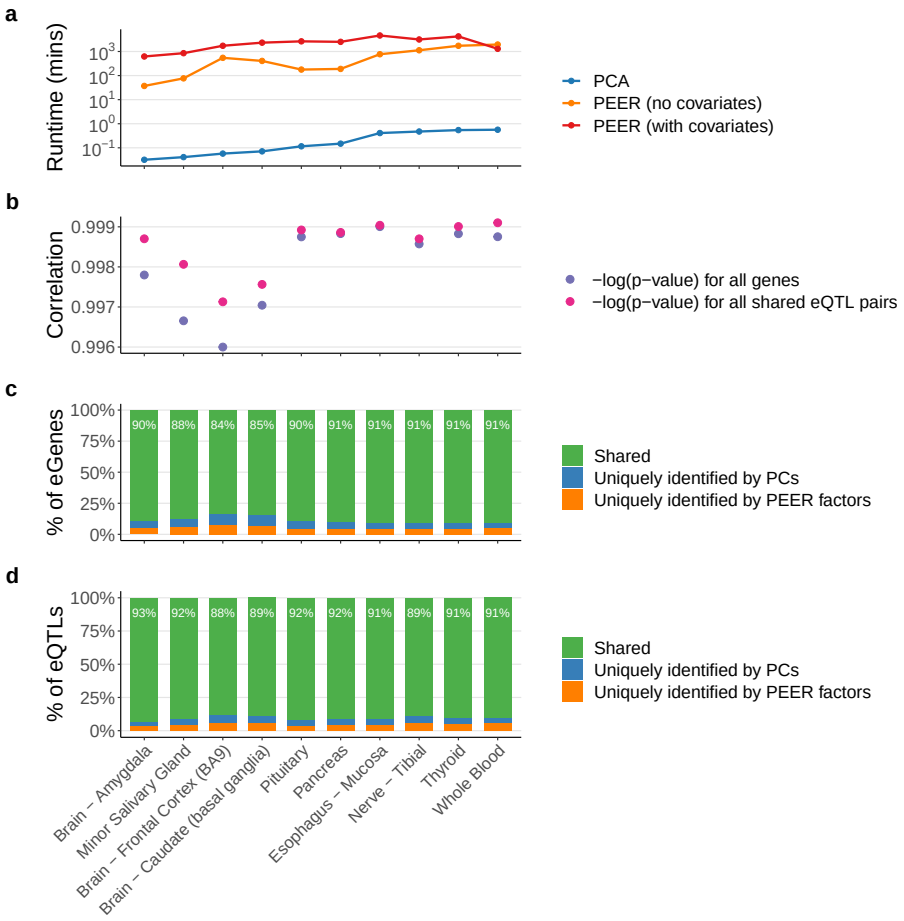

Figure S6: In GTEx eQTL data [10], PEER is at least three orders of magnitude slower than PCA (**a**), and replacing the PEER factors with PCs in GTEx's FastQTL pipeline [10, 13] produces almost identical cis-eQTL results (**b, c, d**). The x-axis shows 10 randomly selected tissue types with increasing sample sizes. **a** For a given gene expression matrix, running PEER without the known covariates (GTEx's approach) takes up to about $1,900$ minutes (equivalent to about 32 hours; Whole Blood), while running PCA (with centering and scaling; our approach) takes no more than a minute. For comparison, we also run PEER *with* the known covariates using the numbers of PEER factors selected by GTEx. This approach takes even longer (up to about $4,600$ minutes, equivalent to about 77 hours; Esophagus - Mucosa). **b** The p-values produced by GTEx's approach and our approach are highly correlated (correlations between the negative common logarithms are shown). **c, d** The overlap of the identified eGenes and eQTL pairs between the two approaches is generally around 90% (see Figure S7 for more detail).

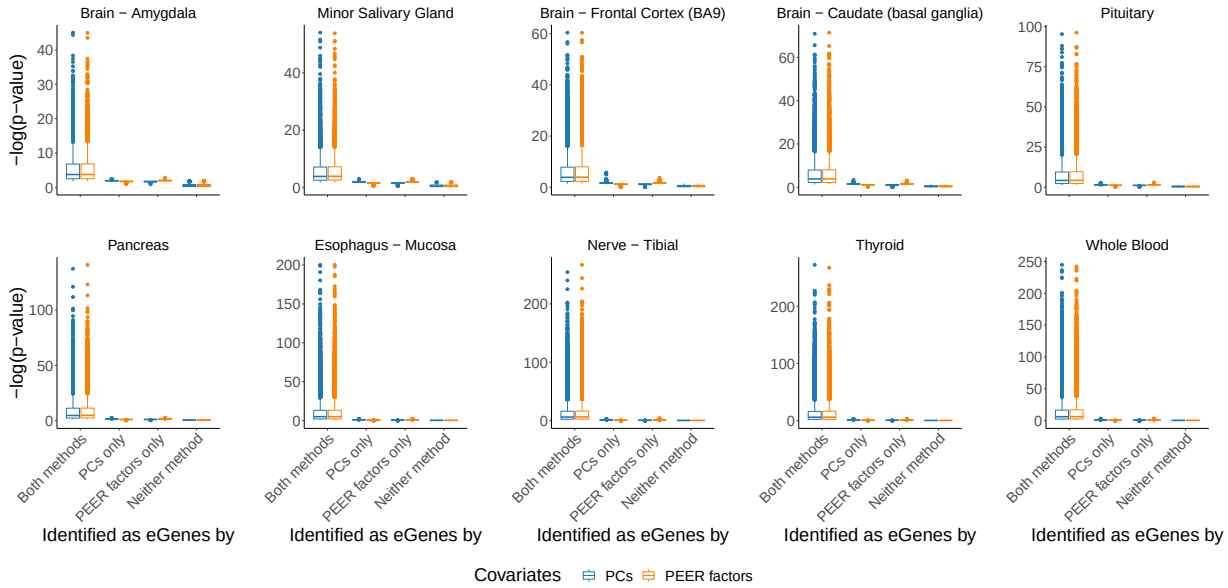

Figure S7: Following the analysis in Figure S6, we find that the eGenes uniquely identified by PCs or PEER factors have marginal p-values compared to those identified by both methods.

S5 Theory of PCA and HCP

S5.1 Principal component analysis (PCA)

Principal component analysis (PCA) [34, 35] is a well-established dimension reduction method with many applications. Here we aim to provide a brief summary of its algorithm, derivation, and interpretation.

Let X denote the $n \times p$ observed data matrix that is observation by feature, e.g., a molecular phenotype matrix. We use X instead of Y here to be consistent with standard PCA notations. We assume that the columns of X have been centered and scaled. That is, X satisfies

$$\frac{1}{n}\sum_{i=1}^{n} x_{ij} = 0\,,\ \ j = 1, \cdots, p \tag{S10}$$

and

$$\frac{1}{n-1}\sum_{i=1}^{n} x_{ij}^2 = 1\,,\ \ j = 1, \cdots, p\,, \tag{S11}$$

where x_{ij} denotes the ij-th entry of X.

The PCA algorithm consists of two steps. In the first step, we calculate the sample covariance matrix $\widehat{\Sigma}$ and perform eigendecomposition on it:

$$\widehat{\Sigma} = \frac{1}{n} X^\top X \qquad \text{definition of sample covariance matrix} \tag{S12}$$

$$:= Q \Lambda Q^\top , \qquad \text{eigendecomposition} \tag{S13}$$

where

$$\underset{p\times p}{Q} = \begin{bmatrix} | & & | \\ q_1 & \cdots & q_p \\ | & & | \end{bmatrix} \tag{S14}$$

is an orthogonal matrix whose columns are eigenvectors of $\widehat{\Sigma}$, and

$$\underset{p\times p}{\Lambda} = \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_p \end{bmatrix}, \ \lambda_1 \geq \cdots \geq \lambda_p \geq 0 , \tag{S15}$$

is a diagonal matrix whose diagonal entries are the corresponding eigenvalues of $\widehat{\Sigma}$. We know that $\widehat{\Sigma}$ is orthogonally diagonalizable because it is a symmetric matrix (recall the spectral theorem for real matrices [42]: a matrix is orthogonally diagonalizable if and only if it is symmetric). The eigenvalues are all non-negative because $\widehat{\Sigma}$ is positive semidefinite.

In the second step, we calculate Z as

$$Z = XQ , \tag{S16}$$

where the columns of Z are called the *principal components* (PCs) or *scores*, and Q is called the *loading matrix* or *rotation matrix*. It is worth noting that some authors may refer to $q_1, \cdots, q_p$ as the PCs. This use of terminology is confusing and should be avoided [43].

The above two steps conclude the PCA algorithm. In practice, however, singular value decomposition (SVD) of the data matrix is often used as a more computationally efficient way of finding the loading matrix and the PCs [34].

The most common derivation of PCA is based on maximum variance [44]. First, we define

$\alpha_1^*, \cdots, \alpha_p^* \in \mathbb{R}^p$ sequentially as

$$\alpha_1^* = \underset{\alpha_1 \in \mathbb{R}^p}{\arg\max}\, \mathrm{Var}\left(X\alpha_1\right) \qquad \text{subject to } \|\alpha_1\|_2 = 1\,, \tag{S17}$$

$$\alpha_2^* = \underset{\alpha_2 \in \mathbb{R}^p}{\arg\max}\, \mathrm{Var}\left(X\alpha_2\right) \qquad \text{subject to } \|\alpha_2\|_2 = 1\,,\ \alpha_2^\top \alpha_1^* = 0\,, \tag{S18}$$

$$\vdots$$

$$\alpha_p^* = \underset{\alpha_p \in \mathbb{R}^p}{\arg\max}\, \mathrm{Var}\left(X\alpha_p\right) \qquad \text{subject to } \|\alpha_p\|_2 = 1\,,\ \alpha_p^\top \alpha_j^* = 0\ \ \forall\, j < p\,. \tag{S19}$$

Then, we define the PCs of X as $X\alpha_1^*, \cdots, X\alpha_p^*$. That is, the PCs are defined sequentially as the linear combinations of the columns of X with maximum variances, subject to certain constraints. It can then be shown that $\alpha_1^*, \cdots, \alpha_p^*$ are given by $q_1, \cdots, q_p$ respectively, where $q_1, \cdots, q_p$ are eigenvectors of $\widehat{\Sigma}$ as defined in (S14).

A complementary property of PCA, which is closely related to the original discussion of Pearson [45], is the minimum reconstruction error property. Given $K < p$, we define Q_K as the matrix that contains the first K columns of Q. That is,

$$\underset{p\times K}{Q_K} := \begin{bmatrix} | & & | \\ q_1 & \cdots & q_K \\ | & & | \end{bmatrix}. \tag{S20}$$

The minimum reconstruction error property of PCA states that Q_K is a global minimizer of the loss function

$$\mathscr{J}\left(\widetilde{Q}_K\right) := \left|\left|\left| X - X\widetilde{Q}_K\widetilde{Q}_K^\top \right|\right|\right|_F^2 \tag{S21}$$

$$= \sum_{i=1}^n \left\| x_i^\top - x_i^\top \widetilde{Q}_K\widetilde{Q}_K^\top \right\|_2^2 = \sum_{i=1}^n \left\| x_i - \widetilde{Q}_K\widetilde{Q}_K^\top x_i \right\|_2^2, \tag{S22}$$

where $\widetilde{Q}_K$ denotes an arbitrary $p \times K$ matrix whose columns are orthonormal, $|||\cdot|||_F$ denotes the Frobenius norm of a matrix, and $x_i^\top$ denotes the ith row of X. Since $\widetilde{Q}_K\widetilde{Q}_K^\top x_i$ represents the (orthogonal) projection of x_i onto the subspace spanned by the columns of $\widetilde{Q}_K$, (S22) measures the total squared ℓ_2 error when approximating each x_i with its projection onto the subspace spanned by the columns of $\widetilde{Q}_K$.

A central idea of PCA is the proportion of variance explained by each PC. To establish this concept, we claim that

$$\sum_{j=1}^p \mathrm{Var}\left(X_j\right) = \sum_{j=1}^p \mathrm{Var}\left(Z_j\right), \tag{S23}$$

$$\mathrm{Var}\left(Z_j\right) = \lambda_j\,,\ j = 1, \cdots, p\,, \tag{S24}$$

and

$$\mathrm{Cov}\left(Z_j\,,\ Z_{j'}\right) = 0\,,\ j, j' = 1, \cdots, p\,,\ j \neq j'\,, \tag{S25}$$

where X_j denotes the jth column of X (the jth original variable) and Z_j denotes the jth column of Z (the jth PC). (S25) means that the PCs are uncorrelated with each other.

We prove (S24) and (S25) by calculating $\widehat{\Sigma}_Z$, the sample covariance matrix of Z (we know that the columns of Z are centered by (S10) and (S16)):

$$\begin{aligned}
\widehat{\Sigma}_Z &= \frac{1}{n} Z^\top Z && \text{definition of sample covariance matrix} && \text{(S26)}\\
&= \frac{1}{n} (XQ)^\top XQ && \text{plugging in (S16)} && \text{(S27)}\\
&= Q^\top \left(\frac{1}{n} X^\top X\right) Q && && \text{(S28)}\\
&= Q^\top \left(Q \Lambda Q^\top\right) Q && \text{plugging in (S13)} && \text{(S29)}\\
&= \Lambda\,. && && \text{(S30)}
\end{aligned}$$

(S23) can be proven by the following:

$$\begin{aligned}
\sum_{j=1}^{p} \mathrm{Var}\left(X_j\right) &= \mathrm{Tr}\left(\widehat{\Sigma}\right) && \text{definition of trace and } \widehat{\Sigma} && \text{(S31)}\\
&= \mathrm{Tr}\left(Q \Lambda Q^\top\right) && \text{plugging in (S13)} && \text{(S32)}\\
&= \mathrm{Tr}\left(\Lambda Q^\top Q\right) && \text{cyclic property of trace} && \text{(S33)}\\
&= \mathrm{Tr}\left(\Lambda\right) && && \text{(S34)}\\
&= \sum_{j=1}^{p} \mathrm{Var}\left(Z_j\right)\,. && \text{by (S24)} && \text{(S35)}
\end{aligned}$$

Because of (S23) and (S24), we may define the proportion of variance in the original data explained by the jth PC as

$$\frac{\lambda_j}{\sum_{j'=1}^{p} \mathrm{Var}\left(X_{j'}\right)} = \frac{\lambda_j}{\sum_{j'=1}^{p} \mathrm{Var}\left(Z_{j'}\right)} = \frac{\lambda_j}{\sum_{j'=1}^{p} \lambda_{j'}}\,, \tag{S36}$$

which provides a basis for deciding the number of PCs to keep (e.g., Algorithms S2 and S3).

S5.2 Hidden covariates with prior (HCP) and its connection to PCA

Hidden covariates with prior (HCP) [33] is a popular hidden variable inference method for QTL mapping defined by minimizing a loss function. Neither Mostafavi et al. [33] nor the HCP package documents the HCP method well. For example, the squares in the loss function (S37) are missing in both Mostafavi et al. [33] and the package documentation, but one can deduce that the squares are there by inspecting the coordinate descent steps in the source code of the R package. Here we aim to provide a better, more accurate documentation of the HCP method and point out its connection to PCA.

Given Y, the molecular phenotype matrix ($n \times p$, sample by feature), X_1, the known covariate matrix ($n \times K_1$, sample by covariate), K, the number of inferred covariates (HCPs), and $\lambda_1,\ \lambda_2,\ \lambda_3 > 0$, the tuning parameters, HCP looks for

$$\underset{X_2,\, W_1,\, W_2}{\arg\min} \left\{ \left|\left|\left| \underset{n\times p}{Y} - \underset{n\times K}{X_2}\ \underset{K\times p}{W_2} \right|\right|\right|_F^2 + \lambda_1 \left|\left|\left| \underset{n\times K}{X_2} - \underset{n\times K_1}{X_1}\ \underset{K_1\times K}{W_1} \right|\right|\right|_F^2 + \lambda_2 |||W_1|||_F^2 + \lambda_3 |||W_2|||_F^2 \right\}, \tag{S37}$$

where $|||\cdot|||_F$ denotes the Frobenius norm of a matrix, X_2 is the hidden covariate matrix, and W_1 and W_2 are weight matrices of the appropriate dimensions. The name of the method, "hidden covariates with prior", comes from the second term in (S37), where the method informs the hidden covariates with the known covariates. The optimization is done through coordinate descent with one deterministic initialization. The columns of the obtained X_2 are reported as the HCPs.

From (S37), we see that the HCP method is closely related to PCA. The first term in (S37) is very similar to (S21), the only difference being that the rows of W_2 in (S37) are not required to be orthonormal and X_2 is not required to be equal to $YW_2^\top$.

Algorithm S2: The elbow method for choosing K in PCA (based on distance to diagonal line).

Input: X, $n \times p$ observed data matrix, observation by feature.
Output: K, the number of PCs selected.

1. Define $d = \texttt{min}(n, p)$. This is the total number of PCs.
2. Run PCA on X with centering and scaling.
3. Obtain the proportion of variance explained by each PC, $t_1, \cdots, t_d$. // $\sum_{j=1}^{d} t_j = 1$.
4. Consider $(1, t_1), \cdots, (d, t_d) \in \mathbb{R}^2$. // d points in $\mathbb{R}^2$.
5. Select K by choosing the point that is the farthest from the diagonal line, i.e., the line that passes through the first point, $(1, t_1)$, and the last point, (d, t_d). Specifically, the distance from (x_0, y_0) to the line that passes through (x_1, y_1) and (x_2, y_2) is given by $|\,(x_2 - x_1)(y_1 - y_0) - (x_1 - x_0)(y_2 - y_1)\,| \,/\, ((x_2 - x_1)^2 + (y_2 - y_1)^2)^{1/2}$.
6. **return** K.

Algorithm S3: The Buja and Eyuboglu (BE) algorithm for choosing K in PCA.

Inputs:

- X, $n \times p$ observed data matrix, observation by feature.
- B, number of permutations (default is 20).
- α, significance level (default is 0.05).

Output: K, the number of PCs selected.

Define $d = \texttt{min}(n, p)$. This is the total number of PCs.
Run PCA on X with centering and scaling.
Obtain the proportion of variance explained (PVE) by each PC, $t_1, \cdots, t_d$. `// ` $\sum_{j=1}^{d} t_j = 1$. `Observed test statistics.`
for $b \leftarrow 1$ **to** B **do**
  Obtain $X^{(b)}$ by permuting each column of X. `// Permute the observations in each feature.`
  Run PCA on $X^{(b)}$ with centering and scaling.
  Obtain the PVE of each PC, $t_1^{(b)}, \cdots, t_d^{(b)}$. `// ` $\sum_{j=1}^{d} t_j^{(b)} = 1$.
end
The p-value for the jth PC is calculated as $p_j = \sum_{b=1}^{B} \mathbb{1}\{t_j^{(b)} \geq t_j\}/B,\ j = 1, \cdots, d$. `// ` p_j `is calculated as the proportion of permutations where the PVE of the` j`th PC is greater than or equal to the PVE of the` j`th original PC. The greater this proportion is, the larger the` p`-value is, and the less significant the PC is.`
for $j \leftarrow 2$ **to** d **do**
  If $p_j \leq p_{j-1}$, then set $p_j = p_{j-1}$. `// Enforce monotone increase of the` p`-values.`
end
Set K to be the maximum j such that $p_j \leq \alpha$.
return K.