

Allele-Specific QTL Fine-Mapping with PLASMA

by Austin T. Wang

Submitted to the
Department of Electrical Engineering and Computer Science
in Partial Fulfillment of the Requirements for the Degree of
Master of Engineering in Computer Science and Molecular Biology
at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2020

© 2020 Massachusetts Institute of Technology. All rights reserved.

Author: _____

Austin T. Wang
Department of Electrical Engineering and Computer Science
January 26, 2020

Certified By: _____

Manolis Kellis: Thesis Supervisor
Professor of Computer Science, MIT
January 26, 2020

Certified By: _____

Alexander Gusev: Thesis Supervisor
Assistant Professor, Dana-Farber Cancer Institute, Harvard Medical School
January 26, 2020

Accepted by: _____

Katrina LaCurts
Chair, Master of Engineering Thesis Committee

Allele-Specific QTL Fine-Mapping with PLASMA

by Austin T. Wang

Submitted to the Department of Electrical Engineering and Computer Science

January 26, 2020

In Partial Fulfillment of the Requirements for the Degree of Master of Engineering in
Computer Science and Molecular Biology

Under joint supervision by Professors Manolis Kellis and Alexander Gusev

Abstract

We introduce PLASMA (PopuLation Allele-Specific MApping), a statistical fine-mapping method that leverages allele-specific (AS) genomic data to improve detection of quantitative trait loci (QTLs) with causal effects on molecular traits. In simulations, PLASMA accurately prioritizes causal QTL variants over a wide range of genetic architectures. Applied to RNA-Seq data from 524 kidney tumor samples, PLASMA achieves a greater power at 50 samples than conventional QTL-based fine-mapping at 500 samples: with over 17% of loci fine-mapped to within 5 causal variants compared to 2% by QTL-based fine-mapping, and a 6.9-fold overall reduction in median credible set size. PLASMA offers high accuracy even at small sample sizes, yielding a 1.3-fold reduction in median credible set size compared to QTL-based fine-mapping when applied to H3K27AC ChIP-Seq from just 28 prostate tumor/normal samples. Our results demonstrate how integrating AS activity can substantially improve the detection of causal variants from existing molecular data.

Contents

1	Introduction	4
2	Methods	5
2.1	Overview of PLASMA	5
2.2	Modeling QTL and AS summary statistics	7
2.3	Inference of credible sets and posterior probabilities	11
2.4	The Jointness Parameter in PLASMA	13
2.5	Generation of simulated loci	13
2.6	Comparison of existing models with PLASMA	14
2.7	Quality control of genotype data	15
2.8	Quality control of RNA-seq data	16
2.9	Quality control of ChIP-seq data	16
2.10	Allele-specific quantification	16
2.11	Functional enrichment analysis	16
3	Results	17
3.1	Simulation framework	17
3.2	Comparison with existing methods in simulation	18
3.3	Inference of multiple causal variants	21
3.4	Fine-mapping of TCGA kidney RNA-Seq data	23
3.5	Fine-mapping of prostate H3k27ac ChIP-Seq data	28
3.6	PLASMA increases functional enrichment of credible set markers	28
4	Discussion	30
5	Appendix	42
5.1	Modeling total expression at a locus (QTL)	42
5.1.1	Modeling genetic effects on total expression	42
5.1.2	Calculation of QTL summary statistics	42
5.2	Modeling allele-specific expression at a locus (AS)	43
5.2.1	Modeling haplotype-specific effects on expression	43
5.2.2	Calculation of AS summary statistics	45
5.2.3	Emprical Analysis of Imperfect Phasing	48
5.3	Inference of causal variants with QTL and AS statistics	48
5.3.1	Modeling the effects of linkage disequilibrium	48
5.3.2	Equivalence in expectation of QTL and AS estimators of LD	48
5.3.3	Jointly modeling total and haplotype-specific effects on expression	50
5.3.4	Modeling summary statistics given a causal configuration	50
5.3.5	The ρ -level credible set	51
5.3.6	The posterior inclusion probability	51
5.4	Computational optimization and implementation	52
5.4.1	Shotgun stochastic search across configurations	52
5.4.2	Calculation of the ρ -level credible set	52
5.4.3	Bayes factor evaluation with matrix reduction	52
5.5	The hyperparameters in terms of heritability	54

1 Introduction

A major open problem in genetics is understanding the biological mechanisms underlying complex traits, which are largely driven by non-coding variants. A widely adopted approach for elucidating these regulatory patterns is the identification of disease variants that also modify molecular phenotypes (such as gene expression) [1–4]. These variants, known as quantitative trait loci (QTLs), are typically single nucleotide polymorphisms (SNPs) that exhibit a statistical association with overall gene expression abundance [5–8]. Although QTL association analysis is now mature, it remains challenging to identify the precise variants that causally influence the molecular trait (as opposed to variants in linkage disequilibrium (LD) with causal variants), a task known as fine-mapping [9]. As only a small subset of QTL-associated markers are estimated to be causal [10,11], direct experimental validation is prohibitive and has motivated statistical fine-mapping solutions [12]. The aim of statistical fine-mapping is to quantify the probability of each marker being causal, allowing one to prioritize the most likely causal markers, and thus formally quantify the effort needed for experimental validation. Recent statistical fine-mapping methods operate on summary QTL statistics and can handle multiple causal variants by modeling the local LD structure [13–16]. These models have two outputs to help guide the prioritization of putative causal SNPs. First, for each marker, a Posterior Inclusion Probability (PIP) is calculated, which corresponds to the marginal probability of causality for the given marker. Second, a $n\%$ -confidence credible set is created: a set of markers with an $n\%$ probability of containing all the causal markers. Although QTL studies have sufficient power to identify thousands of associations, they are typically insufficient for fine-mapping below dozens of credible variants, even for very large studies [5,17]. The need for large studies is severely limits QTL analyses of expensive assays such as ChIP or single-cell RNA-seq, or of difficult to collect tissues.

Here, we sought to improve molecular fine-mapping by leveraging intra-individual allele-specific (AS) signal, which is a measure of cis-regulatory activity that is independent of total, inter-individual variation. For heterozygous variants residing in expressed exons, it is often possible to map expressed reads to each allele and quantify the extent that molecular activity is allele-specific [6,18–21]. AS analysis allows for a precise comparison of the effects on molecular activity that are specific to each allele (cis-effects), while controlling for effects affecting both alleles (trans-

effects). Thus, AS data is inherently less noisy than regular QTL data, which captures total phenotype regardless of source. Allele-specific data has furthermore been used to quantify cis-regulation, implying that both AS and regular QTL features represent the same underlying cis-regulatory patterns [22]. Several methods have recently been developed to robustly identify asQTLs [19, 20, 23], but the calculated association statistics follow a different distribution than QTL summary statistics and cannot be directly integrated into existing fine-mapping software to produce valid posterior measures and credible sets.

To combine the scalability of QTL analysis with the power of AS analysis, we introduce PLASMA (PopuLation Allele-Specific Mapping), a novel fine-mapping method that gains power from both the number of individuals and the number of allelic reads per individual. By modeling each locus across individuals in an allele-specific and LD-aware manner, PLASMA achieves a substantial improvement over existing fine-mapping methods with the same data. We demonstrate through simulations that PLASMA successfully detects causal variants over a wide range of genetic architectures. We apply PLASMA to diverse RNA-Seq data and ChIP-seq data and show a significant improvement in power over conventional QTL-based fine-mapping.

2 Methods

2.1 Overview of PLASMA

PLASMA’s inputs are determined from a given individual-level sequencing-based molecular phenotype (gene or peak) and the corresponding local genotype SNP data (Figure 1a). For each sample, we assume the variant data is phased into haplotypes and expression reads have been mapped to each variant. Reads intersecting heterozygous markers (signified as fSNPs, or feature SNPs, indicated with green or purple on the figure) are then assigned to a particular haplotype, indicated as blue or red on the figure. These reads are then aggregated in an haplotype-specific manner to produce a total expression phenotype and an allelic imbalance phenotype. This aggregation of reads is analogous to the way existing methods such as RASQUAL and WASP calculate allelic fractions and total fragment counts [19, 20]. The total expression phenotype (y) is simply the total number of mapped reads. The allelic imbalance phenotype (w) is defined as the log read ratio between the haplotypes. This log-odds-like phenotype has previously been used to analyze asQTL effect sizes,

showing consistency with conventional QTL analysis [22]. To mitigate the effect of mapping bias, we run state of the art mapping bias and QC pipelines on all RNA-Seq and ChIP-Seq data prior to analysis [19].

PLASMA integrates two statistics computed for each marker to perform fine-mapping: a QTL association statistic (z_β) based on the total phenotype and an AS association statistic (z_ϕ) based on the allelic imbalance phenotype. Figure 1b shows how a causal marker influences total expression and allelic imbalance, and how this effect influences the statistics for the marker. Here, the causal marker’s alternative allele causes higher expression compared to the wild-type allele. Increasing the dosage (x) of the alternative allele increases the total expression (y) at the locus. The effect size (β), consistent with a typical QTL analysis, quantifies the association between a marker’s allelic dosage and the total expression at the locus with a linear relationship with residuals ϵ :

$$\mathbf{y} = \mathbf{x}_i\beta_i + \boldsymbol{\epsilon}_i \tag{1}$$

From this effect size, PLASMA calculates z_β , the QTL association statistic. Note that this statistic is not dependent on haplotype-specific data.

On the other hand, looking at the heterozygotes, the haplotype possessing the alternative allele has a higher expression than the haplotype possessing the wild-type allele. In other words, the direction of imbalance of expression (w) is the same as the phasing (v) of the allele. The ϕ effect size quantifies the association between a marker’s phasing with the imbalance of expression. An important departure from existing methods is that PLASMA models a linear relationship between the phase of a causal marker and the log read ratio, rather than directly relating the genotype to the allelic fraction in a non-linear manner with residuals ζ :

$$\mathbf{w} = \mathbf{v}_i\phi_i + \boldsymbol{\zeta}_i \tag{2}$$

To calculate the AS association statistic z_ϕ , PLASMA models the quality of each sample, taking into account each sample’s read coverage and read overdispersion (Figure 1c).

These QTL and AS association statistics, together with the local LD matrix, are then jointly used to fine-map the locus (Figure 1c). Since PLASMA models both z_β and z_ϕ as a linear combina-

tion of genotypes, z_β and z_ϕ have identical LD (see Appendix for proof). PLASMA assumes that the QTL and AS statistics measure the same underlying cis-regulatory signal and are thus expected to have the same direction and same causal variants (but see Discussion for possible model violations). Although they both measure regulatory effects, the two statistics have independent noise because the haplotype-level variance within individuals is considered only in AS analysis, allowing them to be used jointly in fine-mapping. Furthermore, PLASMA accepts, as a hyperparameter, a correlation between QTL and AS effects, allowing the two sets of statistics to utilize a joint probability distribution (though our analyses show that setting this hyperparameter to zero yields the most power). The distribution is used to assign a probability to a given causal configuration, a binary vector signifying the causal status of each marker in the locus. Although the correlation between QTL and AS causal effects can vary based on the hyperparameter specification, PLASMA assumes that the AS and QTL phenotypes have the same causal variants. PLASMA searches through the space of possible causal configurations, within a constraint on the number of causal variants. This procedure is related to that in CAVIAR, CAVIARBF, and FINEMAP [13–15], but generalized to the two correlated expression phenotypes. From these scored configurations PLASMA computes a posterior inclusion probability (PIP) for each marker, indicating the marginal probability that a marker is causal, and a ρ -level credible set containing the causal variant with ρ probability.

2.2 Modeling QTL and AS summary statistics

Marginal QTL effect sizes for a given locus are calculated under the conventional linear model of total gene expression, with the allelic dosage (\mathbf{x}) as the independent variable, and the total expression (\mathbf{y}) as the dependent variable. Let us consider a QTL study of a given locus with n individuals and m markers. Let \mathbf{y} be an $(n \times 1)$ vector of total expression across the individuals, recentered at zero. Given a marker i , let \mathbf{x}_i be a zero-recentered vector of dosage genotypes. The genetic effect β_i of marker i on total gene expression is defined as follows:

$$\mathbf{y} = \mathbf{x}_i \beta_i + \boldsymbol{\epsilon}_i \tag{3}$$

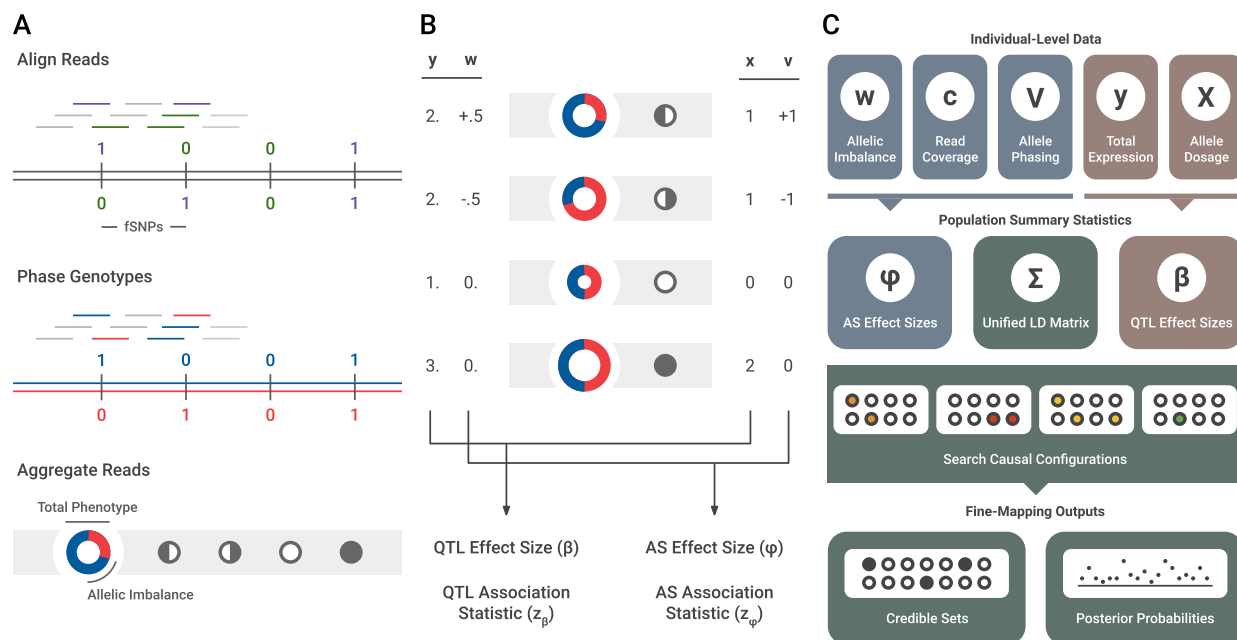


Figure 1: Overview of the PLASMA method. (a) Pre-processing of sequence-based data. First, reads are mapped to the sample's genotype. Reads intersecting markers are colored. Then, the sample's genotype is phased. Reads intersecting heterozygous markers can then be mapped to a particular haplotype. Lastly, reads across the locus are aggregated in an allele-specific manner. To visualize this data, the expression is represented by a ring chart, and the genotypes by pedigree symbols. In the ring chart, the diameter signifies the total read count, and the colors signify the proportion of reads coming from each haplotype. For the pedigree symbols, a white circle signifies a wild-type homozygote, a shaded circle signifies a alternative homozygote, and a half-shaded circle signifies a heterozygote. In heterozygotes, the direction of shading corresponds to the direction of heterozygosity (phasing). (b) Visual representation of QTL and AS statistics under a single causal variant, where the alternative allele increases expression. The total expression (y) is determined by the allelic dosage (x), whereas the allelic imbalance (w) is determined by the phasing (v). These two sets of data are used to calculate QTL and AS association statistics (z_β and z_ϕ). (c) Diagram of PLASMA's fine-mapping process. First, QTL and AS statistics are calculated from read data. Then, these statistics, along with an LD matrix, are used to generate probabilities for causal configurations. By searching through the space of these causal configurations, the model produces credible sets and posterior probabilities for each marker.

The empirical value of β_i is determined with the maximum likelihood estimator, equivalent to the ordinary-least-squares linear regression estimator:

$$\hat{\beta}_i = \left(\mathbf{x}_i^\top \mathbf{x}_i \right)^{-1} \mathbf{x}_i^\top \mathbf{y} \quad (4)$$

The QTL summary statistic (Wald statistic) for marker i is defined as:

$$\hat{z}_{\beta,i} = \frac{\hat{\beta}_i}{\sqrt{\left(\mathbf{x}_i^\top \mathbf{x}_i \right)^{-1} \hat{\sigma}_{y,i}^2}} \quad (5)$$

where $\hat{\sigma}_{y,i}^2$ is calculated from the residuals.

AS effect sizes are calculated under a weighted linear model, with the phasing (\mathbf{v}) as the independent variable, and the allelic imbalance (\mathbf{w}) as the dependent variable. PLASMA models allele-specific expression under the observation that a cis-regulatory variant often has a greater influence on the gene allele of the same haplotype. A marker's phase v is 1 if haplotype A contains the alternative marker allele, -1 if haplotype B contains the alternative marker allele, and 0 if the individual is homozygous for the marker. Let w be the log expression ratio between haplotypes A and B , ϕ_i be the AS effect size of variant i , and ζ_i be the residual, interpreted as the log baseline expression ratio between haplotypes A and B . Additionally, a sampling error $\tau_j = \hat{w}_j - w_j$ is defined for each individual, quantifying the quality of data from the sample. The genetic effect of marker i on allele-specific expression is as follows:

$$\hat{\mathbf{w}} = \mathbf{v}_i \phi_i + \zeta_i + \boldsymbol{\tau} \quad (6)$$

Experimentally-derived AS data, such as RNA-Seq data, yield reads that are mapped to a particular haplotype. For a given individual j , let $c_{A,j}$ be the allele-specific read count from haplotype A . The allele-specific read count is modeled as drawn a beta-binomial distribution, given the total mapped read count c_j :

$$c_{A,j} \sim \text{BB}(\alpha_j, \beta_j, c_j) \quad (7)$$

This beta binomial model is used to estimate the variance of the sampling error τ_j :

$$\hat{\sigma}_{c,j}^2 = \frac{2}{c_j}(1 + \cosh(\hat{w}_j^*))(1 + \rho_{e,j}(c_j - 1)) \quad (8)$$

where $\rho_{e,j}$ is the overdispersion and w_j^* is an adjusted estimator of w_j to reduce the bias of $\hat{\sigma}_{c,j}^2$. (Full derivation in Supplementary Methods).

Due to heteroscedasticity among individuals, the AS effect size ϕ_i is estimated in a weighted manner, giving larger weights to individuals with lower estimated sampling error. Given individual j , the weight for j is set as the inverse of the estimated sampling error variance:

$$\omega_j = \frac{1}{\hat{\sigma}_{c,j}^2} \quad (9)$$

Let weight matrix $\mathbf{\Omega}$ be a diagonal matrix with $\mathbf{\Omega}_{j,j} = \omega_j$. We use the weighted-least-squares estimator for ϕ_i :

$$\hat{\phi}_i = \left(\mathbf{v}_i^\top \mathbf{\Omega} \mathbf{v}_i \right)^{-1} \mathbf{v}_i^\top \mathbf{\Omega} \hat{\mathbf{w}} \quad (10)$$

With this estimator, the AS association statistic for marker i is calculated as the AS effect size divided by the estimated variance of the effect size (full derivation in Appendix):

$$\hat{z}_{\phi,i} = \frac{\hat{\phi}_i}{\sqrt{\left(\mathbf{v}_i^\top \mathbf{\Omega} \mathbf{v}_i \right)^{-2} \left(\left(\mathbf{v}_i^\top \mathbf{\Omega}^2 \mathbf{v}_i \right) \hat{\sigma}_{w,i}^2 + \mathbf{v}_i^\top \mathbf{\Omega} \mathbf{v}_i \right)}} \quad (11)$$

In the case of computationally phased data, there may exist phasing errors that would decrease the accuracy of the estimated effect sizes ($\hat{\phi}$). With imperfect phasing, the observed phasing $\hat{\mathbf{v}}_i$ may differ from the true phasing \mathbf{v}_i , a modified equation may be used to calculate the AS z-score given the per-SNP probability of mis-phasing ψ_i :

$$\hat{z}_{\phi,i} = \frac{\hat{\phi}_i}{\sqrt{\left(\hat{\mathbf{v}}_i^\top \mathbf{\Omega} \hat{\mathbf{v}}_i \right)^{-2} \left(\left(\hat{\mathbf{v}}_i^\top \mathbf{\Omega}^2 \hat{\mathbf{v}}_i \right) \left(\hat{\sigma}_{w,i}^2 + 4\psi_i \hat{\phi}_i^2 \right) + \hat{\mathbf{v}}_i^\top \mathbf{\Omega} \hat{\mathbf{v}}_i \right)}} \quad (12)$$

Fine-mapping simulation results under imperfect phasing are presented in the Appendix and in Figure S5.

2.3 Inference of credible sets and posterior probabilities

PLASMA defines a joint generative model for total (QTL) and haplotype-specific (AS) effects on expression. Let $\hat{\mathbf{z}}$ be the combined vector of AS association statistics and QTL association statistics:

$$\hat{\mathbf{z}} = \begin{bmatrix} \hat{\mathbf{z}}_\phi \\ \hat{\mathbf{z}}_\beta \end{bmatrix} \quad (13)$$

Let \mathbf{R}_z be the genotype LD matrix, and $r_{\beta\phi}$ be a hyperparameter describing the overall correlation between the QTL and AS summary statistics calculated across all loci. Let the combined correlation matrix \mathbf{R} as:

$$\mathbf{R} = \begin{bmatrix} \mathbf{R}_z & r_{\beta\phi}\mathbf{R}_z \\ r_{\beta\phi}\mathbf{R}_z & \mathbf{R}_z \end{bmatrix} \quad (14)$$

The joint z-scores are modeled under a multivariate normal distribution, with covariance \mathbf{R} :

$$\hat{\mathbf{z}} \sim \mathcal{N}_{2m}(\mathbf{z}, \mathbf{R}) \quad (15)$$

PLASMA utilizes a likelihood function that gives the probability of statistics $\hat{\mathbf{z}}$, given a causal configuration. Let a causal configuration \mathbf{c} be a vector of causal statuses corresponding to each marker, with 1 being causal and 0 being non-causal. PLASMA assume that the causal configuration is the same for the QTL and AS signals.

Let hyperparameters $\sigma_{c,\phi}^2$ and $\sigma_{c,\beta}^2$ be the variance of AS and QTL causal effect sizes, respectively. Furthermore, let the jointness parameter $r_{c,\beta\phi}$ be the underlying correlation of the causal QTL and AS effect sizes. (This is not to be confused with $r_{\beta\phi}$, which concerns the correlation between the observed association statistics. See the Appendix for a mathematical relationship between these two hyperparameters.) These three hyperparameters are closely related to the heritability of gene expression (see Appendix). Let $\Sigma_{\mathbf{c}}$ be the covariance matrix of causal effect sizes given a causal configuration:

$$\Sigma_{\mathbf{c}} = \begin{bmatrix} \text{diag}(\mathbf{c})\sigma_{c,\phi}^2 & \text{diag}(\mathbf{c})r_{c,\beta\phi}\sigma_{c,\phi}\sigma_{c,\beta} \\ \text{diag}(\mathbf{c})r_{c,\beta\phi}\sigma_{c,\phi}\sigma_{c,\beta} & \text{diag}(\mathbf{c})\sigma_{c,\beta}^2 \end{bmatrix} \quad (16)$$

PLASMA’s likelihood for a causal configuration is defined as:

$$\mathcal{L}(\mathbf{c}; \hat{\mathbf{z}}) = \mathcal{N}_{2m}(\mathbf{0}, \mathbf{R} + \mathbf{R}\Sigma_{\mathbf{c}}\mathbf{R}) \quad (17)$$

Let γ be the prior probability that a single variant is causal and $1 - \gamma$ as the probability that a variant is not causal. The prior probability of a configuration consisting of m variants is defined as:

$$P(\mathbf{c}) = \prod_{i=1}^m \gamma^{c_i} (1 - \gamma)^{1-c_i} \quad (18)$$

With the prior and likelihood, the posterior probability of a causal configuration, normalized across the set of all possible configurations \mathbb{C} is calculated as:

$$P(\mathbf{c}|\hat{\mathbf{z}}) = \frac{P(\hat{\mathbf{z}}|\mathbf{c})P(\mathbf{c})}{\sum_{\mathbf{c}^* \in \mathbb{C}} P(\hat{\mathbf{z}}|\mathbf{c}^*)P(\mathbf{c}^*)} \quad (19)$$

PLASMA defines the ρ -level credible set \mathbb{K} as the smallest set of markers with a ρ_c probability of including all causal markers. Let $\mathbb{C}_{\mathbb{K}}$ be the set of all causal configurations whose causal markers is a subset of \mathbb{K} , excluding the null set. The credible set confidence level ρ_c is calculated as the sum of the probabilities of the configurations in $\mathbb{C}_{\mathbb{K}}$:

$$\rho_c = \sum_{\mathbf{c} \in \mathbb{C}_{\mathbb{K}}} \Pr(\mathbf{c}|\hat{\mathbf{z}}) \quad (20)$$

Additionally, PLASMA defines a marker’s posterior inclusion probability (PIP) as the probability that a single given marker is causal, marginalized over all other markers. This probability is calculated by summing over all configurations containing the marker.

To reduce the number of configurations to evaluate in the case of multiple causal variants, PLASMA uses the heuristic that configurations with significant probabilities tend to be similar to each other. PLASMA uses a shotgun stochastic search procedure to find all configurations with a significant probability. For each iteration of the algorithm, the next configuration is drawn randomly from the neighborhood of similar configurations, weighted by the posterior probability of each candidate. The search is terminated under the presumption that all configurations with nonzero probability have been uncovered.

Given the large number of configurations evaluated, it is impractical to calculate the best possible credible set satisfying ρ_c . Instead, PLASMA uses a greedy approximation algorithm. At each step, before ρ_c is reached, the algorithm adds the marker that increases the confidence the most.

2.4 The Jointness Parameter in PLASMA

Although PLASMA always assumes the same causal variants for QTL and AS, the correlation between QTL and AS causal effects can be set in PLASMA-J with a jointness hyperparameter $r_{c,\beta\phi}$. A high value (near 1) assumes that the QTL and AS causal effects tend to be consistent in magnitude, while a low value (near zero) assumes more disparity. Note that this is unrelated to choice of causal variants, and PLASMA assumes that QTL and AS share the same causal variants regardless of the jointness parameter.

A previous analysis comparing QTL effects with a similar formulation of AS effects has uncovered a highly nonlinear relationship, especially with QTL effects calculated using untransformed total expression data [22]. As a further complication, this relationship between QTL and AS effects is shown to be highly dependent on allele frequency. Thus, even under the assumption that QTL and AS signals share a causal variant, there is no guarantee of a strong linear correlation between QTL and AS effect sizes. Due to this uncertainty, the jointness parameter is zero by default, making no assumption on the relationship between QTL and AS effect sizes.

To empirically evaluate the effect of the jointness parameter on fine-mapping performance, PLASMA-J was run with different values of the jointness parameter on simulated loci. Figure S2 shows the distribution of PLASMA-J credible sets with different values of jointness, ranging from 0 to 0.99. In the one causal variant case, results are largely invariant to the parameter below a value of 0.99.

2.5 Generation of simulated loci

Genotype data was sampled from phased SNP data using the CEU population in the 1000 Genomes Project. First, a contiguous section of markers is randomly chosen. Next, a random selection of samples are randomly selected from the section. The genotypes corresponding to the chosen samples yield two haplotype matrices, denoted as \mathbf{H}_a and \mathbf{H}_b .

Among the markers, the desired number of causal markers is randomly selected. In the case of multiple causal variants, each causal marker is assigned a relative effect size, sampled from a normal distribution with zero mean and unit variance. For each individual, the ideal un-scaled gene expression for each haplotype \mathbf{q}_a and \mathbf{q}_b is determined by multiplying the relative effect sizes with each haplotype matrix.

Read count data is simulated with this haplotype-specific expression. In real data, only a fraction of the reads can be mapped to a specific haplotype. Due to this difference between total reads and mapped reads, the allelic imbalance and the total read count (QTL) are calculated separately.

To calculate total read count data, the total ideal un-scaled expression \mathbf{q}_t is defined as $\mathbf{q}_a + \mathbf{q}_b$, the sum of the haplotype-specific un-scaled gene expression. Gaussian-distributed noise is then added so that the variance of \mathbf{q}_t is consistent with the total variance across samples as specified by the QTL heritability. Finally, this final expression is scaled so that the total expression across samples is of unit variance. Total read counts are not explicitly generated, since a multiplicative factor across samples does not influence the QTL association statistics calculated by the model. This is reflective of typical QTL study protocols which aggressively rank/quantile normalize the data to fit a normal distribution.

To calculate allele-specific read counts, heritability, mean read coverage, and the total variance of the AS phenotype are taken into account. The ideal allelic imbalance phenotype is determined as $\text{logit}\left(\frac{\mathbf{q}_a}{\mathbf{q}_b}\right)$ (calculated element-wise). Gaussian-distributed noise is then added so that the signal-to-noise ratio of the phenotype’s variance is consistent with the specified AS heritability. This noisy phenotype is then scaled to the specified total variance. The read coverage for each sample is then drawn from a Poisson distribution, given the mean read coverage. Lastly, allele-specific read counts are generated from these phenotypes, with the counts for each sample being drawn from a beta-binomial distribution.

2.6 Comparison of existing models with PLASMA

Our analyses benchmark PLASMA against existing fine-mapping methods. Two distinct versions of PLASMA are tested, “PLASMA-J” and “PLASMA-AS.” The PLASMA-J (Joint-Independent) version looks at both AS and QTL statistics, assuming a shared set of AS and QTL causal variants,

and also that the AS and QTL causal effects are uncorrelated. The “PLASMA-AS” version is restricted to only AS data. As a baseline, we compare PLASMA to a QTL-Only version of PLASMA and to the CAVIAR method (expected to be equivalent to PLASMA QTL-Only) [13]. The behavior and performance of CAVIAR is representative of similar QTL-based methods such as CAVIARBF, FINEMAP, and PAINTOR without functional annotation data [14–16]. The versions of PLASMA are furthermore compared against the only other publicly-released fine-mapping method (to our knowledge) that integrates AS data described in the pre-print of Zou *et al.*, 2018 [24]. This unnamed method, denoted as “AS-Meta,” utilizes the association between SNP heterozygosity and a binary indicator of allelic imbalance. By binarizing allelic imbalance, AS-Meta is expected to lose power relative to treating imbalance as a quantitative phenotype but may be more robust to spurious AS signal. Furthermore, AS-Meta utilizes only indicators of heterozygosity, rather than marker phasing. AS-Meta can therefore be used with unphased genotypes, but at the expense of being unable to leverage the direction of the allelic effect. Lastly, as an additional comparison with an AS-based method, we analyze the performance of RASQUAL, a method for inferring allele-specific genetic effects using both allelic and total expression signal. Note that RASQUAL only computes allele-specific effect sizes for each marker and is not intended to compute credible sets or posterior marginal probabilities. Traditional fine-mapping on RASQUAL statistics is made possible by converting RASQUAL chi-square statistics back to quasi-z-scores with sign based on the direction of the RASQUAL effect-size. These statistics are then fed into standard QTL-only fine-mapping to obtain credible sets and posterior probabilities. We denote the modification of RASQUAL as “RASQUAL+”. This process is comparable to fine-mapping using a combined AS/QTL effect, rather than modeling QTL and AS effects separately.

2.7 Quality control of genotype data

For TCGA data, germline genotype calls are downloaded from the Genomic Data Commons. For PrCa ChIP samples, germline genotypes are called from blood (personal communication). Genotypes were imputed to the Haplotype Reference Consortium [25] using the Michigan Imputation Server [26] and restricted to variants with INFO greater than 0.9 and MAF greater than 0.01. Variants are further restricted to QC-passing SNPs from Ref. [25] which represent common, well-mapped variants from the 1000 Genomes project.

2.8 Quality control of RNA-seq data

Raw RNA-seq BAM files are downloaded from the Genomic Data Commons. Initial RNA-seq mapping and alignment are performed following TCGA parameters for the STAR aligner [27]. Mapping bias is accounted for by re-mapping using the WASP pipeline [19] and the STAR aligner with the same parameters. Reads are randomly de-duplicated as recommended by the WASP pipeline.

Somatic copy number calls are downloaded from FireBrowse and local beta-binomial overdispersion parameters are estimated for each contiguous region of copy number change.

2.9 Quality control of ChIP-seq data

ChIP-seq reads are aligned using bwa and default parameters [28], and peaks are called using MACS2 and default parameters (with DNA-seq input provided as control) [29]. Peaks are then unified across all samples. Mapping bias is accounted for by re-mapping using the WASP pipeline and the bwa aligner with the same parameters. Reads are randomly de-duplicated as recommended by the WASP pipeline. Beta binomial overdispersion parameters are estimated globally for each sample as somatic copy number was expected to be minimal.

2.10 Allele-specific quantification

The StratAS algorithm is used to quantify allele-specific signal and identify initially significant features for fine-mapping [23]. For each peak/gene (the feature) and individual all reads at heterozygous SNPs in the feature are aggregated to compute the haplotype-specific read counts, and summed across the two haplotypes of each individual to compute the QTL read counts. Each QC passing variant within 100kb of the feature are then tested for an allele-specific association with the feature and features significant at a genome-wide false discovery rate (FDR) of 5% are retained for fine-mapping.

2.11 Functional enrichment analysis

For QTLs fine-mapped from RNA-seq we select regions of accessible chromatin in the most relevant tissue as reference the functional feature, reasoning that high-confidence causal variants should be

more abundant in accessible regions. For QTLs fine-mapped from ChIP-seq we select chromosome looping anchors from Hi-ChIP in the relevant tissue as the reference functional feature, reasoning that high-confidence causal variants should be more abundant in regions that are in conformation with promoters.

Enrichment is then estimated by computing the proportion of markers in credible sets that intersect with the functional feature. Controls are calculated as the intersection between all tested markers and the functional feature. Odds ratios and p-values are computed with Fisher’s exact test.

3 Results

3.1 Simulation framework

We evaluate PLASMA with a framework that simulates the expression of whole loci in an allele-specific manner. This simulation framework jointly simulates total reads and allele-specific read counts, under given values of the number of causal variants, the QTL heritability, the AS heritability, the variance of the AS phenotype across samples, and the expected read coverage (see Methods). The variance and heritability of the AS phenotype are handled by two separate parameters, where the former describes the total spread of allelic imbalance, and the latter specifies the fraction of the variance that is due to genetic effects. This allows us to investigate cases where a significant amount of observed imbalance is caused by non-genetic variance in the allelic expression. To quantify the total variance of the AS phenotype in the population, we define the “standard allelic deviation” (d) as the standard deviation of the AS phenotype w , quantified on the allelic fraction scale (between 0.5 and 1). Importantly, this quantity is independent of the genetic effect, which is controlled by the heritability parameter. Simulations were performed using real phased haplotype data from the 1000 Genomes Project European samples. Parameter settings for simulation analyses are shown in Table S1.

As the performance of standard QTL association models is well established, we first focused on performance of our proposed AS statistic. Figure S3a shows how the mean z_ϕ varies as a function of standard allelic deviation and mean read coverage at a fixed AS heritability of 0.5. Second, Figure S3b shows how the mean z_ϕ varies as a function of standard allelic deviation and heritability with

mean coverage fixed at 100. The statistic is the greatest at high read coverage and high heritability, consistent with the degree of experimental and intrinsic signal available to the model. These results hold even at low AS variance ($d = 0.6$) and show that PLASMA does not conflate high AS variance (standard allelic deviation) with high signal (coverage or heritability). This robustness to variance in the AS phenotype makes the model resistant to false-positives driven by non-genetic sources of allelic variance. At very high variance ($d > 0.8$), z_ϕ shows a sharp decrease. This decrease in signal is due to an increase in the sampling error of the AS phenotype (w) at high overall variance, as shown in Equation 47 (See Appendix for a mathematical relationship between total variance and sampling error.)

3.2 Comparison with existing methods in simulation

First, we evaluate how well each PLASMA prioritizes candidate causal markers using simulated loci with one causal variant, compared to existing QTL and AS-based methods. We define the “inclusion curve” for each model, where markers are ranked by posterior probability and added one by one to a cumulative set (note that this set is not dependent on the definition of a credible set). The x axis represents the cumulative number of markers chosen, and the y axis represents the “inclusion rate,” the proportion of true causal markers among the chosen markers. Figures 2a and d show inclusion plots at low and high AS variance, respectively. As expected, FINEMAP, QTL-Only, and the CAVIAR methods are indistinguishable and do not vary with AS variance. Due to this similarity in results, FINEMAP is used as the primary QTL-based methods in subsequent analyses. Furthermore, PLASMA-J and PLASMA-AS perform similarly at both levels of AS variance. Additionally, AS-Meta’s performance exhibits a dependency on the degree of AS variance. Lastly, RASQUAL+ at high AS variance does significantly improve over QTL-based methods, but not as well as PLASMA. At low AS variance (with same amount of signal and noise), RASQUAL+ performs considerably worse, indicating that RASQUAL+ is more sensitive to the genetic architecture of the locus than PLASMA is.

Second, we evaluate the ability of each model to rule out likely non-causal markers in simulated loci with one causal variant. We conduct a direct comparison of the distributions of the 95% confidence credible sets, with smaller sets indicating higher specificity. Figures 2a and d show distribution plots at low and high AS variance, respectively. At low variance, PLASMA-J offers

the smallest median credible 95% set size (3.0), followed by PLASMA-AS (3.0), then AS-Meta (55.0), and lastly the QTL-based methods: FINEMAP (89.0), CAVIAR (89.0), and QTL-Only (91.0). There is some variation due to differences in calibration among the methods, but all QTL-based methods have recall above 0.95. PLASMA appears robust to changes in AS variance; at high AS variance, medians are 3.0 for PLASMA-J and 3.0 for PLASMA-AS. In contrast, the performance of AS-Meta varies significantly with the degree of AS variance, even when the underlying signal (coverage and heritability) is constant, with a median set size of 79.0 at high variance. This sensitivity may be due to the fact that AS-Meta does not incorporate marker phasing, and thus must rely solely on the intensity of imbalance, rather than the direction of imbalance. Here, RASQUAL+ does not generate meaningful credible sets, with 95% credible set recall being 0.06 and 0.58 for low and high AS variance, respectively. RASQUAL+ is therefore not included in further fine-mapping analyses, though we underscore that RASQUAL remains an effective tool for QTL discovery.

Third, we directly compare how PLASMA-J and FINEMAP prioritize a common set of variants pooled from 500 loci, each with 100 total markers and one causal marker. Figure 2g shows a joint histogram of log posterior marginal odds of these 50000 variants, with causal variants highlighted in red. Distributions of posterior log-odds for each method are shown as univariate histograms along each axis. As expected, PLASMA and FINEMAP posterior log-odds are positively correlated. Comparing the distribution of the odds of causal variants to those of the rest, it is furthermore evident that PLASMA more confidently assigns probabilities of causality, and can much more cleanly segregate causal from non-causal variants.

Lastly, we run the AS-based methods across a wide range of coverage and heritability conditions, recording the mean 95% confidence credible sets, shown in Figure 3. Figures 3a-c show mean credible set sizes as a function of AS variance and coverage, and Figures 3d-f show mean credible set sizes as a function of AS variance and AS heritability. In terms of the range of set sizes, PLASMA-J performs the best (3.2 markers on average at best conditions), followed by the PLASMA-AS (3.4 at best conditions), and lastly the AS-Meta method (31 at best conditions). Generally speaking, all methods show results consistent with the behavior of z_ϕ in Figure S3. Although increasing either coverage or heritability results in smaller set sizes, increasing coverage beyond 100 gives diminishing returns as the observed expression levels approach the true expression levels. As expected, AS-Meta

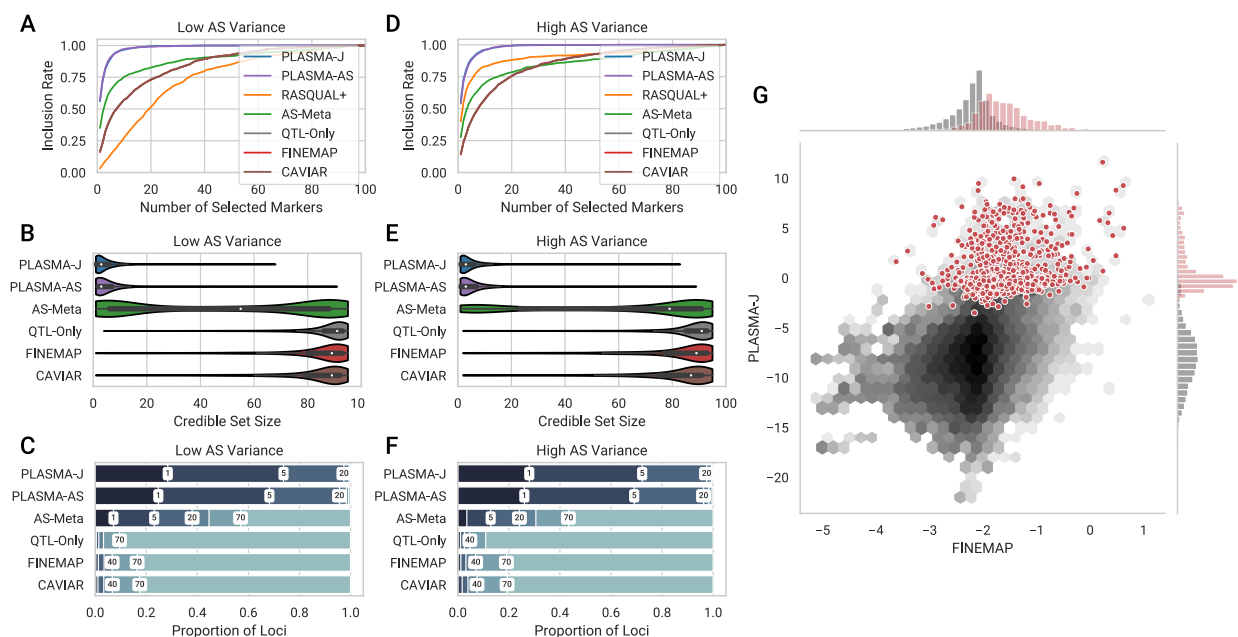


Figure 2: Comparison of fine-mapping methods on two sets of simulated loci of 100 markers and 100 samples each. These two sets differ in the variance of the AS phenotype, but with a fixed mean coverage and AS heritability. (a, d) Inclusion curve at low and high AS variance, respectively (0.6, 0.8 standard allelic deviation). The y-axis shows the expected proportion of markers included in the credible set, and the x-axis shows the number of selected markers by posterior probability. (b, e) Distribution of 95% confidence credible set sizes at low and high AS variance, respectively. (c, f) Proportions of loci with 95% credible set sizes within given thresholds at low and high AS variance, respectively. Thresholds used are 1, 5, 20, 40, 70, and 100 markers. (g) A per-snp fine-mapping comparison between PLASMA-J and FINEMAP across 500 simulated loci with one causal variant each. The axes denote the posterior log-odds of causality for FINEMAP and PLASMA-J, respectively. The black hexagons represent the joint distribution of all markers, while the red dots represent specifically the causal markers. Univariate histograms for PLASMA-J and FINEMAP are plotted along the margins.

tends to struggle at low AS variance, especially apparent at a standard allelic deviation of 0.55, with a mean set size of 78 at best. This may be due to the large majority of samples falling under the threshold for allelic imbalance at 0.65. To verify that PLASMA is calibrated across the full range of conditions, Figure S4 shows that the 95% credible set sizes have at least a 95% chance of including the causal variant.

3.3 Inference of multiple causal variants

To demonstrate PLASMA beyond a one-causal-variant assumption, we fine-mapped sets of simulated loci with 2 causal variants with each version of PLASMA. Figure 4a shows the inclusion curves for each version of PLASMA along with FINEMAP. For these curves, inclusion is defined as the expected proportion of causal variants selected, where an inclusion of 1.0 indicates the identification of both causal variants. Here, PLASMA-J and PLASMA-AS deliver an improvement over conventional QTL fine-mapping. Compared to single causal variant fine-mapping, all methods display a decrease in power, which is consistent with results in earlier QTL fine-mapping analysis [13, 14], where capturing all causal variants becomes increasingly difficult as the number of causal variants increase. The lower power for fine-mapping multiple causal variants may be due to the stringent requirement that a model must identify all causal variants in a locus for an inclusion of 1.0. To evaluate the ability of the models to detect the top causal variant, we relax this requirement from identifying all causal variants per locus to at least one causal variant per locus. Inclusion plots for this scenario are shown in Figure 4b, with PLASMA greatly improving the prioritization of the lead causal variant over existing methods.

We next considered credible set sizes which, unlike the inclusion curves, require accurate calibration to be comparable. Previous analyses have shown that proper calibration of fine-mapping methods is more challenging in the presence of multiple causal variants [30]. Unlike the single causal variant case, where all PLASMA model hyperparameters were inferred from simulation parameters, the causal variance hyperparameters in this case were manually calibrated. This need for calibration may be due to linkage disequilibrium obfuscating the relationship between causal effect sizes and total heritability at a locus, and further complicated by the imperfect estimation of linkage disequilibrium at 100 samples [31]. (See Appendix for information about hyperparameter estimation.) The PLASMA results shown in this section are calibrated such that the recall rates

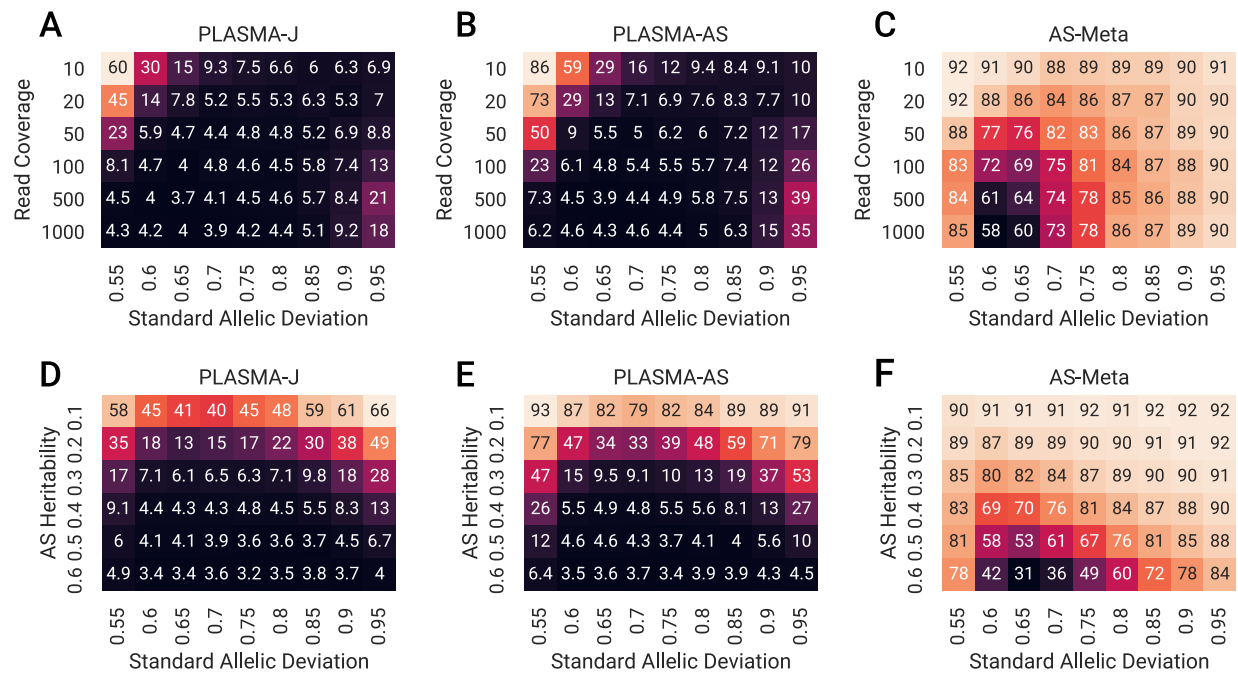


Figure 3: Comparison of the mean 95% confidence credible set sizes across different AS fine-mapping methods. Each square is the mean set size calculated over 500 simulated loci of 100 markers and 100 samples each, with one causal variant. (a, b, c) Mean credible set sizes as a function of standard allelic deviation and mean read coverage, with AS heritability set to 0.4. (d, e, f) Mean credible set sizes as a function of standard allelic deviation and AS heritability, with mean read coverage set to 100 reads.

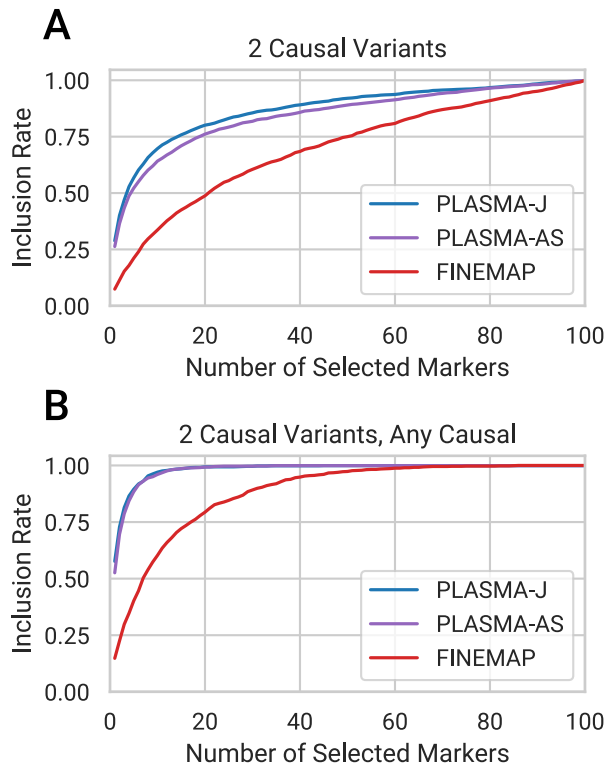


Figure 4: Comparison of fine-mapping methods on sets of simulated loci with two causal variants. Each locus is of 100 markers and 100 samples each, with AS heritability set to 0.4, QTL heritability set to 0.05, and mean read coverage set to 100 reads. (a) Distribution of 95% confidence credible set sizes. (b) Inclusion curve for including both causal variants. (c) Inclusion curve for including one or both causal variants.

for 95% confidence credible sets are 0.95, 0.96 for PLASMA-J and PLASMA-AS, respectively. This calibration yields median credible set sizes of 86.0 and 90.0 for PLASMA-J and PLASMA-AS, respectively. Like PLASMA, FINEMAP requires user-defined hyperparameters on the prior number of causal variants and on the causal effect sizes. These FINEMAP parameters were set to be equivalent to corresponding calibrated PLASMA parameters. Despite this conservative parameter setting, FINEMAP is overconfident on this dataset with a recall rate of 0.86, so the generated credible sets for FINEMAP are not directly comparable to those of PLASMA.

3.4 Fine-mapping of TCGA kidney RNA-Seq data

To evaluate our method on real data, we fine-mapped gene expression data from 524 human kidney tumor samples and 70 matched normal samples collected by TCGA [32]. The data was processed through a rigorous QC pipeline to account for mapping biases based on established best practices

[19, 22]. Figures 5a and c show credible set size distribution plots for tumor and normal data, respectively, under a 1 causal variant assumption. We furthermore analyze how well often each method is able to narrow down credible sets under a certain size in simulated loci with one causal variant, shown in Figures 5b and d. Among the tumor samples ($N=524$, 5652 loci), 27.9% of loci are fine-mapped within 10 variants with PLASMA-J, while 3.4% of loci are fine-mapped within 10 variants with FINEMAP. Furthermore, 263 of these loci can be fine-mapped down to a single causal variant by PLASMA-J. PLASMA-J moreover achieves a median credible set size for 32 variants, whereas FINEMAP achieves a median credible set size of 167 variants. FINEMAP also significantly improves over AS-Meta, which has 6.6% of loci fine-mapped within 10 causal variants, and a median credible set size of 166. Results for normal samples ($N=70$, 2034 loci) have a similar trend, with 23.2%, 2.5%, and 1.3% of loci fine-mapped within 10 causal variants, for PLASMA-J, AS-Meta, and FINEMAP respectively. Corresponding median credible set sizes are 32, 248, and 252 variants, for PLASMA-J, AS-Meta, and FINEMAP respectively. The somewhat lower power for all models is due to having fewer normal samples than tumor samples. To show that these credible set sizes are robust our choice of heritability hyperparameters, fine-mapping analyses were repeated on the full set of tumor genes with the AS heritability hyperparameter set to 0.05 instead of 0.5. A comparison of the credible set sizes with those from the original parameters are shown in Figure S6 .

To investigate how the methods perform at lower sample sizes, we randomly subsample individuals prior to fine-mapping. Figure 5e plots the credible set size distributions for PLASMA-J, AS-Meta, and FINEMAP at various sample sizes of kidney tumor data. In terms of loci fine-mapped to credible set sizes within 10 variants, PLASMA with 50 samples (484 loci within 10 causal variants) has significantly greater power than FINEMAP with 500 samples (193 loci within 10 causal variants) or AS-Meta with 500 samples (371 loci within 10 causal variants). Additionally, in terms of median credible set size, PLASMA with 10 samples (170 median) has about the same power as FINEMAP with 500 samples (167 median). At a given sample size, PLASMA is thus better able to prioritize variants that will be ranked highly in larger studies. Furthermore, as sample size increases, PLASMA increases in power relative to other methods. In tumor samples PLASMA yields a 1.3-fold decrease in median credible set size over FINEMAP at 10 samples, but a 6.9-fold decrease at 500 samples, indicating that PLASMA scales more effectively with sample

size than conventional QTL fine-mapping. Nevertheless, PLASMA yields a substantial reduction of credible set sizes even with sample sizes as low as 10, with a median credible set size of 170, compared to a median set size of 219 with FINEMAP. An analogous down-sampling analysis on the normal data is shown in Figure S7. There, PLASMA has higher power for normal samples than for tumor samples, which may be due to the lower variance in the normal data.

Next, we look at how causal variant prioritization is impacted by sample size in the down-sampled analysis. Because the true causal variants in each locus is not known, we use a proxy of markers with a posterior probability of at least 0.1 when fine-mapped with FINEMAP on all samples. Note that this will strongly bias the credible set in favor of FINEMAP and thus do not compare this proxy to FINEMAP credible sets. In Figure S8, PLASMA is again more effective than AS-Meta at each sample size at prioritizing causal variants.

Lastly, to explore multiple causal variant fine-mapping on real data, we run PLASMA and FINEMAP assuming up to three causal variants on the full tumor and normal kidney RNA-Seq dataset. Figure S9 shows multiple causal variant fine-mapping results for kidney tumor and normal RNA-Seq data. As with the simulations, all methods increase in credible set sizes relative to single-causal-variant fine-mapping. On tumor data, PLASMA-J, PLASMA-AS, and FINEMAP report a median credible set size of 93, 172, and 150 respectively, with the caveat of possibly unstable calibration for multiple causal variants (as seen in simulations). Interestingly, PLASMA-AS displays a larger power drop than FINEMAP does. This difference suggest that allelic imbalance may be less informative when fine-mapping with multiple causal variants. Nevertheless, PLASMA-J performs substantially better than either, suggesting that the joint model is able to combine power from both QTL and AS signals. Regardless, it appears that the majority of loci contain a single causal variant, with FINEMAP estimating this fraction at 68.8%.

Lastly, we look at how PLASMA prioritizes experimentally-verified causal variants at GWAS risk loci. Figure 6 shows the strength AS and QTL associations for DPF3 and SCARB1, genes in two kidney GWAS loci that have verified causal variants [23, 33]. At each sample size threshold, the AS statistic more confidently identifies the true causal variant than the QTL statistic. In the case of DPF3, the AS statistic is able to prioritize the true causal variant at a substantially lower sample size than the QTL statistic. Moreover, the 95% credible sets from the PLASMA-AS model are smaller than those from the QTL-Only model at a given sample size. By producing a more

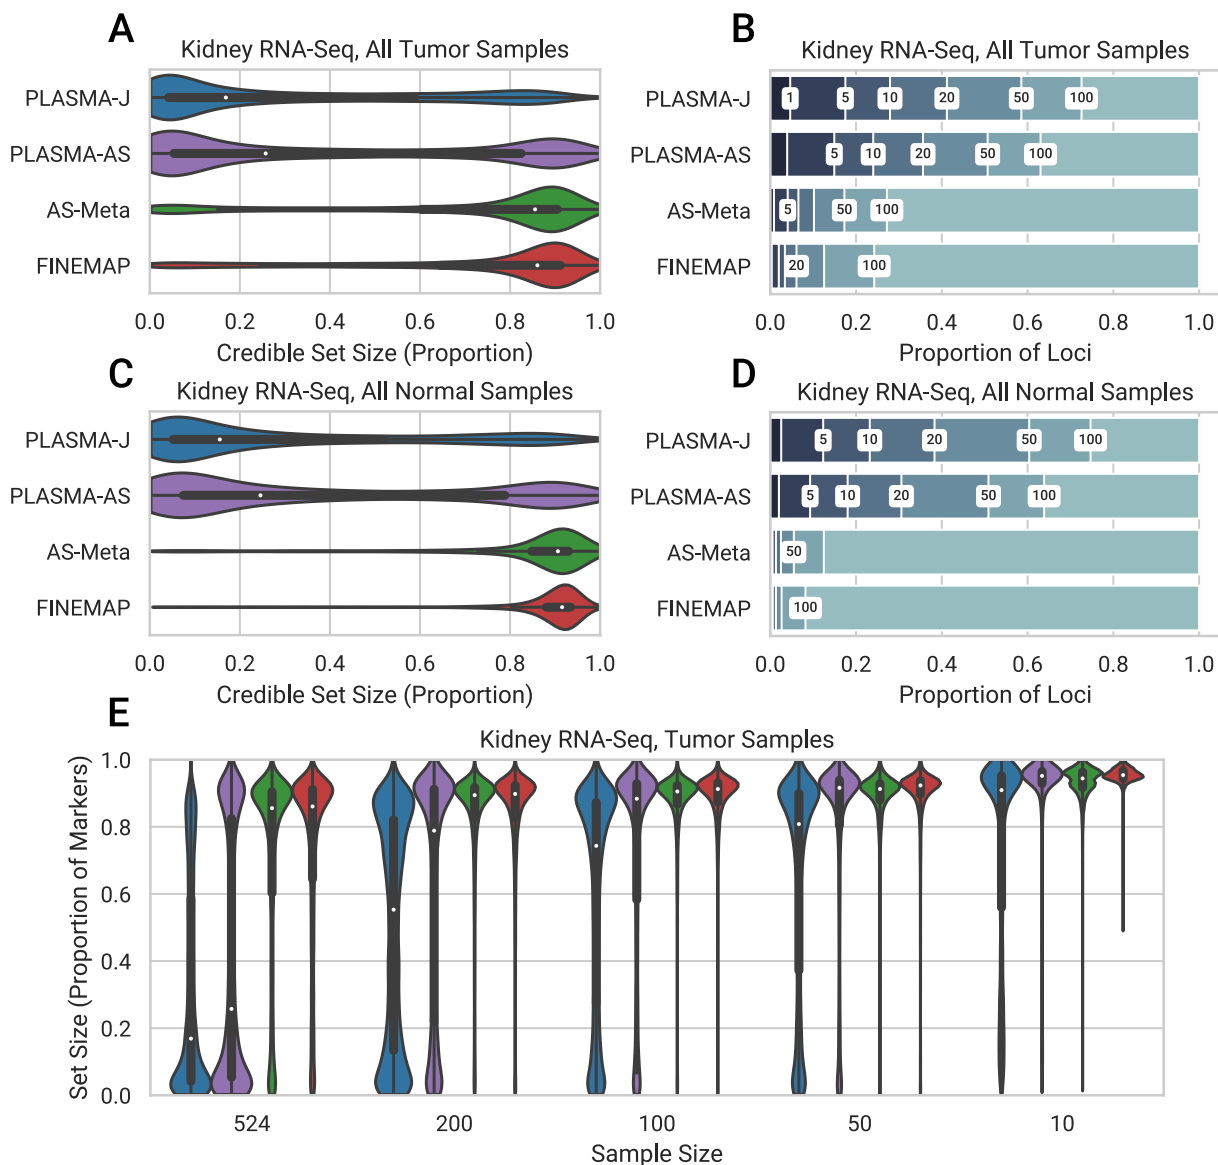


Figure 5: Comparison of the distribution of 95% confidence credible set sizes across loci in kidney tumor and normal samples, with an allelic imbalance false discovery rate of 0.05. Expression was measured as RNA-Seq read counts mapped to phased genotypes. Fine-mapping was conducted with 500 tumor samples and 70 normal samples. Every SNP within 100 kb of each locus was included in the fine-mapping input. (a, c) Distribution of credible set sizes for tumor and normal samples, respectively, under a 1-causal-variant assumption. (b, d) Proportion of loci fine-mapped to 95% credible sets within given thresholds. Thresholds used are 1, 5, 10, 20, 50, and 100. (e) 95% credible set size distributions for randomly down-sampled kidney tumor datasets with decreasing sample sizes.

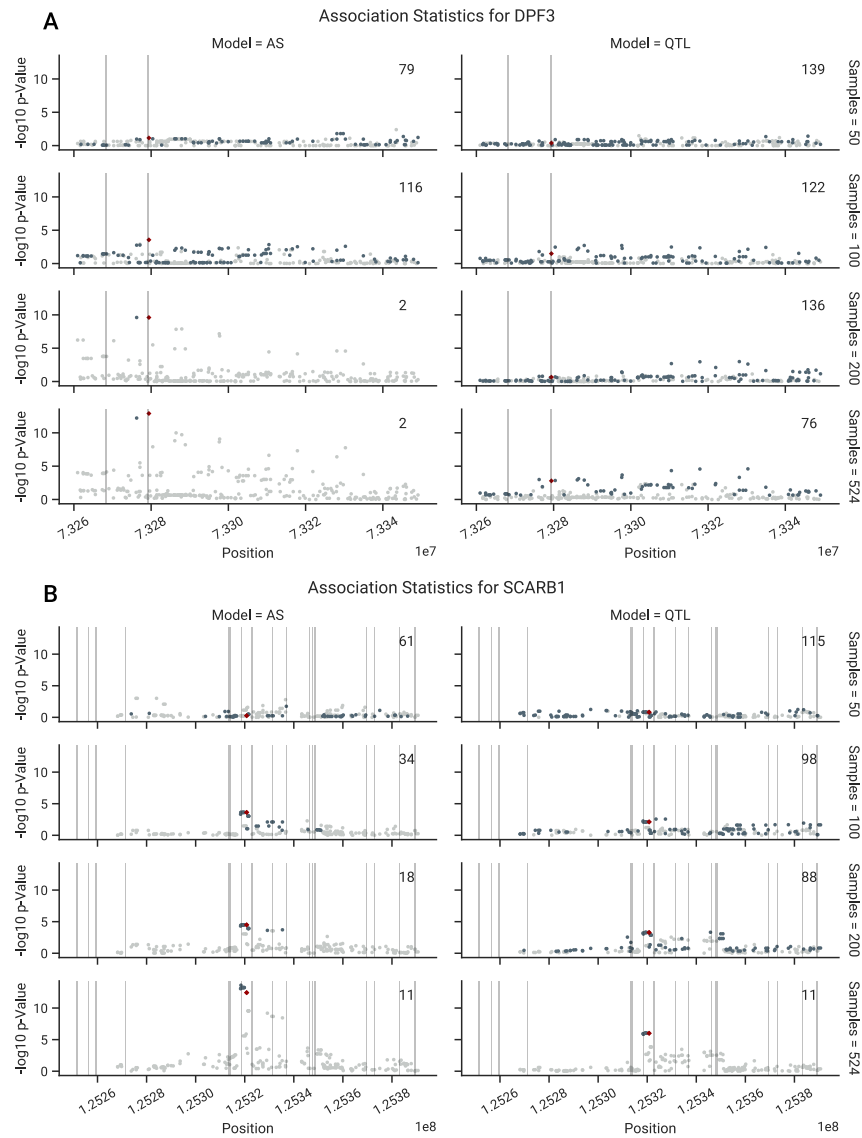


Figure 6: Comparison of AS and QTL associations in the experimentally-verified loci, as a function of sample size. Regions of 70 kb are shown around the causal marker. AS and QTL p-Values were calculated using from z_ϕ and z_β , respectively. Fine-mapping was conducted with the PLASMA-AS and QTL-Only models, respectively. 95% credible set sizes for the whole locus are displayed in the top right of each subplot. Markers in the 95% credible set are shown in dark gray, while markers not in the sets are shown in light gray. The experimentally verified causal marker is shown in red. Regions of open chromatin (DNase-Seq peaks) are shaded in gray. (a) Associations in the DPF3 locus (b) Associations in the SCARB1 locus.

accurate and confident prioritization of causal variants, PLASMA can substantially reduce the difficulty of experimentally validating causal variants.

3.5 Fine-mapping of prostate H3k27ac ChIP-Seq data

To evaluate PLASMA with a different molecular phenotype, we fine-mapped H3k27ac activity measured by ChIP-seq from 24 human prostate tumor samples and 24 matched normals. Although this study measures chromatin activity rather than expression, the nature of the data is nearly identical to that of RNA-Seq and is processed analogously by our QC pipeline and by PLASMA. Instead of fine-mapping eQTLs around gene loci, we fine-mapped chromatin QTLs (cQTLs) around chromatin peaks. Figure 7 shows distribution plots for tumor data (1375 peaks) and normal data (908 peaks) under a 1 causal variant assumption. Among the tumor data, 14.5% of peaks are fine-mapped within 50 variants with PLASMA-J, while 1.9% of loci are fine-mapped within 50 variants with FINEMAP. Furthermore, PLASMA achieves a median credible set size of 236, compared to QTL-Only fine-mapping achieving a size of 318. PLASMA also outperforms AS-Meta, with 1.9% of loci fine-mapped within 50 causal variants (no gain over FINEMAP), and a median credible set size of 310. Results from normal samples are comparable, with 5.2%, 2.5%, and 2.3% of loci fine-mapped within 50 causal variants, for PLASMA-J, AS-Meta, and FINEMAP respectively. These methods achieve a median credible size of 296, 313, and 319 variants, respectively. Overall, these ChIP fine-mapping results are roughly in line with those from RNA-Seq fine-mapping.

3.6 PLASMA increases functional enrichment of credible set markers

To evaluate PLASMA’s ability to select markers in functional regions using kidney RNA-Seq data, we look for enrichment of prioritized variants at open chromatin regions measured with DNase-Seq in a kidney cell-line [34]. Since chromatin accessibility is an indicator of transcription factor binding and regulation [35], an enrichment of credible set markers for open chromatin would indicate that the fine-mapping procedure is prioritizing markers in functionally relevant regions. For instance, the causal variant in the DPF3 locus lies within a DNase-Seq peak (Figure 6a). Note that quantifying overlapping with an independent functional feature such as open chromatin imposes no assumptions on the ground truth, in contrast to comparing to external QTL/GWAS data which may be biased towards conventional QTL analysis. The null distribution is defined as the credible set markers

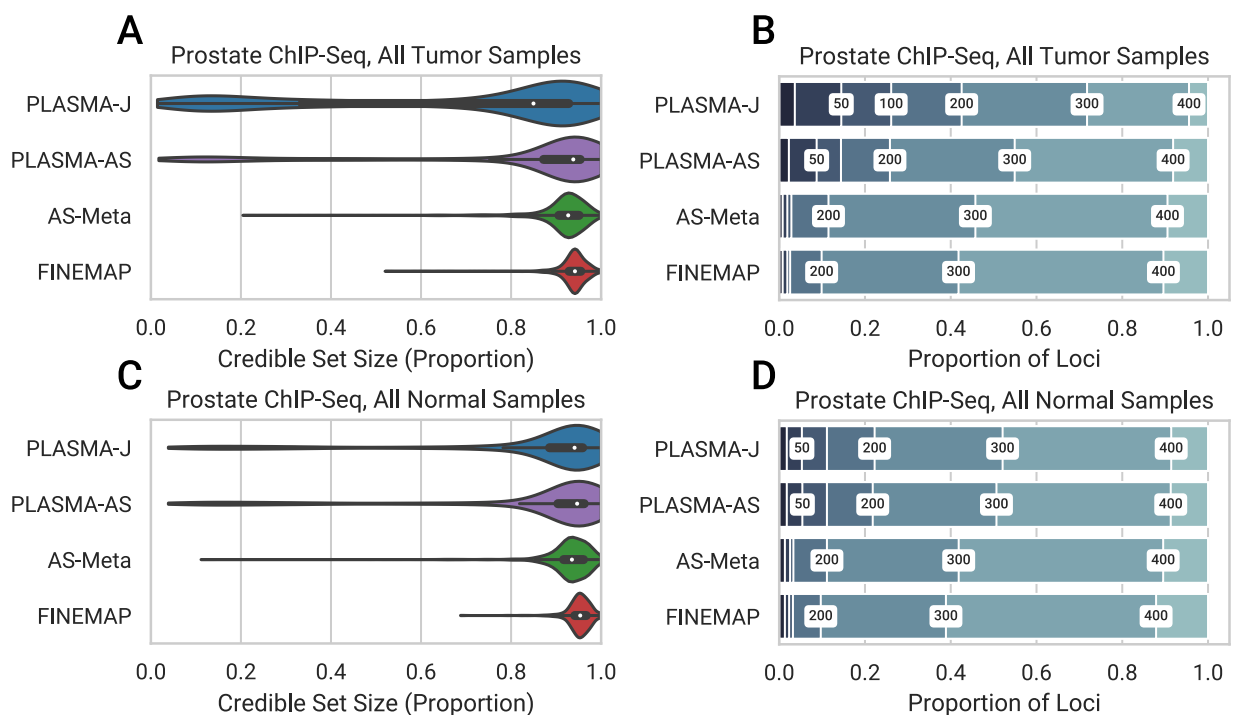


Figure 7: Comparison of 95% confidence credible set sizes for peaks in prostate tumor and normal cells, with an allelic imbalance false discovery rate of 0.05. Presence of H3k27ac histone marks was quantified as ChIP-Seq read counts mapped to phased genotypes. (a, c) Distribution of credible set sizes for tumor and normal samples, respectively, under a 1-causal-variant assumption. (b, d) Proportion of peaks fine-mapped to 95% credible sets within given thresholds. Thresholds used were 20, 50, 100, 200, 300, and 400.

being located independently of open chromatin and use Fisher’s exact test to calculate enrichment as a function of minimum causal variant probability. Figures 8, S10a and S10c,d show the odds ratios and p-values (computed by Fisher’s exact test), respectively, as a function of posterior probability threshold from each fine-mapping method. The credible set markers produced by PLASMA, for the most part, display a significantly stronger enrichment with open chromatin compared to existing methods. For instance, at the $p = 0.1$ threshold for tumor samples, PLASMA’s credible set markers achieve a p-value of 9.26×10^{-52} and an odds ratio of 2.16. In comparison, credible sets from QTL-Only fine-mapping at that threshold achieves a p-value of 2.02×10^{-7} and an odds of 1.62. This enrichment shows that even with far smaller credible sets, PLASMA is able to prioritize markers that fall in regions of likely functional significance. The difference between PLASMA and existing methods is greatest at higher posterior probability thresholds. PLASMA may be assigning a more meaningful number of markers with such high posterior probabilities, compared to existing methods that are rarely so confident about a marker’s causal status.

Similarly, to validate the credible sets computed from prostate ChIP-Seq data, we look for enrichment of credible set markers at chromatin looping anchors measured by Hi-ChIP in a prostate cell-line. Regulatory elements overlapping loops are more likely to be involved in cis-regulation and we reasoned that they should therefore be enriched for true causal cQTLs [36, 37]. Again, we note that this functional feature is independent of the QTL signal or locus LD and is not biased towards a QTL or AS model. Figure S10b and S10e show the odds ratios and p-values, respectively, across models as a function of posterior probability threshold (computed by Fisher’s exact test). The credible set markers produced by PLASMA display a significantly stronger enrichment with looping anchors compared to the other methods. For instance, at the $p = 0.1$ threshold, PLASMA’s credible sets achieve a p-value of 1.05×10^{-6} and an odds of 1.77. In contrast, credible set markers from FINEMAP at that threshold achieves a non-significant p-value of 0.80 and an odds of 0.72.

4 Discussion

We present PLASMA, a statistical fine-mapping method that utilizes allele-specific expression and phased genotypes to select candidate causal variants. By modeling gene expression at a locus in an allele-specific manner, PLASMA scales in power both across individuals and across read

counts. Through read-count-level simulations of loci, we show that PLASMA performs robustly across a wide range of realistic conditions and consistently outperforms existing statistical fine-mapping methods, including cases where a significant amount of observed imbalance is caused by non-genetic factors. We further demonstrate this increased power on experimental data by applying PLASMA to a large RNA-Seq study, as well as a smaller ChIP-Seq study. In both cases, PLASMA achieves substantially smaller credible set sizes compared to existing fine-mapping methods, greatly increasing the number of loci amenable to experimental causal variant validation. Lastly, we show that even with these greatly reduced (more specific) credible set sizes, PLASMA achieves an equivalent or superior degree functional enrichment as existing methods. These results not only present PLASMA as a powerful tool for prioritizing causal variants, but also demonstrate how AS analysis can be directly integrated into statistical fine-mapping. A key benefit of PLASMA is its ability to utilize existing, conventional sequencing-based QTL data, such as RNA-Seq, ChIP-Seq, and ATAC-Seq at low sample size. This allows researchers to gain significant insight simply by revisiting past QTL studies, especially those with sample sizes too low for conventional QTL fine-mapping.

Although it is evident that an AS analysis with PLASMA confers more signal than an equivalently-sized QTL analysis, AS analysis presents additional obstacles and potential confounders. First, unlike conventional QTL fine-mapping methods that rely only on allelic dosage, PLASMA additionally utilizes genotype phasing, making phasing accuracy a potential concern. However, since PLASMA focuses on cis-regulation, the genotypes observed span no more than several hundred kilobases per locus, well within the high accuracy range of modern phasing algorithms [38]. Second, PLASMA depends on having heterozygous individuals in the tested feature and SNP in order to leverage AS signal. In our analyses we focused on features that were testable by AS (10946 of 19645 total genes, 113459 of 525629 total peaks). However, even in the complete absence of heterozygotes, PLASMA can still conduct conventional fine-mapping based on dosage and total expression. Recent technologies that could potentially offer greater signal include RNA-seq with unspliced transcripts [39], and direct allele-specific measurement of expression using single-cell RNA-Seq [40]. Third, PLASMA assumes the same causal configuration underlying both the AS and QTL effects (and is thus able to combine the signals) but the causal effects may differ due to real biological confounding. For example, cis effects on gene *A* followed by (local) trans effects of gene *A* on gene *B* would be identi-

fied as a QTL association, but would not exhibit AS association. This would be a model violation for PLASMA and produce larger credible set sizes. Although PLASMA can consider correlations between causal AS and QTL affect sizes, this parameter is hard to estimate, and we find in real data that the model with correlation set to zero (PLASMA-J) exhibited greater power than a non-zero constant. Future work is required to fully elucidate the relationship between allele-specific and total effects, which likely differs across genes. Fourth, genomic imprinting (where either the maternal or paternal copy of the gene is silenced) or random monoallelic expression would produce the appearance of allelic imbalance within affected individuals in the absence of true cis-regulatory signal [20]. Although PLASMA does not explicitly model such biases, a bias that is independent of genotype will only cause a reduction in power and not produce false-positives. A potential extension would be to model such violations or discrepancies between the QTL and AS models directly, following the lines of methods such as RASQUAL [20]. Fifth, PLASMA currently does not incorporate covariate analysis in the allele-specific model (though the intra-individual nature of the test controls for false positives), which could additionally be used to model environmental confounders and increase power [41]. AS covariate analysis could potentially be achieved through a multivariate likelihood ratio test as in WASP [19].

PLASMA’s approach in combining QTL and AS signals opens up possible future work in two distinct directions. The first direction would be to build upon the generative fine-mapping model to incorporate additional sources of signal. For example, one can incorporate epigenomic annotation data by setting the priors for causality for each marker. Approaches used in existing QTL-based methods such as PAINTOR and RiVIERA-MT [16, 42] could be transferred to PLASMA with relatively little difficulty. Another possibility would be to conduct N -phenotype colocalization by utilizing additional phenotypes in addition to the AS and QTL phenotypes. Generalizing from two to multiple phenotypes would be straightforward, and could utilize the colocalization algorithm first introduced in eCAVIAR [2]. A second, more general direction would be to adapt QTL-based population genetics methods to utilize AS summary statistics. Since both QTL and AS statistics can be characterized as linear combinations of haplotype-level genotypes, they share many distributional properties, including LD, allowing them to be easily interchangeable in many circumstances. One such application would be gene expression prediction for transcriptome-wide association studies (TWAS) [43], where the increased signal of AS statistics could increase power to identify gene-

phenotype relationships. Broadly speaking, the allele-specific model and association statistics that PLASMA introduces will be relevant to any analysis of small sample size or limited tissue.

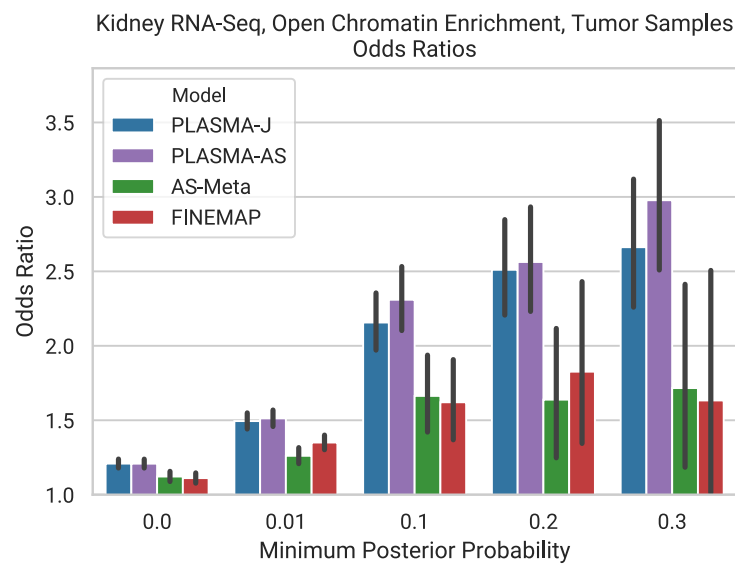


Figure 8: Enrichment of PLASMA kidney RNA-Seq credible sets on open chromatin. The x-axis represents a minimum posterior probability threshold on the markers within credible sets, with a higher threshold indicating more confident but fewer markers. The markers in these thresholded credible sets were intersected with the functional feature in question. 95% confidence intervals were calculated with Fisher's exact test, with a control of all markers in a locus.

Acknowledgments

We thank A. Shetty, E. O'Connor, C. Bell, M. Pomerantz, and M. Freedman for contributions to data collection and processing. We thank F. Hormozdiari for guidance on statistical fine-mapping and C. Kalita for guidance on model validation. We also thank B. Pasanuic, C. Giambartolomei, and M. Kellis for helpful feedback.

References

- [1] Zhu Z, Zhang F, Hu H, Bakshi A, *et al.* Integration of summary data from GWAS and eQTL studies predicts complex trait gene targets. *Nat Genet*, 48(5):481–487, 2016.
- [2] Hormozdiari F, van de Bunt M, Segre AV, Li X, *et al.* Colocalization of GWAS and eQTL Signals Detects Target Genes. *Am J Hum Genet*, 99(6):1245–1260, 2016.
- [3] Ongen H, Brown AA, Delaneau O, Panousis NI, *et al.* Estimating the causal tissues for complex traits and diseases. *Nat Genet*, 49(12):1676–1683, 2017.
- [4] Lappalainen T. Functional genomics bridges the gap between quantitative genetics and molecular biology. *Genome Res*, 25(10):1427–1431, 2015.
- [5] Battle A, Brown CD, Engelhardt BE, Montgomery SB, *et al.* Genetic effects on gene expression across human tissues. *Nature*, 550(7675):204–213, 2017.
- [6] Pickrell JK, Marioni JC, Pai AA, Degner JF, *et al.* Understanding mechanisms underlying human gene expression variation with RNA sequencing. *Nature*, 464(7289):768–772, 2010.
- [7] Veyrieras JB, Kudaravalli S, Kim SY, Dermitzakis ET, *et al.* High-resolution mapping of expression-QTLs yields insight into human gene regulation. *PLoS Genet*, 4(10):e1000214, 2008.
- [8] Lappalainen T, Sammeth M, Friedlander MR, 't Hoen PA, *et al.* Transcriptome and genome sequencing uncovers functional variation in humans. *Nature*, 501(7468):506–511, 2013.
- [9] Schaid DJ, Chen W, Larson NB. From genome-wide associations to candidate causal variants by statistical fine-mapping. *Nat Rev Genet*, 19(8):491–504, 2018.
- [10] Hormozdiari F, Zhu A, Kichaev G, Ju CJ, *et al.* Widespread Allelic Heterogeneity in Complex Traits. *Am J Hum Genet*, 100(5):789–802, 2017.
- [11] Wheeler HE, Shah KP, Brenner J, Garcia T, *et al.* Survey of the Heritability and Sparse Architecture of Gene Expression Traits across Human Tissues. *PLoS Genet*, 12(11):e1006423, 2016.
- [12] Maller JB, McVean G, Byrnes J, Vukcevic D, *et al.* Bayesian refinement of association signals for 14 loci in 3 common diseases. *Nat Genet*, 44(12):1294–1301, 2012.
- [13] Hormozdiari F, Kostem E, Kang EY, Pasaniuc B, *et al.* Identifying causal variants at loci with multiple signals of association. *Genetics*, 198(2):497–508, 2014.
- [14] Chen W, Larrabee BR, Ovsyannikova IG, Kennedy RB, *et al.* Fine Mapping Causal Variants with an Approximate Bayesian Method Using Marginal Test Statistics. *Genetics*, 200(3):719–736, 2015.

- [15] Benner C, Spencer CC, Havulinna AS, Salomaa V, *et al.* FINEMAP: efficient variable selection using summary data from genome-wide association studies. *Bioinformatics*, 32(10):1493–1501, 2016.
- [16] Kichaev G, Yang WY, Lindstrom S, Hormozdiari F, *et al.* Integrating functional data to prioritize causal variants in statistical fine-mapping studies. *PLoS Genet*, 10(10):e1004722, 2014.
- [17] Gupta RM, Hadaya J, Trehan A, Zekavat SM, *et al.* A Genetic Variant Associated with Five Vascular Diseases Is a Distal Regulator of Endothelin-1 Gene Expression. *Cell*, 170(3):522–533, 2017.
- [18] Knowles DA, Davis JR, Edgington H, Raj A, *et al.* Allele-specific expression reveals interactions between genetic variation and environment. *Nat Methods*, 14(7):699–702, 2017.
- [19] van de Geijn B, McVicker G, Gilad Y, Pritchard JK. WASP: allele-specific software for robust molecular quantitative trait locus discovery. *Nat Methods*, 12(11):1061–1063, 2015.
- [20] Kumasaka N, Knights AJ, Gaffney DJ. Fine-mapping cellular QTLs with RASQUAL and ATAC-seq. *Nat Genet*, 48(2):206–213, 2016.
- [21] Moyerbrailean GA, Richards AL, Kurtz D, Kalita CA, *et al.* High-throughput allele-specific expression across 250 environmental conditions. *Genome Res*, 26(12):1627–1638, 2016.
- [22] Mohammadi P, Castel SE, Brown AA, Lappalainen T. Quantifying the regulatory effect size of cis-acting genetic variation using allelic fold change. *Genome Res*, 27(11):1872–1884, 2017.
- [23] Gusev A, Spisak S, Fay AP, Carol H, *et al.* Allelic imbalance reveals widespread germline-somatic regulatory differences and prioritizes risk loci in Renal Cell Carcinoma. *bioRxiv*, 2019. doi: 10.1101/631150.
- [24] Zou J, Hormozdiari F, Jew B, Ernst J, *et al.* Leveraging allele-specific expression to refine fine-mapping for eQTL studies. *bioRxiv*, 2018. doi: 10.1101/257279.
- [25] Moyerbrailean GA, Kalita CA, Harvey CT, Wen X, *et al.* Which Genetics Variants in DNase-Seq Footprints Are More Likely to Alter Binding? *PLoS Genet*, 12(2):e1005875, 2016.
- [26] Das S, Forer L, Schonherr S, Sidore C, *et al.* Next-generation genotype imputation service and methods. *Nat Genet*, 48(10):1284–1287, 2016.
- [27] Dobin A, Davis CA, Schlesinger F, Drenkow J, *et al.* STAR: ultrafast universal RNA-seq aligner. *Bioinformatics*, 29(1):15–21, 2013.
- [28] Li H, Durbin R. Fast and accurate short read alignment with Burrows-Wheeler transform. *Bioinformatics*, 25(14):1754–1760, 2009.
- [29] Zhang Y, Liu T, Meyer CA, Eeckhoute J, *et al.* Model-based analysis of ChIP-Seq (MACS). *Genome Biol*, 9(9):R137, 2008.
- [30] Wang G, Sarkar A, Carbonetto P, Stephens M. A simple new approach to variable selection in regression, with application to genetic fine-mapping. *bioRxiv*, 2019. doi: 10.1101/501114.

- [31] Benner C, Havulinna AS, Jarvelin MR, Salomaa V, *et al.* Prospects of Fine-Mapping Trait-Associated Genomic Regions by Using Summary Statistics from Genome-wide Association Studies. *Am J Hum Genet*, 101(4):539–551, 2017.
- [32] Creighton CJ, Morgan M, Gunaratne PH, Wheeler DA, *et al.* Comprehensive molecular characterization of clear cell renal cell carcinoma. *Nature*, 499(7456):43–49, 2013.
- [33] Colli LM, Jessop L, Myers TA, Machiela MJ, *et al.* Abstract 401: Functional characterization of the 14q24 renal cancer susceptibility locus implicates SWI/SNF complex member DPF3 via inhibition of apoptosis. *Cancer Research*, 78(13 Supplement):401–401, 2018. ISSN 0008-5472. doi: 10.1158/1538-7445.AM2018-401.
- [34] Kundaje A, Meuleman W, Ernst J, Bilenky M, *et al.* Integrative analysis of 111 reference human epigenomes. *Nature*, 518(7539):317–330, 2015.
- [35] Degner JF, Pai AA, Pique-Regi R, Veyrieras JB, *et al.* DNase I sensitivity QTLs are a major determinant of human expression variation. *Nature*, 482(7385):390–394, 2012.
- [36] Grubert F, Zaugg JB, Kasowski M, Ursu O, *et al.* Genetic Control of Chromatin States in Humans Involves Local and Distal Chromosomal Interactions. *Cell*, 162(5):1051–1065, 2015.
- [37] Mumbach MR, Satpathy AT, Boyle EA, Dai C, *et al.* Enhancer connectome in primary human cells identifies target genes of disease-associated DNA elements. *Nat Genet*, 49(11):1602–1612, 2017.
- [38] Loh PR, Danecek P, Palamara PF, Fuchsberger C, *et al.* Reference-based phasing using the Haplotype Reference Consortium panel. *Nat Genet*, 48(11):1443–1448, 2016.
- [39] Herzog L, Straube K, Neugebauer KM. Long-read sequencing of nascent RNA reveals coupling among RNA processing events. *Genome Res*, 28(7):1008–1019, 2018.
- [40] van der Wijst MGP, Brugge H, de Vries DH, Deelen P, *et al.* Single-cell RNA sequencing identifies cell-type-specific cis-eQTLs and co-expression QTLs. *Nat Genet*, 50(4):493–497, 2018.
- [41] Stegle O, Parts L, Piipari M, Winn J, *et al.* Using probabilistic estimation of expression residuals (PEER) to obtain increased power and interpretability of gene expression analyses. *Nat Protoc*, 7(3):500–507, 2012.
- [42] Li Y, Kellis M. RiVIERA-MT: A Bayesian model to infer risk variants in related traits using summary statistics and functional genomic annotations. *bioRxiv*, 2016. doi: 10.1101/059345.
- [43] Gusev A, Ko A, Shi H, Bhatia G, *et al.* Integrative approaches for large-scale transcriptome-wide association studies. *Nat Genet*, 48(3):245–252, 2016.

List of Figures

- 1 Overview of the PLASMA method. (a) Pre-processing of sequence-based data. First, reads are mapped to the sample's genotype. Reads intersecting markers are colored. Then, the sample's genotype is phased. Reads intersecting heterozygous markers can then be mapped to a particular haplotype. Lastly, reads across the locus are aggregated in an allele-specific manner. To visualize this data, the expression is represented by a ring chart, and the genotypes by pedigree symbols. In the ring chart, the diameter signifies the total read count, and the colors signify the proportion of reads coming from each haplotype. For the pedigree symbols, a white circle signifies a wild-type homozygote, a shaded circle signifies a alternative homozygote, and a half-shaded circle signifies a heterozygote. In heterozygotes, the direction of shading corresponds to the direction of heterozygosity (phasing). (b) Visual representation of QTL and AS statistics under a single causal variant, where the alternative allele increases expression. The total expression (y) is determined by the allelic dosage (x), whereas the allelic imbalance (w) is determined by the phasing (v). These two sets of data are used to calculate QTL and AS association statistics (z_β and z_ϕ). (c) Diagram of PLASMA's fine-mapping process. First, QTL and AS statistics are calculated from read data. Then, these statistics, along with an LD matrix, are used to generate probabilities for causal configurations. By searching through the space of these causal configurations, the model produces credible sets and posterior probabilities for each marker. 8

- 2 Comparison of fine-mapping methods on two sets of simulated loci of 100 markers and 100 samples each. These two sets differ in the variance of the AS phenotype, but with a fixed mean coverage and AS heritability. (a, d) Inclusion curve at low and high AS variance, respectively (0.6, 0.8 standard allelic deviation). The y-axis shows the expected proportion of markers included in the credible set, and the x-axis shows the number of selected markers by posterior probability. (b, e) Distribution of 95% confidence credible set sizes at low and high AS variance, respectively. (c, f) Proportions of loci with 95% credible set sizes within given thresholds at low and high AS variance, respectively. Thresholds used are 1, 5, 20, 40, 70, and 100 markers. (g) A per-snp fine-mapping comparison between PLASMA-J and FINEMAP across 500 simulated loci with one causal variant each. The axes denote the posterior log-odds of causality for FINEMAP and PLASMA-J, respectively. The black hexagons represent the joint distribution of all markers, while the red dots represent specifically the causal markers. Univariate histograms for PLASMA-J and FINEMAP are plotted along the margins. 20

- 3 Comparison of the mean 95% confidence credible set sizes across different AS fine-mapping methods. Each square is the mean set size calculated over 500 simulated loci of 100 markers and 100 samples each, with one causal variant. (a, b, c) Mean credible set sizes as a function of standard allelic deviation and mean read coverage, with AS heritability set to 0.4. (d, e, f) Mean credible set sizes as a function of standard allelic deviation and AS heritability, with mean read coverage set to 100 reads. 22

- 4 Comparison of fine-mapping methods on sets of simulated loci with two causal variants. Each locus is of 100 markers and 100 samples each, with AS heritability set to 0.4, QTL heritability set to 0.05, and mean read coverage set to 100 reads. (a) Distribution of 95% confidence credible set sizes. (b) Inclusion curve for including both causal variants. (c) Inclusion curve for including one or both causal variants. 23
- 5 Comparison of the distribution of 95% confidence credible set sizes across loci in kidney tumor and normal samples, with an allelic imbalance false discovery rate of 0.05. Expression was measured as RNA-Seq read counts mapped to phased genotypes. Fine-mapping was conducted with 500 tumor samples and 70 normal samples. Every SNP within 100 kb of each locus was included in the fine-mapping input. (a, c) Distribution of credible set sizes for tumor and normal samples, respectively, under a 1-causal-variant assumption. (b, d) Proportion of loci fine-mapped to 95% credible sets within given thresholds. Thresholds used are 1, 5, 10, 20, 50, and 100. (e) 95% credible set size distributions for randomly down-sampled kidney tumor datasets with decreasing sample sizes. 26
- 6 Comparison of AS and QTL associations in the experimentally-verified loci, as a function of sample size. Regions of 70 kb are shown around the causal marker. AS and QTL p-Values were calculated using from z_ϕ and z_β , respectively. Fine-mapping was conducted with the PLASMA-AS and QTL-Only models, respectively. 95% credible set sizes for the whole locus are displayed in the top right of each subplot. Markers in the 95% credible set are shown in dark gray, while markers not in the sets are shown in light gray. The experimentally verified causal marker is shown in red. Regions of open chromatin (DNase-Seq peaks) are shaded in gray. (a) Associations in the DPF3 locus (b) Associations in the SCARB1 locus. 27
- 7 Comparison of 95% confidence credible set sizes for peaks in prostate tumor and normal cells, with an allelic imbalance false discovery rate of 0.05. Presence of H3k27ac histone marks was quantified as ChIP-Seq read counts mapped to phased genotypes. (a, c) Distribution of credible set sizes for tumor and normal samples, respectively, under a 1-causal-variant assumption. (b, d) Proportion of peaks fine-mapped to 95% credible sets within given thresholds. Thresholds used were 20, 50, 100, 200, 300, and 400. 29
- 8 Enrichment of PLASMA kidney RNA-Seq credible sets on open chromatin. The x-axis represents a minimum posterior probability threshold on the markers within credible sets, with a higher threshold indicating more confident but fewer markers. The markers in these thresholded credible sets were intersected with the functional feature in question. 95% confidence intervals were calculated with Fisher's exact test, with a control of all markers in a locus. 34
- S1 Comparison of estimators for linkage disequilibrium (LD) at various sample sizes. Estimators used were the dosage estimator, the phasing estimator, and the haplotype-level estimator. At each sample size, 100 markers from 1000-Genomes data were used to calculate LD using all three estimators. Plots of LD correlations are shown for each pair of estimators. (a-c) Dosage LD vs. phasing LD at 50, 200 and 1000 samples, respectively. (d-f) Haplotype-Specific LD vs. dosage LD at 50, 200 and 1000 samples, respectively. (g-i) Haplotype-Specific LD vs. phasing LD at 50, 200 and 1000 samples, respectively. 58

S2	Effect of the jointness hyperparameter in PLASMA on fine-mapping. Fine-mapping of simulated loci was conducted with PLASMA-J under various values of the jointness parameter. Each value of the jointness parameter was evaluated using 500 simulated loci with one causal variant per locus. Plotted are the distributions of the 95% confidence credible sets.	59
S3	Mean top AS association statistics (z_ϕ) under various conditions of mean read coverage, heritability, and standard allelic deviation (a metric of the variance of the AS phenotype). Standard allelic deviation is defined as the allelic fraction f such that the range between $1 - f$ and f captures all loci with an AS phenotype within one standard deviation of the mean (balanced expression). Each square is the mean statistic over 500 simulated loci of 100 markers, with one causal variant. (a) Mean top z_ϕ as a function of standard allelic deviation and mean read coverage, with AS heritability set to 0.5. (b) Mean top z_ϕ as a function of standard allelic deviation and AS heritability, with mean read coverage set to 100 reads.	59
S4	Recall rates for 95% credible sets under various conditions of mean read coverage, heritability, and standard allelic deviation (intensity of AS variance). Recall rate is defined as the expected proportion of causal markers included in the credible set. Each square is the mean statistic over 500 simulated loci of 100 markers, with one causal variant. (a, c, e) Recall rates as a function of standard allelic deviation and mean read coverage, with AS heritability set to 0.4. (d, f, h) Recall rates as a function of standard allelic deviation and AS heritability, with mean read coverage set to 100 reads.	60
S5	Impact of imperfect phasing on AS-based methods. Each dataset consists of 500 loci with one causal variant per locus. (a, b) Recall and 95% confidence credible sets under perfect phasing. (c, d) Recall and 95% confidence credible sets under imperfect phasing, with a phasing error z-score correction. The switch error and blip error were set to 0.152% and 0.165%, respectively.	61
S6	Effect of the AS heritability hyperparameter on 95% credible set sizes with tumor RNA-Seq data. (a) Set size distribution with the heritability hyperparameter at 50% (b) Cumulative set size distribution with the heritability hyperparameter at 50% (c) Set size distribution with the heritability hyperparameter at 5% (d) Cumulative set size distribution with the heritability hyperparameter at 5%	62
S7	95% credible set size distributions for randomly down-sampled kidney normal data with decreasing sample sizes.	62
S8	Inclusion curve for fine-mapped kidney tumor RNA-Seq data for PLASMA and AS-Meta across sample sizes. Inclusion was evaluated against a gold-standard set of markers, defined as all markers with at least 0.1 posterior probability under FINEMAP with all samples.	63
S9	Fine-mapping of kidney RNA-Seq loci under the assumption of up to three causal variants. (Note that FINEMAP credible sets may not be calibrated under multiple causal variants.) (a) 95% confidence credible sets. (b) Proportions of loci with 95% credible set sizes within given thresholds. Thresholds used were 1, 5, 10, 20, 50, and 100.	63

S10	Enrichment odds ratios of PLASMA-J credible sets with functional features. The x-axis represents a minimum posterior probability threshold on the markers within credible sets, with a higher threshold indicating more confident but fewer markers. The markers in these thresholded credible sets were intersected with the functional feature in question. 95 % confidence intervals and p-values were calculated under Fisher's exact test, with a control of all markers in a locus. (a) Enrichment odds ratios of fine-mapped credible sets with open chromatin for kidney normal RNA-Seq data. (b) Enrichment odds ratios of fine-mapped credible sets with chromatin looping for prostate tumor H3K27ac ChIP-Seq data. (c, d) Enrichment p-values of kidney tumor and normal credible sets, respectively, with open chromatin. (e) Enrichment of fine-mapped credible sets with chromatin looping for prostate tumor H3K27ac ChIP-Seq data.	64
-----	---	----

List of Tables

S1	Summary of hyperparameter values used in each set of simulation analyses, referenced by figure legend. Parentheses indicate multiple values tested.	65
----	---	----

5 Appendix

5.1 Modeling total expression at a locus (QTL)

5.1.1 Modeling genetic effects on total expression

We calculate marginal effect sizes for a given locus under the conventional linear model of total gene expression. Let us consider a QTL study of a given locus with n individuals and m markers. Let \mathbf{y} be an $(n \times 1)$ vector of total expression across the individuals, recentered at zero. Given a marker i , let \mathbf{x}_i be an $(n \times 1)$ zero-recentered vector of genotypes. We define β_i , the genetic effect of marker i on total gene expression as follows:

$$\mathbf{y} = \mathbf{x}_i \beta_i + \boldsymbol{\epsilon}_i \quad (21)$$

We model the residuals $\boldsymbol{\epsilon}_i$ as normally distributed with variance $\sigma_{y,i}^2$.

5.1.2 Calculation of QTL summary statistics

We use the maximum likelihood estimator of β_i , equivalent to the ordinary-least-squares linear regression estimator:

$$\hat{\beta}_i = \left(\mathbf{x}_i^\top \mathbf{x}_i \right)^{-1} \mathbf{x}_i^\top \mathbf{y} \quad (22)$$

Under the null model where i is not causal, i does not explain any amount of variation of the phenotype, and the variance of \mathbf{y} is simply $\sigma_{y,i}^2$. Thus, under the null:

$$\begin{aligned} \text{Var} \left(\hat{\beta}_i \right) &= \left(\mathbf{x}_i^\top \mathbf{x}_i \right)^{-2} \text{Var} \left(\mathbf{x}_i^\top \mathbf{y} \right) \\ &= \left(\mathbf{x}_i^\top \mathbf{x}_i \right)^{-2} \left(\mathbf{x}_i^\top \mathbf{x}_i \right) \text{Var} (y) \\ &= \left(\mathbf{x}_i^\top \mathbf{x}_i \right)^{-1} \text{Var} (y) \\ &= \left(\mathbf{x}_i^\top \mathbf{x}_i \right)^{-1} \sigma_{y,i}^2 \end{aligned} \quad (23)$$

We estimate $\sigma_{\beta,i}^2$ from the residuals:

$$\hat{\sigma}_{\beta,i}^2 = \frac{\boldsymbol{\epsilon}_i^\top \boldsymbol{\epsilon}_i}{n-1} \quad (24)$$

We thus define our QTL summary statistic (Wald statistic) for marker i as:

$$\hat{z}_{\beta,i} = \frac{\hat{\beta}_i}{\sqrt{\left(\mathbf{x}_i^\top \mathbf{x}_i \right)^{-1} \hat{\sigma}_{y,i}^2}} \quad (25)$$

We assume that the number of individuals is enough such that the observed statistic is normally distributed with unit variance:

$$\hat{z}_{\beta,i} \sim \mathcal{N}(z_{\beta,i}, 1) \quad (26)$$

In the case where \mathbf{x}_i is of unit variance, the statistic simplifies to:

$$\hat{z}_{\beta,i} = \frac{\hat{\beta}_i \sqrt{n}}{\sqrt{\hat{\sigma}_{y,i}^2}} \quad (27)$$

5.2 Modeling allele-specific expression at a locus (AS)

5.2.1 Modeling haplotype-specific effects on expression

We model allele-specific expression under the observation that a cis-regulatory variant often has a greater influence on the gene allele of the same haplotype. Under this model, an individual who is heterozygous for one or more cis-regulatory markers will show an imbalance in expression between the alleles.

From a quantitative perspective, let us consider a single locus in a single individual who is heterozygous for marker i . Let 0 and 1 represent the wild-type and alternative marker alleles, respectively. We define e_0 as the expression of the gene allele on the same phase as marker allele 0, and e_1 as the expression of the gene allele on the same phase as marker allele 1. Let e'_0 and e'_1 be baseline expressions without the effect of marker i . We define δ_i as the cis-regulatory strength of marker allele 1 over marker allele 0 such that:

$$\frac{e_1}{e_0} = \delta_i \frac{e'_1}{e'_0} \quad (28)$$

If we define i 's phase, v_i , we can arbitrarily assign haplotypes A and B . The above equation then becomes:

$$\frac{e_A}{e_B} = (\delta_i)^{v_i} \frac{e'_A}{e'_B} \quad (29)$$

The marker's phase is 1 if haplotype A contains the alternative marker allele, -1 if haplotype B contains the alternative marker allele, and 0 if the individual is homozygous for the marker.

We now re-write Equation 29 as a linear model. Let w be the log expression ratio between haplotypes A and B:

$$w = \log \left(\frac{e_A}{e_B} \right) \quad (30)$$

Let ϕ_i be the log allelic fold change (logAFC) caused by variant i :

$$\phi_i = \log(\delta_i) \quad (31)$$

Let ζ_i be the log baseline expression ratio between haplotypes A and B:

$$\zeta_i = \log \left(\frac{e'_A}{e'_B} \right) \quad (32)$$

With these parameters we rewrite Equation 29 as:

$$w = v_i \phi_i + \zeta_i \quad (33)$$

Given n individuals, this expression becomes:

$$\mathbf{w} = \mathbf{v}_i \phi_i + \boldsymbol{\zeta}_i \quad (34)$$

We assume that $\boldsymbol{\zeta}_i$ is drawn from a normal distribution with variance $\sigma_{w,i}^2$. Note that under this model, ϕ_i can be interpreted as the effect size of marker i on allelic imbalance, with $\boldsymbol{\zeta}_i$ as the residuals. Furthermore, assuming no haplotype bias, both \mathbf{w} and \mathbf{v}_i are zero-centered in expectation.

Experimentally-derived AS data, such as RNA-Seq data, yield reads that are mapped to a

particular haplotype. Given c_A and c_B , the read counts mapped to haplotypes A and B respectively, we define our estimator of w as:

$$\hat{w} = \log \left(\frac{c_A}{c_B} \right) \quad (35)$$

For a given individual j , we define $c_{A,j}$ as the allele-specific read count from haplotype A . We model the allele-specific read count as drawn a beta-binomial distribution, given the total mapped read count c_j :

$$c_{A,j} \sim \text{BB}(\alpha_j, \beta_j, c_j) \quad (36)$$

We define π_j as the expected proportion of read counts (allelic fraction) from haplotype A :

$$\pi_j = \frac{\mathbb{E}[c_{A,j}]}{c_j} = \frac{\alpha_j}{\alpha_j + \beta_j} \quad (37)$$

α_j and β_j can be re-parameterized in terms of π_j and the sampling overdispersion ρ_e .

$$\rho_e = \frac{1}{\alpha_j + \beta_j + 1} \quad (38)$$

With this re-paramaterization, the mean and variance of $c_{A,i}$ is given as follow:

$$\mathbb{E}[c_{A,j}] = c_j \pi_j \quad (39)$$

$$\text{Var}(c_{A,j}) = c_j \pi_j (1 - \pi_j) (1 + \rho_e (c_j - 1)) \quad (40)$$

We use this beta binomial model to estimate the variance of \hat{w}_i . We scale the distribution by $\frac{1}{c_i}$ to get the mean and variance for the read count proportion:

$$\mathbb{E} \left[\frac{c_{A,j}}{c_j} \right] = \pi_j \quad (41)$$

$$\text{Var} \left(\frac{c_{A,j}}{c_j} \right) = \frac{1}{c_j} \pi_j (1 - \pi_j) (1 + \rho_e (c_j - 1)) \quad (42)$$

We define w^* as the logit-transformed allelic fraction:

$$w_j^* = \text{logit}(\pi_j) = \log \frac{\pi_j}{1 - \pi_j} \quad (43)$$

$$\frac{dw_j^*}{d\pi_j} = \frac{1}{\pi_j(1 - \pi_j)} \quad (44)$$

$$\frac{d^2 w_j^*}{d\pi_j^2} = \frac{2\pi_j - 1}{\pi_j^2(1 - \pi_j)^2} \quad (45)$$

We can thus find the approximate mean and variance of \hat{w}_j given w_j^* using Taylor expansions:

$$\begin{aligned}
\mathbb{E}[\hat{w}_j] &= \mathbb{E}\left[\text{logit}\left(\frac{c_{A,j}}{c_j}\right)\right] \\
&\approx \text{logit}\left(\mathbb{E}\left[\frac{c_{A,j}}{c_j}\right]\right) + \frac{1}{2}\text{Var}\left(\frac{c_{A,j}}{c_j}\right)\frac{d^2}{d\pi_j^2}\text{logit}(\pi_j) \\
&\approx \text{logit}(\pi_j) + \frac{1}{2}\left(\frac{1}{c_j}\pi_j(1-\pi_j)(1+\rho_e(c_j-1))\right)\left(\frac{2\pi_j-1}{\pi_j^2(1-\pi_j)^2}\right) \\
&\approx \text{logit}(\pi_j) + \frac{2\pi_j-1}{2c_j\pi_j(1-\pi_j)}(1+\rho_e(c_j-1)) \\
&\approx w_j^* + \frac{1}{c_j}\sinh(w_j^*)(1+\rho_e(c_j-1))
\end{aligned} \tag{46}$$

$$\begin{aligned}
\text{Var}(\hat{w}_j) &= \text{Var}\left(\text{logit}\left(\frac{c_{A,j}}{c_j}\right)\right) \\
&\approx \text{Var}\left(\frac{c_{A,j}}{c_j}\right)\left(\frac{d}{d\pi_j}\text{logit}(\pi_j)\right)^2 \\
&\approx \left(\frac{1}{c_j}\pi_j(1-\pi_j)(1+\rho_e(c_j-1))\right)\left(\frac{1}{\pi_j(1-\pi_j)}\right)^2 \\
&\approx \frac{1+\rho_e(c_j-1)}{c_j\pi_j(1-\pi_j)} \\
&\approx \frac{2}{c_j}(1+\cosh(w_j^*))(1+\rho_e(c_j-1))
\end{aligned} \tag{47}$$

Note that w and w^* are not equivalent because $\mathbb{E}[\text{logit}(c_A/c)] \neq \text{logit}(\mathbb{E}[c_A/c])$. Equation 46 implies that \hat{w} is a biased estimator of w^* , especially at low read counts and/or high overdispersion. To get an estimator of w^* with reduced bias, we take the approximation that $\sinh(w^*) \approx w^*$ around zero:

$$\hat{w}_j^* = \frac{\hat{w}_j}{1 + \frac{1}{c_j}(1 + \rho_e(c_j - 1))} \tag{48}$$

We use \hat{w}^* to find an estimator of $\sigma_{c,j}^2$, the variance of \hat{w} :

$$\hat{\sigma}_{c,j}^2 = \frac{2}{c_j}(1 + \cosh(\hat{w}_j^*))(1 + \rho_e(c_j - 1)) \tag{49}$$

Given our estimator \hat{w}_j , we quantify the sampling error $\tau_j = \hat{w}_j - w_j$, with $\mathbb{E}[\tau_j] = 0$ and $\text{Var}(\tau_j) = \sigma_{c,j}^2$. Thus, across individuals:

$$\hat{\mathbf{w}} = \mathbf{v}_i\phi_i + \boldsymbol{\zeta}_i + \boldsymbol{\tau} \tag{50}$$

5.2.2 Calculation of AS summary statistics

Due to heteroscedasticity among individuals, we estimate the AS effect size ϕ_i in a weighted manner, giving larger weights to individuals with lower expected sampling error. Given individual j , we

define the weight for j as the inverse of the estimated read count variance:

$$\omega_j = \frac{1}{\hat{\sigma}_{c,j}^2} \quad (51)$$

We define our weight matrix $\mathbf{\Omega}$ as a diagonal matrix with $\mathbf{\Omega}_{j,j} = \omega_j$.

We use the weighted-least-squares estimator for ϕ_i :

$$\hat{\phi}_i = \left(\mathbf{v}_i^\top \mathbf{\Omega} \mathbf{v}_i \right)^{-1} \mathbf{v}_i^\top \mathbf{\Omega} \hat{\mathbf{w}} \quad (52)$$

Under the null model where i is not causal, the variance of w_j is $\sigma_{w,i}^2$, and the variance of \hat{w}_j is $\sigma_{w,i}^2 + \sigma_{c,j}^2$. We assume that the experimental errors τ and biological residuals ζ_i are uncorrelated. Thus, under the null:

$$\begin{aligned} \text{Var}(\hat{\phi}_i) &= \mathbb{E} \left[(\hat{\phi}_i - \phi_i)^2 \right] \\ &= \left(\mathbf{v}_i^\top \mathbf{\Omega} \mathbf{v}_i \right)^{-1} \mathbf{v}_i^\top \mathbf{\Omega} (\zeta_i + \boldsymbol{\tau}) (\zeta_i + \boldsymbol{\tau})^\top \mathbf{\Omega}^\top \mathbf{v}_i \left(\mathbf{v}_i^\top \mathbf{\Omega} \mathbf{v}_i \right)^{-1} \\ &= \left(\mathbf{v}_i^\top \mathbf{\Omega} \mathbf{v}_i \right)^{-2} \left(\mathbf{v}_i^\top \mathbf{\Omega} \zeta_i^\top \zeta_i \mathbf{\Omega}^\top \mathbf{v}_i + \mathbf{v}_i^\top \mathbf{\Omega} \boldsymbol{\tau} \boldsymbol{\tau}^\top \mathbf{\Omega}^\top \mathbf{v}_i \right) \\ &= \left(\mathbf{v}_i^\top \mathbf{\Omega} \mathbf{v}_i \right)^{-2} \left(\mathbf{v}_i^\top \mathbf{\Omega} \mathbf{I} \sigma_{w,i}^2 \mathbf{\Omega}^\top \mathbf{v}_i + \mathbf{v}_i^\top \mathbf{\Omega} \mathbf{I} \mathbf{v}_i \right) \\ &= \left(\mathbf{v}_i^\top \mathbf{\Omega} \mathbf{v}_i \right)^{-2} \left(\left(\mathbf{v}_i^\top \mathbf{\Omega}^2 \mathbf{v}_i \right) \sigma_{w,i}^2 + \mathbf{v}_i^\top \mathbf{\Omega} \mathbf{v}_i \right) \end{aligned} \quad (53)$$

We now estimate $\sigma_{w,i}^2$ from the residuals. Note that we are estimating the variance of the biological residuals $\text{Var}(\zeta_i)$, which is distinct from the total residuals are $\zeta_i + \boldsymbol{\tau}$, so we cannot directly use the variance of the total residuals. We instead use the following estimator for $\sigma_{w,i}^2$:

$$\hat{\sigma}_{w,i}^2 = \frac{\sum_{j=1}^n \left(\omega_j (\zeta_{i,j} + \tau_j)^2 - 1 \right)}{\sum_{j=1}^n \omega_j} \quad (54)$$

We show that this estimator is equal to $\sigma_{\phi,i}^2$ in expectation:

$$\begin{aligned} \mathbb{E}[\hat{\sigma}_{w,i}^2] &= \frac{\sum_{j=1}^n \left(\omega_j \mathbb{E} \left[(\zeta_{i,j} + \tau_j)^2 \right] - 1 \right)}{\sum_{j=1}^n \omega_j} \\ &= \frac{\sum_{j=1}^n \left(\omega_j \text{Var}(\zeta_{i,j} + \tau_j) - 1 \right)}{\sum_{j=1}^n \omega_j} \\ &= \frac{\sum_{j=1}^n \left(\omega_j \text{Var}(\zeta_{i,j}) + \omega_j \text{Var}(\tau_j) - 1 \right)}{\sum_{j=1}^n \omega_j} \\ &= \frac{\sum_{j=1}^n \omega_j \text{Var}(\zeta_{i,j})}{\sum_{j=1}^n \omega_j} \\ &= \text{Var}(\zeta_i) \\ &= \sigma_{w,i}^2 \end{aligned} \quad (55)$$

With this estimator, we define the AS association statistic for marker i as follows:

$$\hat{z}_{\phi,i} = \frac{\hat{\phi}_i}{\sqrt{(\mathbf{v}_i^\top \boldsymbol{\Omega} \mathbf{v}_i)^{-2} \left((\mathbf{v}_i^\top \boldsymbol{\Omega}^2 \mathbf{v}_i) \hat{\sigma}_{w,i}^2 + \mathbf{v}_i^\top \boldsymbol{\Omega} \mathbf{v}_i \right)}} \quad (56)$$

We assume that the observed statistic is normally distributed with unit variance:

$$\hat{z}_{\phi,i} \sim \mathcal{N}(z_{\phi,i}, 1) \quad (57)$$

To gain an intuitive understanding of the association statistic, let us examine it under simplifying conditions. We assume that \mathbf{v}_i is of unit variance, that read count overdispersion is negligible, and that allelic imbalance and read coverage are fixed across individuals. Under these conditions, let $\boldsymbol{\Omega} = \frac{c}{k} \mathbf{I}$ for coverage c and some constant k . Equation 56 simplifies to:

$$\hat{z}_{\phi,i} = \frac{\hat{\phi}_i \sqrt{n}}{\sqrt{\hat{\sigma}_{w,i}^2 + \frac{k}{c}}} \quad (58)$$

We can see that under high experimental noise (k/c), the denominator is dominated by the quality of data (read coverage). In contrast, when experimental noise is low, the denominator is dominated by $\hat{\sigma}_{w,i}^2$, determined by the inherent heritability of the locus's AS phenotype.

In the case where phasing error is significant, we would expect the estimated AS effects ($\hat{\phi}$) to have more deviation from the true effects. We derive a correction for the AS z-score, given a per-marker probability of mis-phasing ψ_i . We define $\hat{\mathbf{v}}_i$ as the imperfect observed phasing for marker i , and we define the phasing error vector $\boldsymbol{\delta}_i$ such that $\boldsymbol{\delta}_i = \mathbf{v}_i - \hat{\mathbf{v}}_i$. Note that each δ is a ternary -2/0/2 indicator, with each δ^2 being a binary 0/4 indicator of a phasing error. We assume that the occurrence of a phasing error is independent of which haplotype the alternative allele is one, so that $\mathbb{E}[\boldsymbol{\delta}_i] = 0$. We now derive the variance of $\hat{\phi}_i$ under imperfect phasing:

$$\begin{aligned} \hat{\phi}_i - \phi_i &= \left(\hat{\mathbf{v}}_i^\top \boldsymbol{\Omega} \hat{\mathbf{v}}_i \right)^{-1} \hat{\mathbf{v}}_i^\top \boldsymbol{\Omega} \hat{\mathbf{w}} - \phi_i \\ &= \left(\hat{\mathbf{v}}_i^\top \boldsymbol{\Omega} \hat{\mathbf{v}}_i \right)^{-1} \hat{\mathbf{v}}_i^\top \boldsymbol{\Omega} (\mathbf{v}_i \phi_i + \boldsymbol{\zeta}_i + \boldsymbol{\tau}) - \phi_i \\ &= \left(\hat{\mathbf{v}}_i^\top \boldsymbol{\Omega} \hat{\mathbf{v}}_i \right)^{-1} \hat{\mathbf{v}}_i^\top \boldsymbol{\Omega} (\hat{\mathbf{v}}_i \phi_i + \boldsymbol{\delta}_i \phi_i + \boldsymbol{\zeta}_i + \boldsymbol{\tau}) - \phi_i \\ &= \left(\hat{\mathbf{v}}_i^\top \boldsymbol{\Omega} \hat{\mathbf{v}}_i \right)^{-1} \hat{\mathbf{v}}_i^\top \boldsymbol{\Omega} (\boldsymbol{\delta}_i \phi_i + \boldsymbol{\zeta}_i + \boldsymbol{\tau}) \end{aligned} \quad (59)$$

$$\begin{aligned} \text{Var}(\hat{\phi}_i) &= \mathbb{E} \left[(\hat{\phi}_i - \phi_i)^2 \right] \\ &= \left(\hat{\mathbf{v}}_i^\top \boldsymbol{\Omega} \hat{\mathbf{v}}_i \right)^{-1} \hat{\mathbf{v}}_i^\top \boldsymbol{\Omega} (\boldsymbol{\delta}_i \phi_i + \boldsymbol{\zeta}_i + \boldsymbol{\tau}) (\boldsymbol{\delta}_i \phi_i + \boldsymbol{\zeta}_i + \boldsymbol{\tau})^\top \boldsymbol{\Omega}^\top \hat{\mathbf{v}}_i \left(\hat{\mathbf{v}}_i^\top \boldsymbol{\Omega} \hat{\mathbf{v}}_i \right)^{-1} \\ &= \left(\hat{\mathbf{v}}_i^\top \boldsymbol{\Omega} \hat{\mathbf{v}}_i \right)^{-2} \left(\hat{\mathbf{v}}_i^\top \boldsymbol{\Omega} \boldsymbol{\delta}_i^\top \boldsymbol{\delta}_i \phi_i^2 \boldsymbol{\Omega}^\top \hat{\mathbf{v}}_i + \hat{\mathbf{v}}_i^\top \boldsymbol{\Omega} \boldsymbol{\zeta}_i^\top \boldsymbol{\zeta}_i \boldsymbol{\Omega}^\top \hat{\mathbf{v}}_i + \hat{\mathbf{v}}_i^\top \boldsymbol{\Omega} \boldsymbol{\tau} \boldsymbol{\tau}^\top \boldsymbol{\Omega}^\top \hat{\mathbf{v}}_i \right) \\ &= \left(\hat{\mathbf{v}}_i^\top \boldsymbol{\Omega} \hat{\mathbf{v}}_i \right)^{-2} \left(\hat{\mathbf{v}}_i^\top \boldsymbol{\Omega} (\mathbf{I} 4\psi_i \phi_i^2 + \mathbf{I} \sigma_{w,i}^2) \boldsymbol{\Omega}^\top \hat{\mathbf{v}}_i + \hat{\mathbf{v}}_i^\top \boldsymbol{\Omega} \mathbf{I} \hat{\mathbf{v}}_i \right) \\ &= \left(\hat{\mathbf{v}}_i^\top \boldsymbol{\Omega} \hat{\mathbf{v}}_i \right)^{-2} \left(\left(\hat{\mathbf{v}}_i^\top \boldsymbol{\Omega}^2 \hat{\mathbf{v}}_i \right) (\sigma_{w,i}^2 + 4\psi_i \phi_i^2) + \hat{\mathbf{v}}_i^\top \boldsymbol{\Omega} \hat{\mathbf{v}}_i \right) \end{aligned} \quad (60)$$

When calculating $\text{Var}(\hat{\phi}_i)$, we approximate the ϕ_i^2 term with the observed ϕ_i^2 . We thus define

the corrected z-score:

$$\hat{z}_{\phi,i} = \frac{\hat{\phi}_i}{\sqrt{(\hat{\mathbf{v}}_i^\top \boldsymbol{\Omega} \hat{\mathbf{v}}_i)^{-2} \left((\hat{\mathbf{v}}_i^\top \boldsymbol{\Omega}^2 \hat{\mathbf{v}}_i) \left(\hat{\sigma}_{w,i}^2 + 4\psi_i \hat{\phi}_i^2 \right) + \hat{\mathbf{v}}_i^\top \boldsymbol{\Omega} \hat{\mathbf{v}}_i \right)}} \quad (61)$$

5.2.3 Empirical Analysis of Imperfect Phasing

We evaluate the AS-based methods, PLASMA and RASQUAL+, on simulated loci with imperfect phasing. (We do not include AS-Meta, since the method does not depend on phased genotypes.) We modelled our imperfect phasing simulation on the EAGLE phasing paper [cite], which quantified two sources of phasing error: switch error and blip error. Switch error is the probability that the phasing swaps at a given marker (affecting all downstream markers) and was set to 0.152%. Blip error is the probability that a given marker is misphased, with downstream markers unaffected, and was set to 0.165%.

Results for these imperfectly-phased loci are shown in Figure S5, with comparisons to perfectly phased loci. Comparing the inclusion curves for each method, we see only a very small difference in power with perfect and imperfect phasing, which is expected given the low error rates reported by EAGLE. Similarly, we see a small increase in credible set sizes for each model (RASQUAL+ not included due to non-meaningful sets). However, we do see a modest decrease in recall (0.98 to 0.89 for PLASMA-J, and 0.98 to 0.92 for PLASMA-AS), indicating that imperfect phasing does have some impact on accuracy. Nevertheless, it appears that the overall effect of imperfect phasing is small at real-world error rates.

5.3 Inference of causal variants with QTL and AS statistics

5.3.1 Modeling the effects of linkage disequilibrium

PLASMA assumes that the QTL and AS statistics are drawn from the same marker linkage disequilibrium (LD), a property that we verify under the assumption that the two alleles in an individual are distributed independently (See following section for proof). This equivalence yields three LD estimators: using dosage (0/1/2) genotypes x , phasing (-1/0/1 genotypes v), or haplotype-specific (0/1) genotypes. To evaluate each of these three estimators and to empirically demonstrate their equivalence in expectation, we sample a contiguous block of 100 SNPs from phased 1000-Genomes genotype data at sample sizes ranging from 50 to 1000. We then calculate LD matrices using each estimator and compare the estimated correlations for each pair of markers. In Figure S1, we see that correlation among the estimators increases as a function of sample size, and also that a sample size of 50 is more than sufficient for high concordance among the three estimators. Since the haplotype-specific estimator appears to be more correlated to both the dosage estimator and the phasing estimator than the latter two are to each other, we use the haplotype-specific estimator in our analyses. We believe that the haplotype-specific estimator is the most accurate of the three since it as effectively double the sample size as the dosage and phasing estimators.

5.3.2 Equivalence in expectation of QTL and AS estimators of LD

Due to linkage disequilibrium, there exist significant correlations of genotypes among markers. This correlation is reflected in the correlations in the association statistics. PLASMA models the QTL (dosage) and AS (phasing) genotypes having an identical LD structure. Here, we show that the LD calculated from the dosage, phasing, and haplotype-specific genotypes are equivalent in expectation, under the assumption that the two alleles in an individual are independent.

Let a be a binary genetic marker, and let p_a be a random variable for the (0/1) genotype of a on haplotype I (arbitrarily chosen), and let q_a be a random variable for the (0/1) genotype of a on haplotype II.

We define the dosage genotype of marker a as the sum of the centered haplotypes:

$$x_a = (p_a - \bar{p}_a) + (q_a - \bar{q}_a) \quad (62)$$

Likewise, we define the phasing genotype of marker a as the difference of the centered haplotypes:

$$v_a = (p_a - \bar{p}_a) - (q_a - \bar{q}_a) \quad (63)$$

(Note that since $\bar{p}_a = \bar{q}_a$, the phasing is also equivalent to the difference of the uncentered haplotypes.)

We now calculate the covariance of x_a and x_b for two markers a and b . Since these variables are zero-centered, we can express the covariance as $\mathbb{E}[x_a x_b]$:

$$\begin{aligned} \mathbb{E}[x_a x_b] &= \mathbb{E}[(p_a - \bar{p}_a + q_a - \bar{q}_a)(p_b - \bar{p}_b + q_b - \bar{q}_b)] \\ &= \mathbb{E}[(p_a - \bar{p}_a)(p_b - \bar{p}_b)] + \mathbb{E}[(q_a - \bar{q}_a)(p_b - \bar{p}_b)] \\ &\quad + \mathbb{E}[(p_a - \bar{p}_a)(q_b - \bar{q}_b)] + \mathbb{E}[(q_a - \bar{q}_a)(q_b - \bar{q}_b)] \end{aligned} \quad (64)$$

We make the assumption that the two alleles in an individual are independent, so that the covariance of any p and q is zero. The covariance simplifies to:

$$\begin{aligned} \mathbb{E}[x_a x_b] &= \mathbb{E}[(p_a - \bar{p}_a)(p_b - \bar{p}_b)] + \mathbb{E}[(q_a - \bar{q}_a)(q_b - \bar{q}_b)] \\ &= 2\mathbb{E}[(p_a - \bar{p}_a)(p_b - \bar{p}_b)] \end{aligned} \quad (65)$$

We calculate the correlation of x_a and x_b from this correlation, and we see that it is equivalent to the correlation of p_a and p_b :

$$\begin{aligned} \frac{\mathbb{E}[x_a x_b]}{\sqrt{\mathbb{E}[x_a x_a] \mathbb{E}[x_b x_b]}} &= \frac{2\mathbb{E}[(p_a - \bar{p}_a)(p_b - \bar{p}_b)]}{\sqrt{2\mathbb{E}[(p_a - \bar{p}_a)(p_a - \bar{p}_a)] 2\mathbb{E}[(p_b - \bar{p}_b)(p_b - \bar{p}_b)]}} \\ &= \frac{\mathbb{E}[(p_a - \bar{p}_a)(p_b - \bar{p}_b)]}{\sqrt{\mathbb{E}[(p_a - \bar{p}_a)(p_a - \bar{p}_a)] \mathbb{E}[(p_b - \bar{p}_b)(p_b - \bar{p}_b)]}} \end{aligned} \quad (66)$$

Under the same assumptions, we calculate the covariance of v_a and v_b .

$$\begin{aligned} \mathbb{E}[v_a v_b] &= \mathbb{E}[(p_a - \bar{p}_a - (q_a - \bar{q}_a))(p_b - \bar{p}_b - (q_b - \bar{q}_b))] \\ &= \mathbb{E}[(p_a - \bar{p}_a)(p_b - \bar{p}_b)] - \mathbb{E}[(q_a - \bar{q}_a)(p_b - \bar{p}_b)] \\ &\quad - \mathbb{E}[(p_a - \bar{p}_a)(q_b - \bar{q}_b)] + \mathbb{E}[(q_a - \bar{q}_a)(q_b - \bar{q}_b)] \\ &= \mathbb{E}[(p_a - \bar{p}_a)(p_b - \bar{p}_b)] + \mathbb{E}[(q_a - \bar{q}_a)(q_b - \bar{q}_b)] \\ &= 2\mathbb{E}[(p_a - \bar{p}_a)(p_b - \bar{p}_b)] \end{aligned} \quad (67)$$

We see that the covariance of v_a and v_b is equal to the covariance of x_a and x_b , implying that the correlation of v_a and v_b is equal to the correlation of x_a and x_b , and thus p_a and p_b . Therefore, under the assumption that the two alleles in an individual are independent, the LD calculated from dosage genotypes x , phasing genotypes v , and haplotype-specific genotypes p are equivalent in expectation. (Note that this assumption can theoretically be violated if a marker is under selection, but we assume that selection is insignificant due to the weak effects of common markers.)

5.3.3 Jointly modeling total and haplotype-specific effects on expression

We define $\hat{\mathbf{z}}$ as the combined vector of AS association statistics and QTL association statistics:

$$\hat{\mathbf{z}} = \begin{bmatrix} \hat{\mathbf{z}}_\phi \\ \hat{\mathbf{z}}_\beta \end{bmatrix} \quad (68)$$

Let $r_{\beta\phi}$ be the overall correlation between the QTL and AS summary statistics calculated across all loci. We define the combined correlation matrix \mathbf{R} as:

$$\mathbf{R} = \begin{bmatrix} \mathbf{R}_z & r_{\beta\phi}\mathbf{R}_z \\ r_{\beta\phi}\mathbf{R}_z & \mathbf{R}_z \end{bmatrix} \quad (69)$$

We model the joint distribution as multivariate normal, with covariance \mathbf{R} :

$$\hat{\mathbf{z}} \sim \mathcal{N}_{2m}(\mathbf{z}, \mathbf{R}) \quad (70)$$

5.3.4 Modeling summary statistics given a causal configuration

The goal of this method is to infer the causal markers, given QTL and AS association statistics. To this end, we introduce a likelihood function that gives the probability of statistics $\hat{\mathbf{z}}$, given a causal configuration. We define a causal configuration \mathbf{c} as a vector of causal statuses corresponding to each marker, with 1 being causal and 0 being non-causal.

Let $\mathbf{z}_{\mathbf{c},\phi}$ and $\mathbf{z}_{\mathbf{c},\beta}$ be the underlying causal AS and QTL effects, respectively, across markers such that:

$$\mathbf{z}_{\mathbf{c}} = \begin{bmatrix} \mathbf{z}_{\mathbf{c},\phi} \\ \mathbf{z}_{\mathbf{c},\beta} \end{bmatrix} \quad (71)$$

We define hyperparameters $\sigma_{c,\phi}^2$ and $\sigma_{c,\beta}^2$ as the variance of AS and QTL causal effect sizes, respectively $r_{c,\beta\phi}$ as the underlying correlation of the causal QTL and AS effect sizes. (This is not to be confused with $r_{\beta\phi}$, which concerns the correlation between the association statistics.) We define $\Sigma_{\mathbf{c}}$, the covariance matrix of causal effect sizes, given a causal configuration:

$$\Sigma_{\mathbf{c}} = \begin{bmatrix} \text{diag}(\mathbf{c})\sigma_{c,\phi}^2 & \text{diag}(\mathbf{c})r_{c,\beta\phi}\sigma_{c,\phi}\sigma_{c,\beta} \\ \text{diag}(\mathbf{c})r_{c,\beta\phi}\sigma_{c,\phi}\sigma_{c,\beta} & \text{diag}(\mathbf{c})\sigma_{c,\beta}^2 \end{bmatrix} \quad (72)$$

We model the causal effect sizes, given a causal configuration, as drawn from a multivariate normal distribution:

$$\mathbf{z}_{\mathbf{c}}|\mathbf{c} \sim \mathcal{N}_{2m}(\mathbf{0}, \Sigma_{\mathbf{c}}) \quad (73)$$

Furthermore, we model the expected association statistic for a given marker as a linear combination of all effects correlated to the marker.

$$\mathbf{z} = \mathbf{R}\mathbf{z}_{\mathbf{c}} \quad (74)$$

$$\mathbf{z}|\mathbf{c} \sim \mathcal{N}_{2m}(\mathbf{0}, \mathbf{R}\Sigma_{\mathbf{c}}\mathbf{R}) \quad (75)$$

Combining Equations 70 and 75, we get a probability distribution for the observed association statistics given a causal configuration. This is our likelihood for a causal configuration.

$$\begin{aligned} \hat{\mathbf{z}}|\mathbf{c} &\sim \mathcal{N}_{2m}(\mathbf{0}, \mathbf{R} + \mathbf{R}\Sigma_{\mathbf{c}}\mathbf{R}) \\ P(\hat{\mathbf{z}}|\mathbf{c}) &= \mathcal{N}_{2m}(\mathbf{0}, \mathbf{R} + \mathbf{R}\Sigma_{\mathbf{c}}\mathbf{R}) = \mathcal{L}(\mathbf{c}; \hat{\mathbf{z}}) \end{aligned} \quad (76)$$

To get a prior distribution for the causal configuration \mathbf{c} , we define the hyperparameter γ as

the prior probability that a single variant is causal and $1 - \gamma$ as the probability that a variant is not causal. The probability of a configuration consisting of m variants thus becomes:

$$P(\mathbf{c}) = \prod_{i=1}^m \gamma^{c_i} (1 - \gamma)^{1 - c_i} \quad (77)$$

We can view the prior as a regularization term by taking the negative log:

$$\begin{aligned} -\log P(\mathbf{c}) &= \sum_{i=1}^m (-c_i \log(\gamma) - \log(1 - \gamma)) \\ &= -\text{logit}(\gamma) \|\mathbf{c}\|_k - m \log(1 - \gamma) \end{aligned} \quad (78)$$

Since \mathbf{c} is a binary vector, $\|\mathbf{c}\|_k$ is the same for all positive k . Thus, the prior imposes L_k regularization with $\lambda = -\text{logit}(\gamma)$. In practice, this regularization favors causal configurations with fewer causal variants.

With the prior and likelihood, we define the posterior probability of a causal configuration, normalized across the set of all possible configurations \mathbb{C} :

$$P(\mathbf{c}|\hat{\mathbf{z}}) = \frac{P(\hat{\mathbf{z}}|\mathbf{c})P(\mathbf{c})}{\sum_{\mathbf{c}^* \in \mathbb{C}} P(\hat{\mathbf{z}}|\mathbf{c}^*)P(\mathbf{c}^*)} \quad (79)$$

This posterior probability can be alternatively expressed with Bayes Factors. We define the null model as the scenario where all markers are non-causal, so that $\mathbf{c} = \mathbf{0}$. The Bayes Factor for a particular \mathbf{c} would thus be:

$$\text{BF}(\mathbf{c} : \mathbf{0}) = \frac{P(\hat{\mathbf{z}}|\mathbf{c})}{P(\hat{\mathbf{z}}|\mathbf{0})} \quad (80)$$

We rewrite Equation 79 with Bayes Factors:

$$P(\mathbf{c}|\hat{\mathbf{z}}) = \frac{\text{BF}(\mathbf{c} : \mathbf{0})P(\mathbf{c})}{\sum_{\mathbf{c}^* \in \mathbb{C}} \text{BF}(\mathbf{c}^* : \mathbf{0})P(\mathbf{c}^*)} \quad (81)$$

5.3.5 The ρ -level credible set

In practice, due to the large number of possible configurations, the probability of any given configuration will likely be small. For more meaningful probabilities, we calculate the total probability of the possible non-null configurations from a set of markers.

We define \mathbb{K} as a set of markers that putatively includes all causal markers. We define $\mathbb{C}_{\mathbb{K}}$ as the set of all causal configurations whose causal markers is a subset of \mathbb{K} , excluding the null set. Thus, the probability that \mathbb{K} includes all causal markers is the sum of the probabilities of the configurations in $\mathbb{C}_{\mathbb{K}}$.

$$P(\mathbb{K}|\hat{\mathbf{z}}) = \sum_{\mathbf{c} \in \mathbb{C}_{\mathbb{K}}} \text{Pr}(\mathbf{c}|\hat{\mathbf{z}}) \quad (82)$$

We set this probability as ρ_c , the confidence level of \mathbb{K} . Given a value for ρ_c , commonly 0.95, we seek to find \mathbb{K} that minimizes the number of causal variants.

5.3.6 The posterior inclusion probability

An alternative way of summarizing the configurations is to calculate a marker's posterior inclusion probability (PIP), also known as the posterior probability of association. We define the PIP as the

probability that a single given marker is causal, marginalized over all other markers. We calculate this probability by summing over all configurations containing the marker.

5.4 Computational optimization and implementation

5.4.1 Shotgun stochastic search across configurations

The computation of the probability of a given configuration requires knowledge of the Bayes Factor for every possible configuration. As there are 2^m possible configurations, traversing this whole space is intractable. To reduce the number of configurations to evaluate, we use the heuristic that configurations with significant probabilities tend to be similar to each other.

We use a shotgun stochastic search procedure based on that of FINEMAP to find all configurations with a significant probability [15]. Given a selected configuration \mathbf{c} , we define $\mathbb{G}_{\mathbf{c}}$, the neighborhood of \mathbf{c} , as follows:

- All configurations resulting from setting a causal marker in \mathbf{c} to non-causal
- All configurations resulting from setting a non-causal marker in \mathbf{c} to causal
- All configurations resulting from swapping the casual statuses of two markers in \mathbf{c}

For each iteration of the algorithm, the next configuration is drawn randomly from $\mathbb{G}_{\mathbf{c}}$, weighted by the posterior probability of each candidate. Upon termination, we assume that all configurations with nonzero probability have been uncovered.

5.4.2 Calculation of the ρ -level credible set

Given the large number of configurations evaluated, it is impractical to calculate the best possible credible set satisfying ρ_c . Instead, we use a greedy approximation algorithm utilized in CAVIAR and CAVIARBF [13, 14]. At each step, before ρ_c is reached, the algorithm adds the marker that increases the confidence the most.

5.4.3 Bayes factor evaluation with matrix reduction

Direct calculation of the Bayes Factor for a configuration requires the manipulation of $m \times m$ matrices, resulting in an $O(m^3)$ runtime per configuration. We now show that it is sufficient to evaluate the Bayes Factor using only the elements corresponding to causal SNPs, an optimization introduced by FINEMAP [15]. This reduces complexity from $O(m^3)$ to $O(k^3)$, where k is the number of causal variants.

We expand the MVN probability density functions in equation 80 and use the binomial inverse

theorem:

$$\begin{aligned}
\text{BF}(\mathbf{c} : \mathbf{0}) &= \frac{|\mathbf{R} + \mathbf{R}\boldsymbol{\Sigma}_c\mathbf{R}|^{-(1/2)} \exp\left(-\frac{1}{2}\hat{\mathbf{z}}^\top (\mathbf{R} + \mathbf{R}\boldsymbol{\Sigma}_c\mathbf{R})^{-1} \hat{\mathbf{z}}\right)}{|\mathbf{R}|^{-(1/2)} \exp\left(-\frac{1}{2}\hat{\mathbf{z}}^\top (\mathbf{R})^{-1} \hat{\mathbf{z}}\right)} \\
&= \frac{|\mathbf{I} + \boldsymbol{\Sigma}_c\mathbf{R}|^{-(1/2)} \exp\left(-\frac{1}{2}\hat{\mathbf{z}}^\top (\mathbf{R} + \mathbf{R}\boldsymbol{\Sigma}_c\mathbf{R})^{-1} \hat{\mathbf{z}}\right)}{\exp\left(-\frac{1}{2}\hat{\mathbf{z}}^\top (\mathbf{R})^{-1} \hat{\mathbf{z}}\right)} \\
&= \frac{|\mathbf{I} + \boldsymbol{\Sigma}_c\mathbf{R}|^{-(1/2)} \exp\left(-\frac{1}{2}\hat{\mathbf{z}}^\top \left(\mathbf{R}^{-1} - (\mathbf{I} + \boldsymbol{\Sigma}_c\mathbf{R})^{-1} \boldsymbol{\Sigma}_c\right) \hat{\mathbf{z}}\right)}{\exp\left(-\frac{1}{2}\hat{\mathbf{z}}^\top (\mathbf{R})^{-1} \hat{\mathbf{z}}\right)} \\
&= |\mathbf{I} + \boldsymbol{\Sigma}_c\mathbf{R}|^{-(1/2)} \exp\left(\frac{1}{2}\hat{\mathbf{z}}^\top (\mathbf{I} + \boldsymbol{\Sigma}_c\mathbf{R})^{-1} \boldsymbol{\Sigma}_c \hat{\mathbf{z}}\right)
\end{aligned} \tag{83}$$

We permute s to separate causal and non-causal SNPs:

$$\hat{\mathbf{z}} = \begin{bmatrix} \hat{\mathbf{z}}_C \\ \hat{\mathbf{z}}_N \end{bmatrix} \tag{84}$$

We likewise permute the rows and columns of \mathbf{R} and $\boldsymbol{\Sigma}_c$ such that:

$$\mathbf{R} = \begin{bmatrix} \mathbf{R}_{CC} & \mathbf{R}_{CN} \\ \mathbf{R}_{NC} & \mathbf{R}_{NN} \end{bmatrix} \tag{85}$$

$$\boldsymbol{\Sigma}_c = \begin{bmatrix} \boldsymbol{\Sigma}_{c,CC} & \boldsymbol{\Sigma}_{c,CN} \\ \boldsymbol{\Sigma}_{c,NC} & \boldsymbol{\Sigma}_{c,NN} \end{bmatrix} = \begin{bmatrix} \boldsymbol{\Sigma}_{c,CC} & \mathbf{0}_{k \times (m-k)} \\ \mathbf{0}_{(m-k) \times k} & \mathbf{0}_{(m-k) \times (m-k)} \end{bmatrix} \tag{86}$$

Note that $\boldsymbol{\Sigma}_c$ can be nonzero only among causal markers since \mathbf{c} is 0 for non-causal markers. Furthermore:

$$\begin{aligned}
\boldsymbol{\Sigma}_c\mathbf{R} &= \begin{bmatrix} \boldsymbol{\Sigma}_{c,CC} & \mathbf{0}_{k \times (m-k)} \\ \mathbf{0}_{(m-k) \times k} & \mathbf{0}_{(m-k) \times (m-k)} \end{bmatrix} \begin{bmatrix} \mathbf{R}_{CC} & \mathbf{R}_{CN} \\ \mathbf{R}_{NC} & \mathbf{R}_{NN} \end{bmatrix} \\
&= \begin{bmatrix} \boldsymbol{\Sigma}_{c,CC}\mathbf{R}_{CC} & \boldsymbol{\Sigma}_{c,CC}\mathbf{R}_{CN} \\ \mathbf{0}_{(m-k) \times k} & \mathbf{0}_{(m-k) \times (m-k)} \end{bmatrix} \\
\mathbf{I} + \boldsymbol{\Sigma}_c\mathbf{R} &= \begin{bmatrix} \mathbf{I}_{k \times k} + \boldsymbol{\Sigma}_{c,CC}\mathbf{R}_{CC} & \boldsymbol{\Sigma}_{c,CC}\mathbf{R}_{CN} \\ \mathbf{0}_{(m-k) \times k} & \mathbf{I}_{(m-k) \times (m-k)} \end{bmatrix}
\end{aligned} \tag{87}$$

Blockwise inversion yields:

$$\begin{aligned}
(\mathbf{I} + \boldsymbol{\Sigma}_c\mathbf{R})^{-1} &= \begin{bmatrix} (\mathbf{I}_{k \times k} + \boldsymbol{\Sigma}_{c,CC}\mathbf{R}_{CC})^{-1} & -(\mathbf{I}_{k \times k} + \boldsymbol{\Sigma}_{c,CC}\mathbf{R}_{CC})^{-1} \boldsymbol{\Sigma}_{c,CC}\mathbf{R}_{CN} \\ \mathbf{0}_{(m-k) \times k} & \mathbf{I}_{(m-k) \times (m-k)} \end{bmatrix} \\
(\mathbf{I} + \boldsymbol{\Sigma}_c\mathbf{R})^{-1} \boldsymbol{\Sigma}_c &= \begin{bmatrix} (\mathbf{I}_{k \times k} + \boldsymbol{\Sigma}_{c,CC}\mathbf{R}_{CC})^{-1} & -(\mathbf{I}_{k \times k} + \boldsymbol{\Sigma}_{c,CC}\mathbf{R}_{CC})^{-1} \boldsymbol{\Sigma}_{c,CC}\mathbf{R}_{CN} \\ \mathbf{0}_{(m-k) \times k} & \mathbf{I}_{(m-k) \times (m-k)} \end{bmatrix} \begin{bmatrix} \boldsymbol{\Sigma}_{c,CC} & \mathbf{0}_{k \times (m-k)} \\ \mathbf{0}_{(m-k) \times k} & \mathbf{0}_{(m-k) \times (m-k)} \end{bmatrix} \\
&= \begin{bmatrix} (\mathbf{I}_{k \times k} + \boldsymbol{\Sigma}_{c,CC}\mathbf{R}_{CC})^{-1} \boldsymbol{\Sigma}_{c,CC} & \mathbf{0}_{k \times (m-k)} \\ \mathbf{0}_{(m-k) \times k} & \mathbf{0}_{(m-k) \times (m-k)} \end{bmatrix}
\end{aligned} \tag{88}$$

We can simplify this equation since $\Sigma_{\mathbf{c},\mathbf{CC}}$ is of full rank and is thus invertible:

$$\begin{aligned}
(\mathbf{I} + \Sigma_{\mathbf{c}}\mathbf{R})^{-1} \Sigma_{\mathbf{c}} &= \begin{bmatrix} \left(\Sigma_{\mathbf{c},\mathbf{CC}}^{-1} (\mathbf{I}_{k \times k} + \Sigma_{\mathbf{c},\mathbf{CC}}\mathbf{R}_{\mathbf{CC}})\right)^{-1} & \mathbf{0}_{k \times (m-k)} \\ \mathbf{0}_{(m-k) \times k} & \mathbf{0}_{(m-k) \times (m-k)} \end{bmatrix} \\
&= \begin{bmatrix} \left(\Sigma_{\mathbf{c},\mathbf{CC}}^{-1} + \mathbf{R}_{\mathbf{CC}}\right)^{-1} & \mathbf{0}_{k \times (m-k)} \\ \mathbf{0}_{(m-k) \times k} & \mathbf{0}_{(m-k) \times (m-k)} \end{bmatrix} \\
\hat{\mathbf{z}}^\top (\mathbf{I} + \Sigma_{\mathbf{c}}\mathbf{R})^{-1} \Sigma_{\mathbf{c}} \hat{\mathbf{z}} &= \begin{bmatrix} \hat{\mathbf{z}}_{\mathbf{C}}^\top & \mathbf{s}_{\mathbf{N}}^\top \end{bmatrix} \begin{bmatrix} \left(\Sigma_{\mathbf{c},\mathbf{CC}}^{-1} + \mathbf{R}_{\mathbf{CC}}\right)^{-1} & \mathbf{0}_{k \times (m-k)} \\ \mathbf{0}_{(m-k) \times k} & \mathbf{0}_{(m-k) \times (m-k)} \end{bmatrix} \begin{bmatrix} \hat{\mathbf{z}}_{\mathbf{C}} \\ \hat{\mathbf{z}}_{\mathbf{N}} \end{bmatrix} \\
&= \hat{\mathbf{z}}_{\mathbf{C}}^\top \left(\Sigma_{\mathbf{c},\mathbf{CC}}^{-1} + \mathbf{R}_{\mathbf{CC}}\right)^{-1} \hat{\mathbf{z}}_{\mathbf{C}}
\end{aligned} \tag{89}$$

We can also simplify the determinant in Equation 83:

$$\begin{aligned}
|\mathbf{I} + \Sigma_{\mathbf{c}}\mathbf{R}| &= \begin{vmatrix} \mathbf{I}_{k \times k} + \Sigma_{\mathbf{c},\mathbf{CC}}\mathbf{R}_{\mathbf{CC}} & \Sigma_{\mathbf{c},\mathbf{CC}}\mathbf{R}_{\mathbf{CN}} \\ \mathbf{0}_{(m-k) \times k} & \mathbf{I}_{(m-k) \times (m-k)} \end{vmatrix} \\
&= |\mathbf{I}_{k \times k} + \Sigma_{\mathbf{c},\mathbf{CC}}\mathbf{R}_{\mathbf{CC}}| |\mathbf{I}_{(m-k) \times (m-k)}| \\
&= |\mathbf{I}_{k \times k} + \Sigma_{\mathbf{c},\mathbf{CC}}\mathbf{R}_{\mathbf{CC}}|
\end{aligned} \tag{90}$$

We have thus shown that evaluating the Bayes Factor with only putative causal markers is mathematically equivalent to evaluating with all markers. Thus:

$$\text{BF}(\mathbf{c} : \mathbf{0}) = |\mathbf{I}_{k \times k} + \Sigma_{\mathbf{c},\mathbf{CC}}\mathbf{R}_{\mathbf{CC}}|^{-(1/2)} \exp\left(\frac{1}{2} \hat{\mathbf{z}}_{\mathbf{C}}^\top \left(\Sigma_{\mathbf{c},\mathbf{CC}}^{-1} + \mathbf{R}_{\mathbf{CC}}\right)^{-1} \hat{\mathbf{z}}_{\mathbf{C}}\right) \tag{91}$$

5.5 The hyperparameters in terms of heritability

The model takes a number of hyperparameters specifying the variances and covariances of the association statistics. We reparameterize the hyperparameters in terms of the QTL heritability of the locus h_β^2 and the AS heritability of the locus h_ϕ^2 .

First, we look at $\sigma_{c,\beta}^2$ and $\sigma_{c,\phi}^2$, given k expected causal variants. Let σ_y^2 be the overall variance of the QTL phenotype across the individuals. Since the heritability h_β is the proportion of the variance attributed to the causal variance, the average variance of a causal marker's QTL effect size is given by:

$$\sigma_{c,\beta}^2 = \frac{h_\beta^2}{lk} \sigma_y^2 \tag{92}$$

where l is the average LD score between the given marker and the other causal markers. Similarly, the variance of the AS effect size is given by:

$$\sigma_{c,\phi}^2 = \frac{h_\phi^2}{lk} \sigma_w^2 \tag{93}$$

However, in the case of the AS phenotype, where the quality of data varies considerably among individuals, we must take into account the variance introduced by sampling. As we recall, the

variance of the observed phenotype for a given individual j under a beta-binomial model is:

$$\begin{aligned}\text{Var}(\hat{w}_j) &\approx \frac{2}{c_j}(1 + \cosh(w_j^*))(1 + \rho_e(c_j - 1)) \\ \hat{w}_j^* &\approx \frac{\hat{w}_j}{1 + \frac{1}{c_j}(1 + \rho_e(c_j - 1))}\end{aligned}\tag{94}$$

We now derive an estimator $\hat{\sigma}_w^2 = \mathbb{E}[\text{Var}(\hat{w}_j)]$ for the total expected variance of the observed phenotype, across individuals. As an approximation, we substitute the individual read coverage for the expected read coverage \bar{c} , assuming that $\mathbb{E}[\text{Var}(\hat{w}_j|c_j)] \approx \mathbb{E}[\text{Var}(\hat{w}_j|\mathbb{E}[c])]$. Thus:

$$\hat{\sigma}_w^2 = \frac{2}{\bar{c}}(1 + \mathbb{E}[\cosh(w_j^*)])(1 + \rho_e(\bar{c} - 1))\tag{95}$$

Since we model \hat{w}_j as normally distributed, \hat{w}_j^* is approximately normally distributed. We now find $\mathbb{E}[\cosh(w_j^*)]$. Given a normally distributed zero-mean variable $A \sim \mathcal{N}(0, \sigma^2)$, the even-numbered moments are given by:

$$\mathbb{E}[A^{2p}] = \sigma^{2p}(2p - 1)!!\tag{96}$$

Taking the Taylor expansion of cosh and using linearity of expectation:

$$\begin{aligned}\mathbb{E}[\cosh(A)] &= \sum_{p=0}^{\infty} \frac{\mathbb{E}[A^{2p}]}{2p!} \\ &= \sum_{p=0}^{\infty} \frac{\sigma^{2p}(2p - 1)!!}{2p!} \\ &= \sum_{p=0}^{\infty} \frac{\sigma^{2p}}{2^p p!} \\ &= \sum_{p=0}^{\infty} \left(\frac{\sigma^2}{2}\right)^p \frac{1}{p!} \\ &= \exp\left(\frac{\sigma^2}{2}\right)\end{aligned}\tag{97}$$

Substituting this result back into the formula for $\hat{\sigma}_w^2$:

$$\hat{\sigma}_w^2 = \frac{2}{\bar{c}} \left(1 + \exp\left(\frac{\sigma_w^2}{2\left(1 + \frac{1}{\bar{c}}(1 + \rho_e(\bar{c} - 1))\right)^2}\right) \right) (1 + \rho_e(\bar{c} - 1))\tag{98}$$

Thus, the variance of the calculated AS effect size is given by:

$$\sigma_{c,\hat{\phi}}^2 = \left(\frac{h_{\hat{\phi}}^2}{lk}\right) \frac{2}{\bar{c}} \left(1 + \exp\left(\frac{\sigma_w^2}{2\left(1 + \frac{1}{\bar{c}}(1 + \rho_e(\bar{c} - 1))\right)^2}\right) \right) (1 + \rho_e(\bar{c} - 1))\tag{99}$$

We also define the observed AS heritability $h_{\hat{\phi}}^2$ such that

$$h_{\hat{\phi}}^2 = h_{\hat{\phi}}^2 \frac{\sigma_w^2}{\hat{\sigma}_w^2}\tag{100}$$

We now derive an expression for $r_{\beta\phi}$, the overall correlation between the QTL and AS statistics for a casual variant. We first find an expression, given causal variant i for the variance of $\hat{z}_{\beta,i}$. Since i is causal, the variance of $\hat{\beta}_i$ is a combination of the variance of a causal variant and the average phenotypic variation across individuals:

$$\begin{aligned}\text{Var}(\hat{\beta}_i) &= \sigma_{c,\beta}^2 + \frac{\sigma_y^2 - \sigma_{c,\beta}^2}{n} \\ \text{Var}(\hat{z}_{\beta,i}) &= \frac{\text{Var}(\hat{\beta}_i)}{(\sigma_y^2 - \sigma_{c,\beta}^2)/n} = n \frac{\sigma_{c,\beta}^2}{\sigma_y^2 - \sigma_{c,\beta}^2} + 1 \\ &= n \frac{h_\beta^2/(lk)}{1 - h_\beta^2/(lk)} + 1\end{aligned}\tag{101}$$

We model the QTL statistic of a causal variant as a zero-mean normal distribution:

$$\hat{z}_{\beta,i} \sim \mathcal{N}\left(0, n \frac{h_\beta^2/(lk)}{1 - h_\beta^2/(lk)} + 1\right)\tag{102}$$

We model the AS statistic in a similar manner:

$$\hat{z}_{\phi,i} \sim \mathcal{N}\left(0, n \frac{h_\phi^2/(lk)}{1 - h_\phi^2/(lk)} + 1\right)\tag{103}$$

We model the noise as independently distributed between the two statistics, but the causal variance as correlated with coefficient $r_{c,\beta\phi}$. Thus:

$$\begin{aligned}\hat{z}_{\beta,i} + \hat{z}_{\phi,i} &\sim \mathcal{N}\left(0, n \frac{h_\beta^2/(lk)}{1 - h_\beta^2/(lk)} + n \frac{h_\phi^2/(lk)}{1 - h_\phi^2/(lk)} + 2 + 2r_{c,\beta\phi} \sqrt{\left(n \frac{h_\beta^2/(lk)}{1 - h_\beta^2/(lk)}\right) \left(n \frac{h_\phi^2/(lk)}{1 - h_\phi^2/(lk)}\right)}\right) \\ \hat{z}_{\beta,i} - \hat{z}_{\phi,i} &\sim \mathcal{N}\left(0, n \frac{h_\beta^2/(lk)}{1 - h_\beta^2/(lk)} + n \frac{h_\phi^2/(lk)}{1 - h_\phi^2/(lk)} + 2 - 2r_{c,\beta\phi} \sqrt{\left(n \frac{h_\beta^2/(lk)}{1 - h_\beta^2/(lk)}\right) \left(n \frac{h_\phi^2/(lk)}{1 - h_\phi^2/(lk)}\right)}\right)\end{aligned}\tag{104}$$

We now find the covariance between \hat{z}_β and \hat{z}_ϕ . Since the distributions are zero-mean, the covariance is just $\mathbb{E}[\hat{z}_\beta \hat{z}_\phi]$. Expanding out this product:

$$\begin{aligned}\mathbb{E}[\hat{z}_\beta \hat{z}_\phi] &= \frac{1}{4} \left(\mathbb{E}[(\hat{z}_\beta + \hat{z}_\phi)^2] - \mathbb{E}[(\hat{z}_\beta - \hat{z}_\phi)^2] \right) \\ &= \frac{1}{4} (\text{Var}(\hat{z}_\beta + \hat{z}_\phi) - \text{Var}(\hat{z}_\beta - \hat{z}_\phi)) \\ &= r_{c,\beta\phi} \sqrt{\left(n \frac{h_\beta^2/(lk)}{1 - h_\beta^2/(lk)}\right) \left(n \frac{h_\phi^2/(lk)}{1 - h_\phi^2/(lk)}\right)}\end{aligned}\tag{105}$$

The correlation is thus:

$$\begin{aligned}
r_{\beta\phi} &= \frac{\mathbb{E}[\hat{z}_\beta \hat{z}_\phi]}{\sqrt{\text{Var}(\hat{z}_\beta) \text{Var}(\hat{z}_\phi)}} \\
&= r_{c,\beta\phi} \sqrt{\left(n \frac{h_\beta^2/(lk)}{1 - h_\beta^2/(lk)} \right) \left(n \frac{h_\phi^2/(lk)}{1 - h_\phi^2/(lk)} \right)} \\
&= \sqrt{\left(n \frac{h_\beta^2/(lk)}{1 - h_\beta^2/(lk)} + 1 \right) \left(n \frac{h_\phi^2/(lk)}{1 - h_\phi^2/(lk)} + 1 \right)} \\
&= r_{c,\beta\phi} \sqrt{\frac{n^2 h_\beta^2 h_\phi^2}{(lk + h_\beta^2(n-1))(lk + h_\phi^2(n-1))}}
\end{aligned} \tag{106}$$

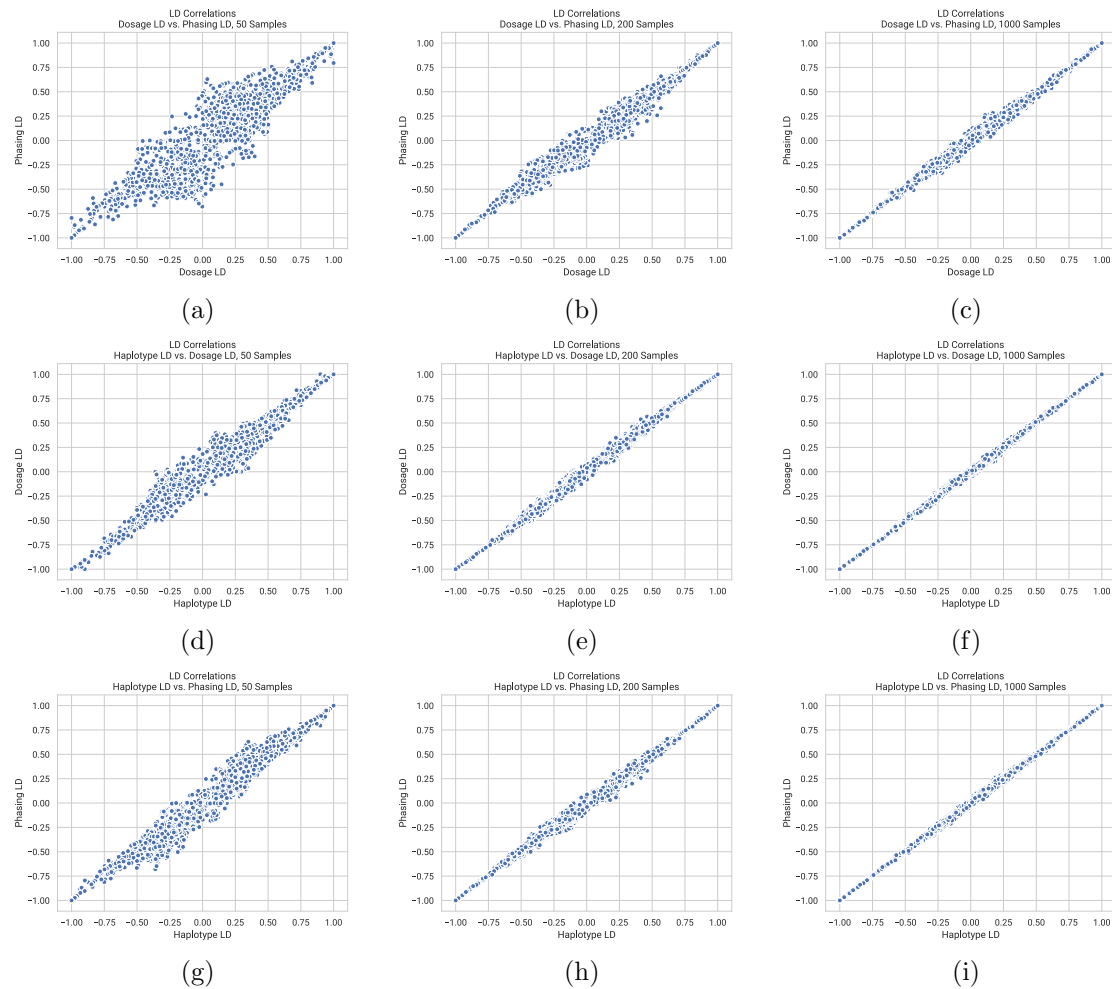


Figure S1: Comparison of estimators for linkage disequilibrium (LD) at various sample sizes. Estimators used were the dosage estimator, the phasing estimator, and the haplotype-level estimator. At each sample size, 100 markers from 1000-Genomes data were used to calculate LD using all three estimators. Plots of LD correlations are shown for each pair of estimators. (a-c) Dosage LD vs. phasing LD at 50, 200 and 1000 samples, respectively. (d-f) Haplotype-Specific LD vs. dosage LD at 50, 200 and 1000 samples, respectively. (g-i) Haplotype-Specific LD vs. phasing LD at 50, 200 and 1000 samples, respectively.

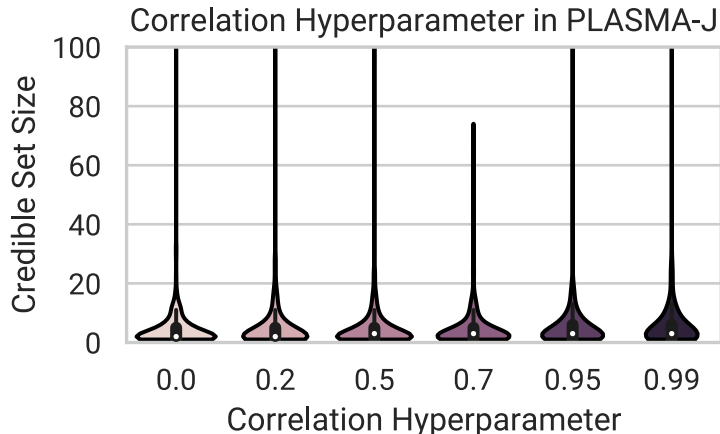


Figure S2: Effect of the jointness hyperparameter in PLASMA on fine-mapping. Fine-mapping of simulated loci was conducted with PLASMA-J under various values of the jointness parameter. Each value of the jointness parameter was evaluated using 500 simulated loci with one causal variant per locus. Plotted are the distributions of the 95% confidence credible sets.

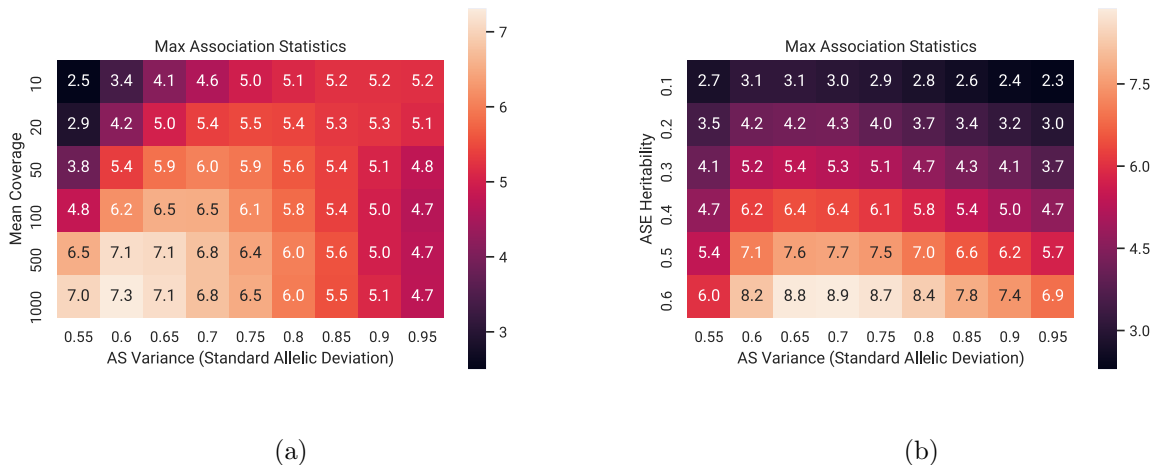


Figure S3: Mean top AS association statistics (z_ϕ) under various conditions of mean read coverage, heritability, and standard allelic deviation (a metric of the variance of the AS phenotype). Standard allelic deviation is defined as the allelic fraction f such that the range between $1 - f$ and f captures all loci with an AS phenotype within one standard deviation of the mean (balanced expression). Each square is the mean statistic over 500 simulated loci of 100 markers, with one causal variant. (a) Mean top z_ϕ as a function of standard allelic deviation and mean read coverage, with AS heritability set to 0.5. (b) Mean top z_ϕ as a function of standard allelic deviation and AS heritability, with mean read coverage set to 100 reads.

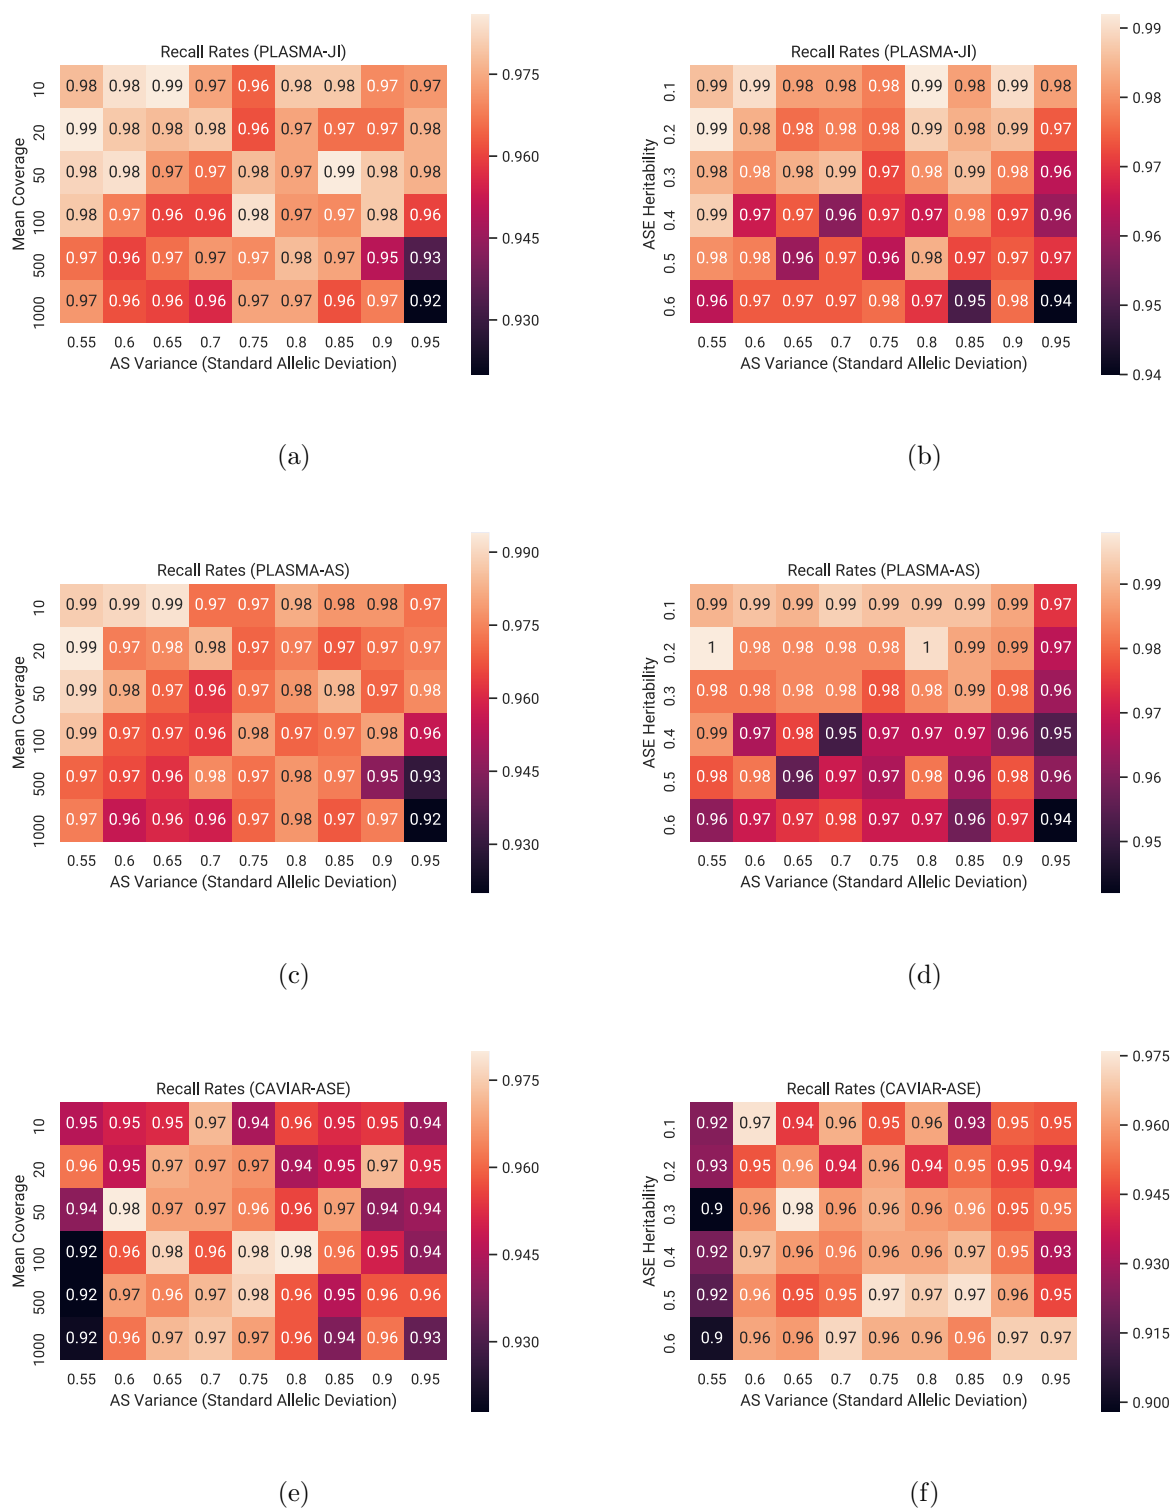


Figure S4: Recall rates for 95% credible sets under various conditions of mean read coverage, heritability, and standard allelic deviation (intensity of AS variance). Recall rate is defined as the expected proportion of causal markers included in the credible set. Each square is the mean statistic over 500 simulated loci of 100 markers, with one causal variant. (a, c, e) Recall rates as a function of standard allelic deviation and mean read coverage, with AS heritability set to 0.4. (d, f, h) Recall rates as a function of standard allelic deviation and AS heritability, with mean read coverage set to 100 reads.

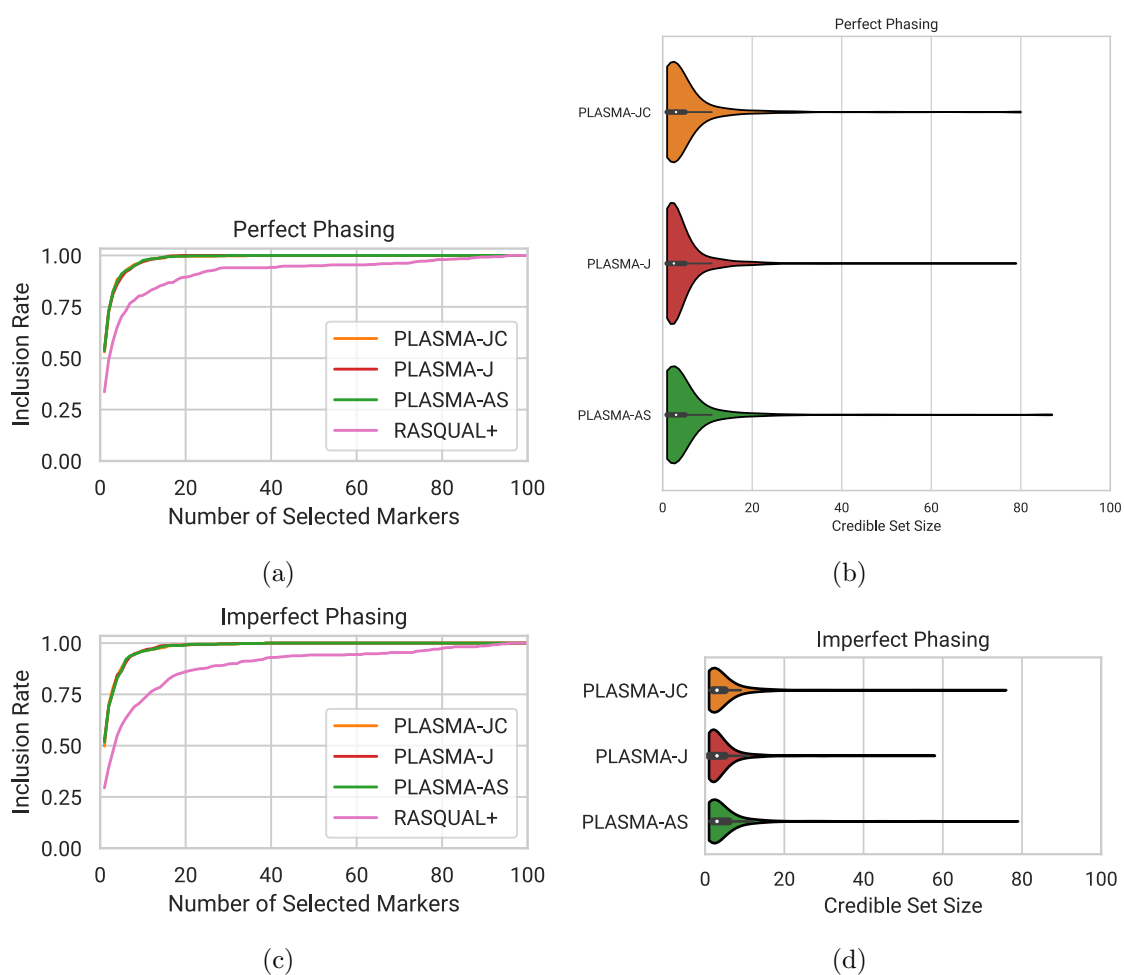


Figure S5: Impact of imperfect phasing on AS-based methods. Each dataset consists of 500 loci with one causal variant per locus. (a, b) Recall and 95% confidence credible sets under perfect phasing. (c, d) Recall and 95% confidence credible sets under imperfect phasing, with a phasing error z-score correction. The switch error and blip error were set to 0.152% and 0.165%, respectively.

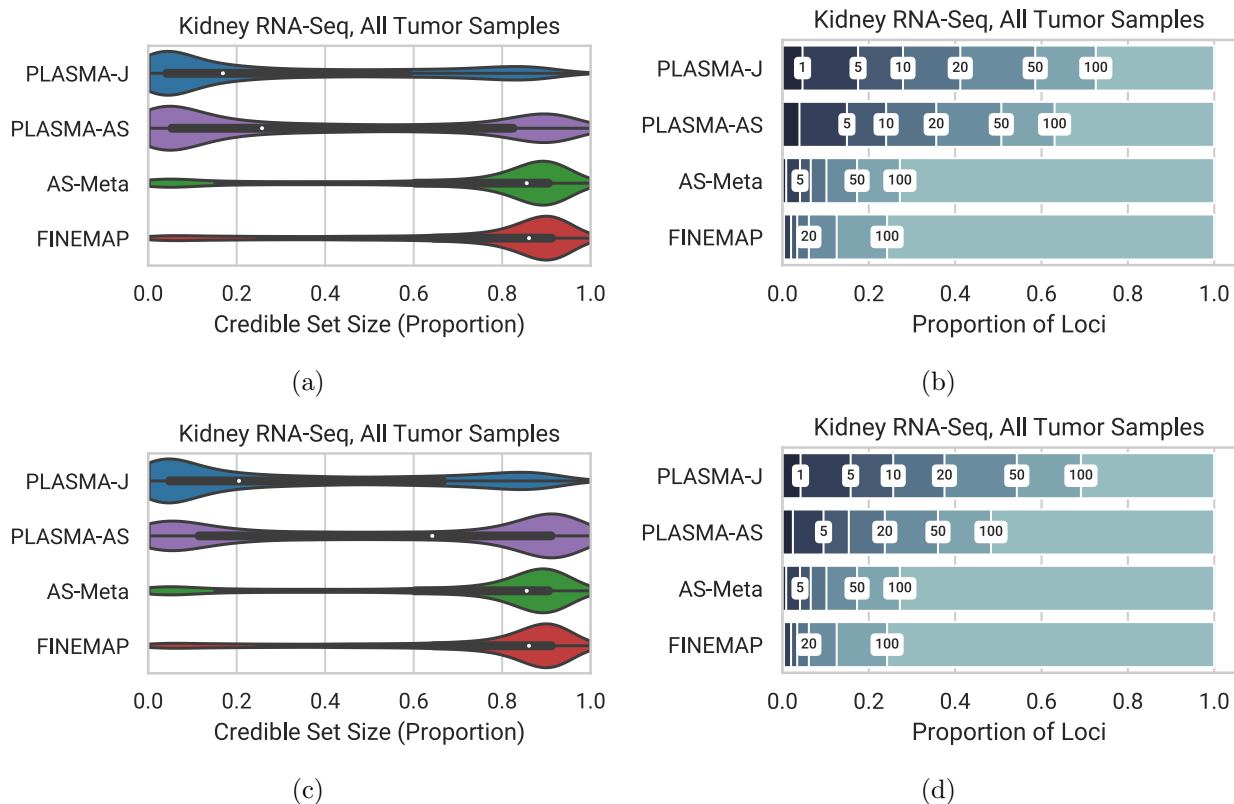


Figure S6: Effect of the AS heritability hyperparameter on 95% credible set sizes with tumor RNA-Seq data. (a) Set size distribution with the heritability hyperparameter at 50% (b) Cumulative set size distribution with the heritability hyperparameter at 50% (c) Set size distribution with the heritability hyperparameter at 5% (d) Cumulative set size distribution with the heritability hyperparameter at 5%

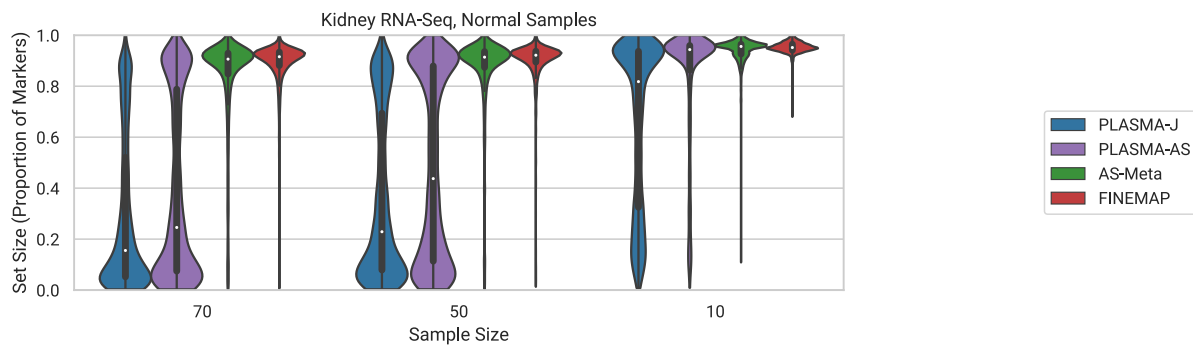


Figure S7: 95% credible set size distributions for randomly down-sampled kidney normal data with decreasing sample sizes.

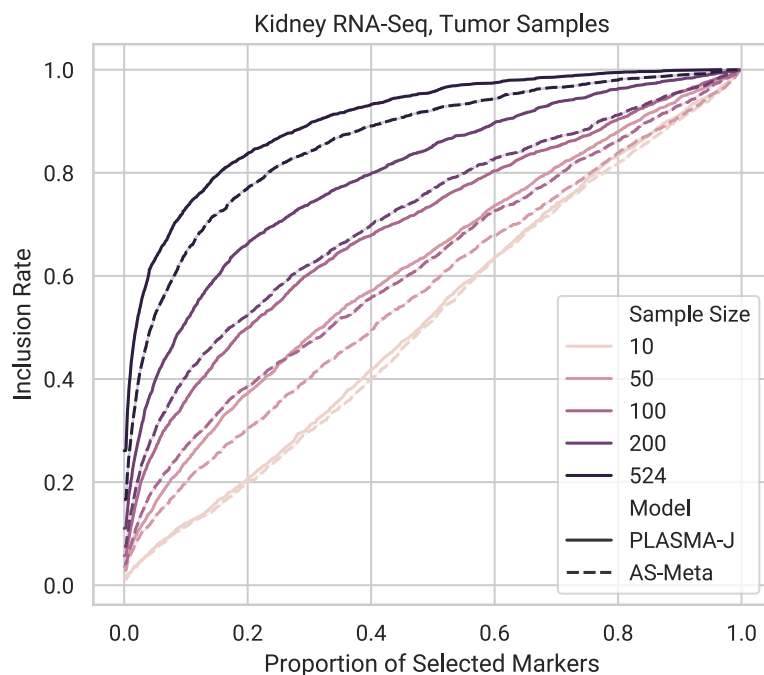


Figure S8: Inclusion curve for fine-mapped kidney tumor RNA-Seq data for PLASMA and AS-Meta across sample sizes. Inclusion was evaluated against a gold-standard set of markers, defined as all markers with at least 0.1 posterior probability under FINEMAP with all samples.

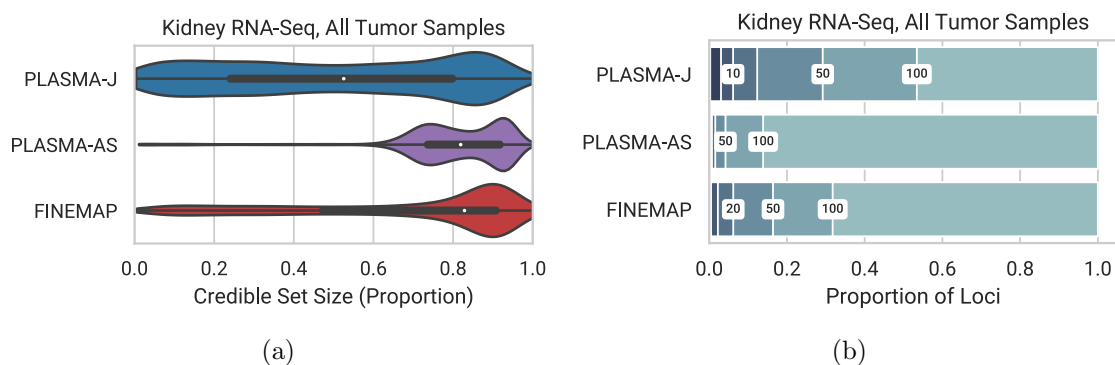


Figure S9: Fine-mapping of kidney RNA-Seq loci under the assumption of up to three causal variants. (Note that FINEMAP credible sets may not be calibrated under multiple causal variants.) (a) 95% confidence credible sets. (b) Proportions of loci with 95% credible set sizes within given thresholds. Thresholds used were 1, 5, 10, 20, 50, and 100.

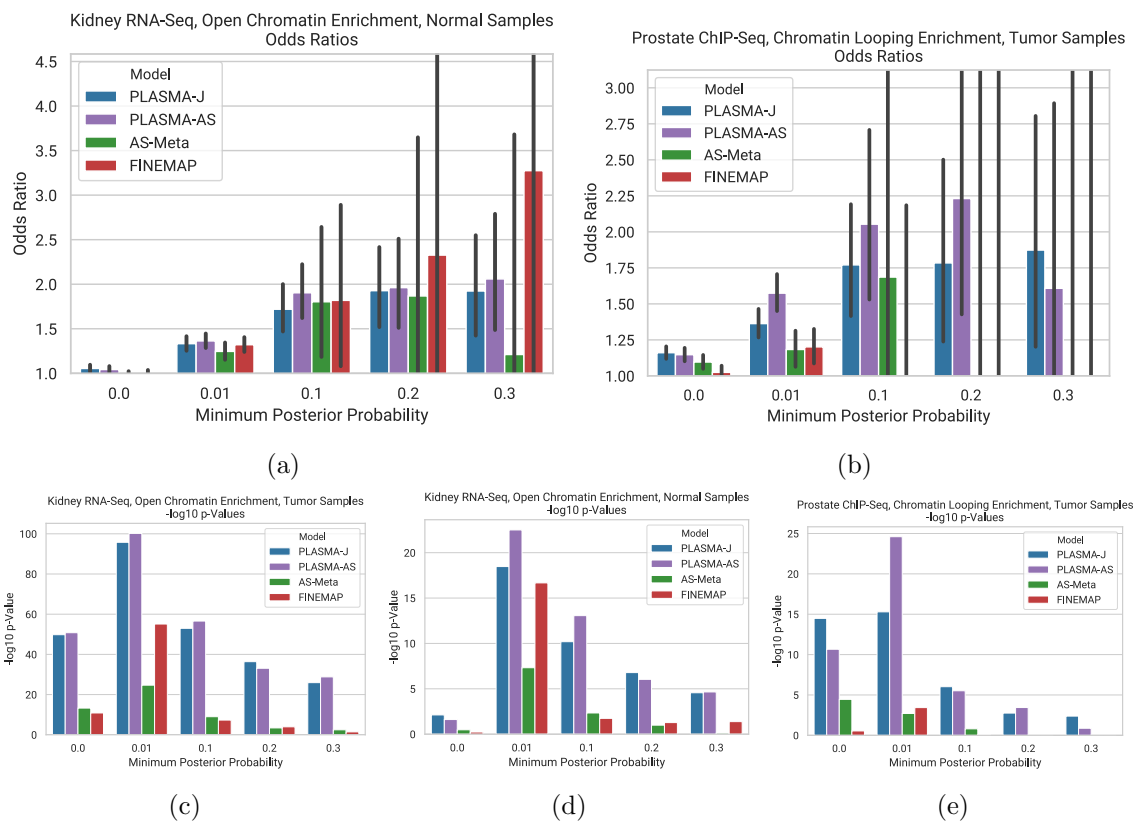


Figure S10: Enrichment odds ratios of PLASMA-J credible sets with functional features. The x-axis represents a minimum posterior probability threshold on the markers within credible sets, with a higher threshold indicating more confident but fewer markers. The markers in these thresholded credible sets were intersected with the functional feature in question. 95 % confidence intervals and p-values were calculated under Fisher's exact test, with a control of all markers in a locus. (a) Enrichment odds ratios of fine-mapped credible sets with open chromatin for kidney normal RNA-Seq data. (b) Enrichment odds ratios of fine-mapped credible sets with chromatin looping for prostate tumor H3K27ac ChIP-Seq data. (c, d) Enrichment p-values of kidney tumor and normal credible sets, respectively, with open chromatin. (e) Enrichment of fine-mapped credible sets with chromatin looping for prostate tumor H3K27ac ChIP-Seq data.

Table S1: Summary of hyperparameter values used in each set of simulation analyses, referenced by figure legend. Parentheses indicate multiple values tested.

Parameter	Figure 2	Figures 3a-c	Figures 3d-f
Number of markers per locus	100	100	100
Number of samples	100	100	100
Minimum minor allele frequency	0.01	0.01	0.01
AS overdispersion (ρ_e)	0.05	0.05	0.05
QTL Heritability (h_β^2)	0.05	0.05	0.05
AS Heritability (h_ϕ^2)	0.4	0.4	(0.1 to 0.6)
Manual PLASMA QTL Heritability Setting	N/A	N/A	N/A
Manual PLASMA AS Heritability Setting	N/A	N/A	N/A
Standard Allelic Deviation (d)	(0.6, 0.8)	(0.55 to 0.95)	(0.55 to 0.95)
Number of causal variants	1	1	1
Mean read coverage	100	(10 to 1000)	100
Min. causal variants in searched configurations	1	1	1
Max. causal variants in searched configs	1	1	1
Search probability closeness threshold	0.001	0.001	0.001
Search iterations convergence threshold	1000	1000	1000
Maximum search iterations	100000	100000	100000
Jointness parameter ($r_{c,\beta\phi}$)	0.	0.	0.
Phasing switch error	0.	0.	0.
Phasing blip error	0.	0.	0.

Parameter	Figure 4	Figure S2	Figure S5
Number of markers per locus	100	100	100
Number of samples	100	100	100
Minimum minor allele frequency	0.01	0.01	0.01
AS overdispersion (ρ_e)	0.05	0.05	0.05
QTL Heritability (h_β^2)	0.05	0.05	0.05
AS Heritability (h_ϕ^2)	0.4	0.4	0.4
Manual PLASMA QTL Heritability Setting	0.02	N/A	N/A
Manual PLASMA AS Heritability Setting	0.2	N/A	N/A
Standard Allelic Deviation (d)	0.7	0.7	0.7
Number of causal variants	2	1	1
Mean read coverage	100	100	100
Min. causal variants in searched configurations	1	1	1
Max. causal variants in searched configs	3	1	1
Search probability closeness threshold	0.001	0.001	0.001
Search iterations convergence threshold	1000	1000	1000
Maximum search iterations	100000	100000	100000
Jointness parameter ($r_{c,\beta\phi}$)	0.	(0. to 0.99)	0.
Phasing switch error	0.	0.	0.00152
Phasing blip error	0.	0.	0.00165