

Allele-Specific QTL Fine-Mapping with PLASMA

Austin T. Wang^{1,2,3}, Anamay Shetty^{3,4}, Edward O'Connor³, Connor Bell³, Mark M. Pomerantz³, Matthew L. Freedman^{3,5,6}, Alexander Gusev^{3,7}

¹Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02142, USA

²Department of Biology, Massachusetts Institute of Technology, Cambridge, MA 02142, USA

³Department of Medical Oncology, Dana-Farber Cancer Institute, Boston, MA 02215, USA

⁴Cambridge University, Cambridge, UK

⁵The Eli and Edythe L. Broad Institute, Cambridge, MA 02142, USA

⁶Center for Functional Cancer Epigenetics, Dana-Farber Cancer Institute, Boston, MA 02215, USA

⁷Brigham & Women's Hospital, Division of Genetics, Boston, MA 02215, USA

Abstract

Although quantitative trait locus (QTL) associations have been identified for many molecular traits such as gene expression, it remains challenging to distinguish the causal nucleotide from nearby variants. In addition to traditional QTLs by association, allele-specific (AS) QTLs are a powerful measure of cis-regulation that are largely concordant with traditional QTLs, and can be less susceptible to technical/environmental noise. However, existing asQTL analysis methods do not produce probabilities of causality for each marker, and do not take into account correlations among markers at a locus in linkage disequilibrium (LD). We introduce PLASMA (PopuLation Allele-Specific Mapping), a novel, LD-aware method that integrates QTL and asQTL information to fine-map causal regulatory variants while drawing power from both the number of individuals and the number of allelic reads per individual. We demonstrate through simulations that PLASMA successfully detects causal variants over a wide range of genetic architectures. We apply PLASMA to RNA-Seq data from 524 kidney tumor samples and show that over 13 percent of loci can be fine-mapped to within 5 causal variants, compared less than 2 percent of loci using existing QTL-based fine-mapping. PLASMA furthermore achieves a greater power at 50 samples than conventional QTL fine-mapping does at over 500 samples. Overall, PLASMA achieves a 6.4-fold reduction in median 95% credible set size compared to existing QTL-based fine-mapping. We additionally apply PLASMA to H3K27AC ChIP-Seq from 28 prostate tumor/normal samples and demonstrate that PLASMA is able to prioritize markers even at small samples, with PLASMA achieving a 1.4-fold reduction in median 95% credible set sizes over existing QTL-based fine-mapping. Variants in the PLASMA credible sets for RNA-Seq and ChIP-Seq were enriched for open chromatin and chromatin looping (respectively) at a comparable or greater degree than credible variants from existing methods, while containing far fewer markers. Our results demonstrate how integrating AS activity can substantially improve the detection of causal variants from existing molecular data and at low sample size.

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 individual-level molecular activity (such as gene expression) in the population [1–4]. These 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 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 expression regardless of source. The AS effect-size has also been shown to be highly correlated to conventional QTL effect-sizes, implying that both features typically share 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 Results

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, indi-

cated 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]. In practice, we also mitigate the effect of mapping bias by running 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 (see Methods for quantitative explanations of the statistics). Here, the causal marker’s alternative allele causes higher expression compared to the wild-type allele. We see that 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:

$$\mathbf{y} = \mathbf{x}_i\beta_i + \boldsymbol{\epsilon}_i \quad (1)$$

From this effect size, PLASMA calculates z_β , the QTL association statistic (See Methods for the precise relationship with β). Note that this statistic is not dependent on haplotype-specific data.

On the other hand, looking at the heterozygotes, we see that 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 direction of the phase (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:

$$\mathbf{w} = \mathbf{v}_i\phi_i + \boldsymbol{\zeta}_i \quad (2)$$

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, see Methods for the precise weighing scheme).

These QTL and AS association statistics, together with the local LD matrix, are then jointly used to fine-map the locus (Figure 1c, see Methods for the fine-mapping model). Since PLASMA models both z_β and z_ϕ as a linear combination of genotypes, z_β and z_ϕ have identical LD (see Supplemental Methods 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 measure the same underlying effect, the two statistics have independent noise because the intra-heterozygous variance 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 we find 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 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” as the standard deviation of the AS phenotype w , quantified on the allelic fraction scale (between 0.5 and 1). Importantly, this metric 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.

As the performance of standard QTL association models is well established, we first focused on performance of our proposed AS statistic. Figure S1a 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 S1b 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 (standard allelic deviation of 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 (standard allelic deviation greater than 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 46 (See Supplemental Methods for a mathematical relationship between total variance and sampling error.)


2.3 Comparison with existing methods in simulation

Next, we compare PLASMA’s fine-mapping performance with existing fine-mapping methods. We test two different “flavors” of PLASMA, “PLASMA-JI” and “PLASMA-AS.” The PLASMA-JI (Joint-Independent) flavor 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” flavor is restricted to only AS data. As a baseline, we also compared 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


Align Reads




Phase Genotypes



Aggregate Reads



(a)



(b)

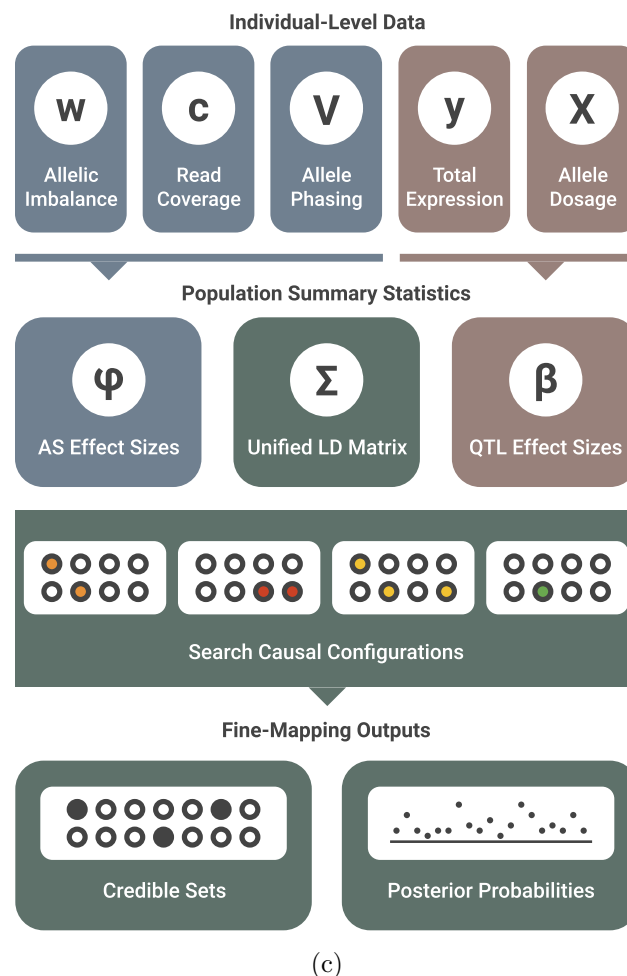


Figure 1: Overview of the PLASMA method. (a) Pre-processing of sequence-based data. First, reads are mapped to the sample's genotype. Reads intersecting markers are colored. Then, the sample's genotype is phased. Reads intersecting heterozygous markers can then be mapped to a particular haplotype. Lastly, reads across the locus are aggregated in an allele-specific manner. We visualize this data by a schematic, where 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.

data [14–16]. We furthermore compare the flavors of PLASMA against the only other publicly-released fine-mapping method (to our knowledge) that integrates AS data described in the preprint of Zou *et al.*, 2018 [24]. This unnamed method, which we denote as “CAVIAR-ASE,” utilizes the association between SNP heterozygosity and a binary indicator of allelic imbalance. By binarizing allelic imbalance, CAVIAR-ASE is expected to lose power relative to treating imbalance as a quantitative phenotype but may be more robust to spurious AS signal. Furthermore, CAVIAR-ASE utilizes only indicators of heterozygosity, rather than marker phasing. CAVIAR-ASE can therefore be used with unphased genotypes, but at the expense of being unable to leverage the direction of the allelic effect.

First, we evaluate how well each model prioritizes candidate causal markers using simulated loci with one causal variant. 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 c show inclusion plots at low and high AS variance, respectively. As expected, the QTL-Only flavor and the CAVIAR methods are indistinguishable and do not vary with AS variance (thus, we do not include CAVIAR in further results). Furthermore, we see that PLASMA-JI and PLASMA-AS perform similarly at both levels of AS variance. Lastly, we see a dependency of CAVIAR-ASE’s performance on the degree on AS variance.

Second, we evaluate the ability of each model to rule out likely non-causal markers in simulated loci with one causal variant. To do so, we directly compare the distributions of the 95% confidence credible sets, with smaller sets indicating higher specificity. Figures 2b and d show distribution plots at low and high AS variance, respectively. At low variance, PLASMA-JI offers the smallest mean credible set size (10.4), followed by PLASMA-AS (10.6), then CAVIAR-ASE (44.6), and lastly CAVIAR (78.2) and QTL-Only (77.0). The AS-based flavors of PLASMA are resistant to changes in AS variance, with mean credible set sizes at high AS variance of 9.6 for PLASMA-JI and 10.4 for PLASMA-AS. In contrast, the performance of CAVIAR-ASE varies significantly with the degree of AS variance, even when the underlying signal (coverage and heritability) is constant, with a mean set size of 67.4 at high variance. This sensitivity may be due to the fact that CAVIAR-ASE does not incorporate marker phasing, and thus must rely solely on the intensity of imbalance, rather than the direction of imbalance.

Third, 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-d show mean credible set sizes as a function of AS variance and coverage, and Figures 3e-h show mean credible set sizes as a function of AS variance and AS heritability. In terms of the range of set sizes, the PLASMA-JI flavor performs the best (7.4 variants on average at best conditions), followed by the PLASMA-AS flavor (7.4 at best conditions), and lastly the CAVIAR-ASE method (18.0 at best conditions). Generally speaking, all methods show results consistent with the behavior of z_ϕ in Figure S1. 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, CAVIAR-ASE tends to struggle at low AS variance, especially apparent at a standard allelic deviation of 0.55, with a mean set size of 52.3 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 S2 shows that the 95% credible set sizes have at least a 95% chance of including the causal variant.

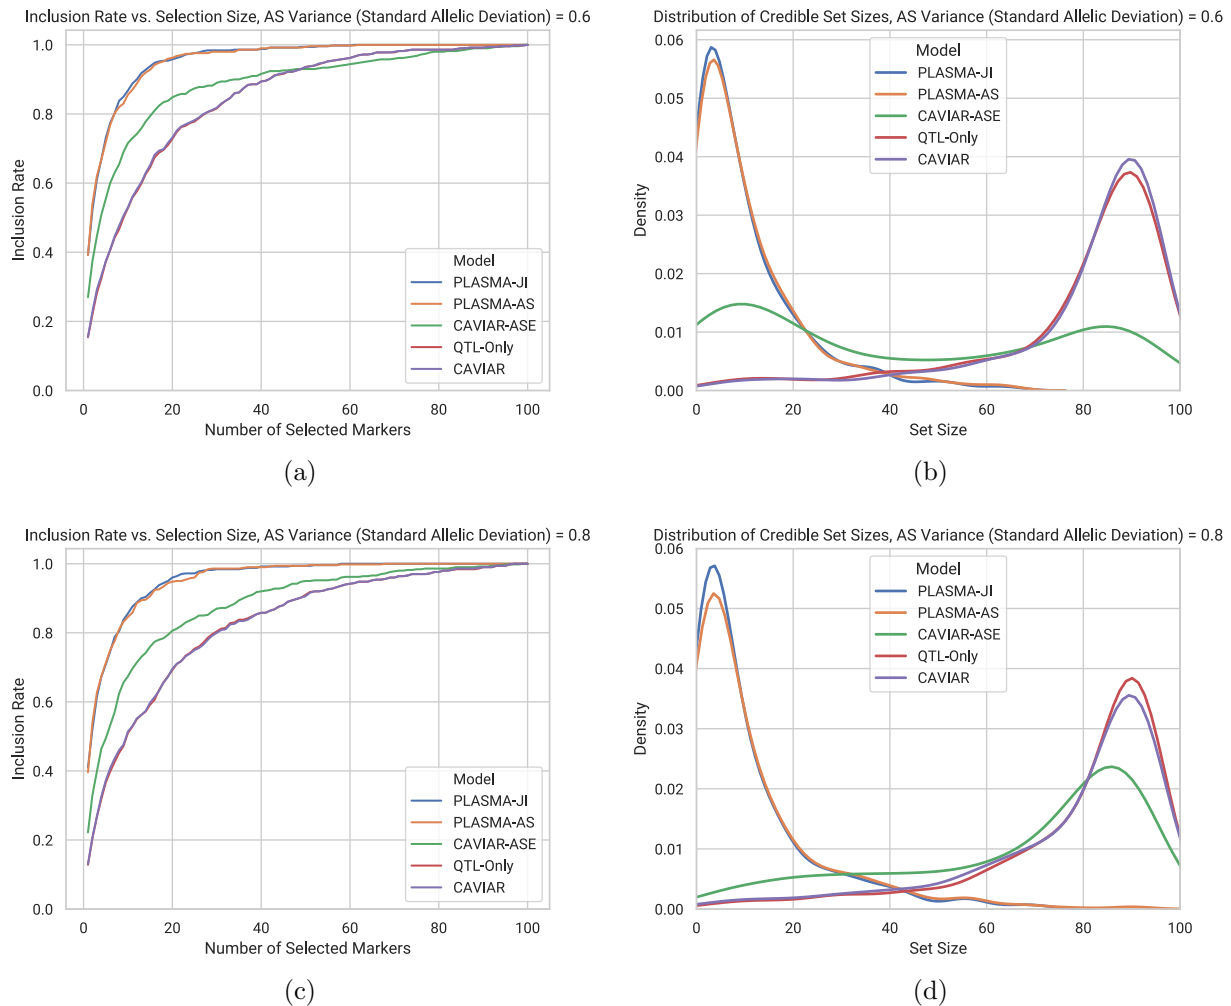


Figure 2: Comparison of fine-mapping methods on two sets of simulated loci of 90 markers and 100 samples each. These two sets differ in the variance of the AS phenotype, but with a fixed mean coverage and AS heritability. (a) Inclusion curve at AS variance (0.6 standard allelic deviation). The y-axis 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) Distribution of 95% confidence credible set sizes at low AS variance (0.6 standard allelic deviation) (c) Inclusion curve at high AS variance (0.8 standard allelic deviation) (d) 95% confidence credible set size distribution at high AS variance (0.8 standard allelic deviation)

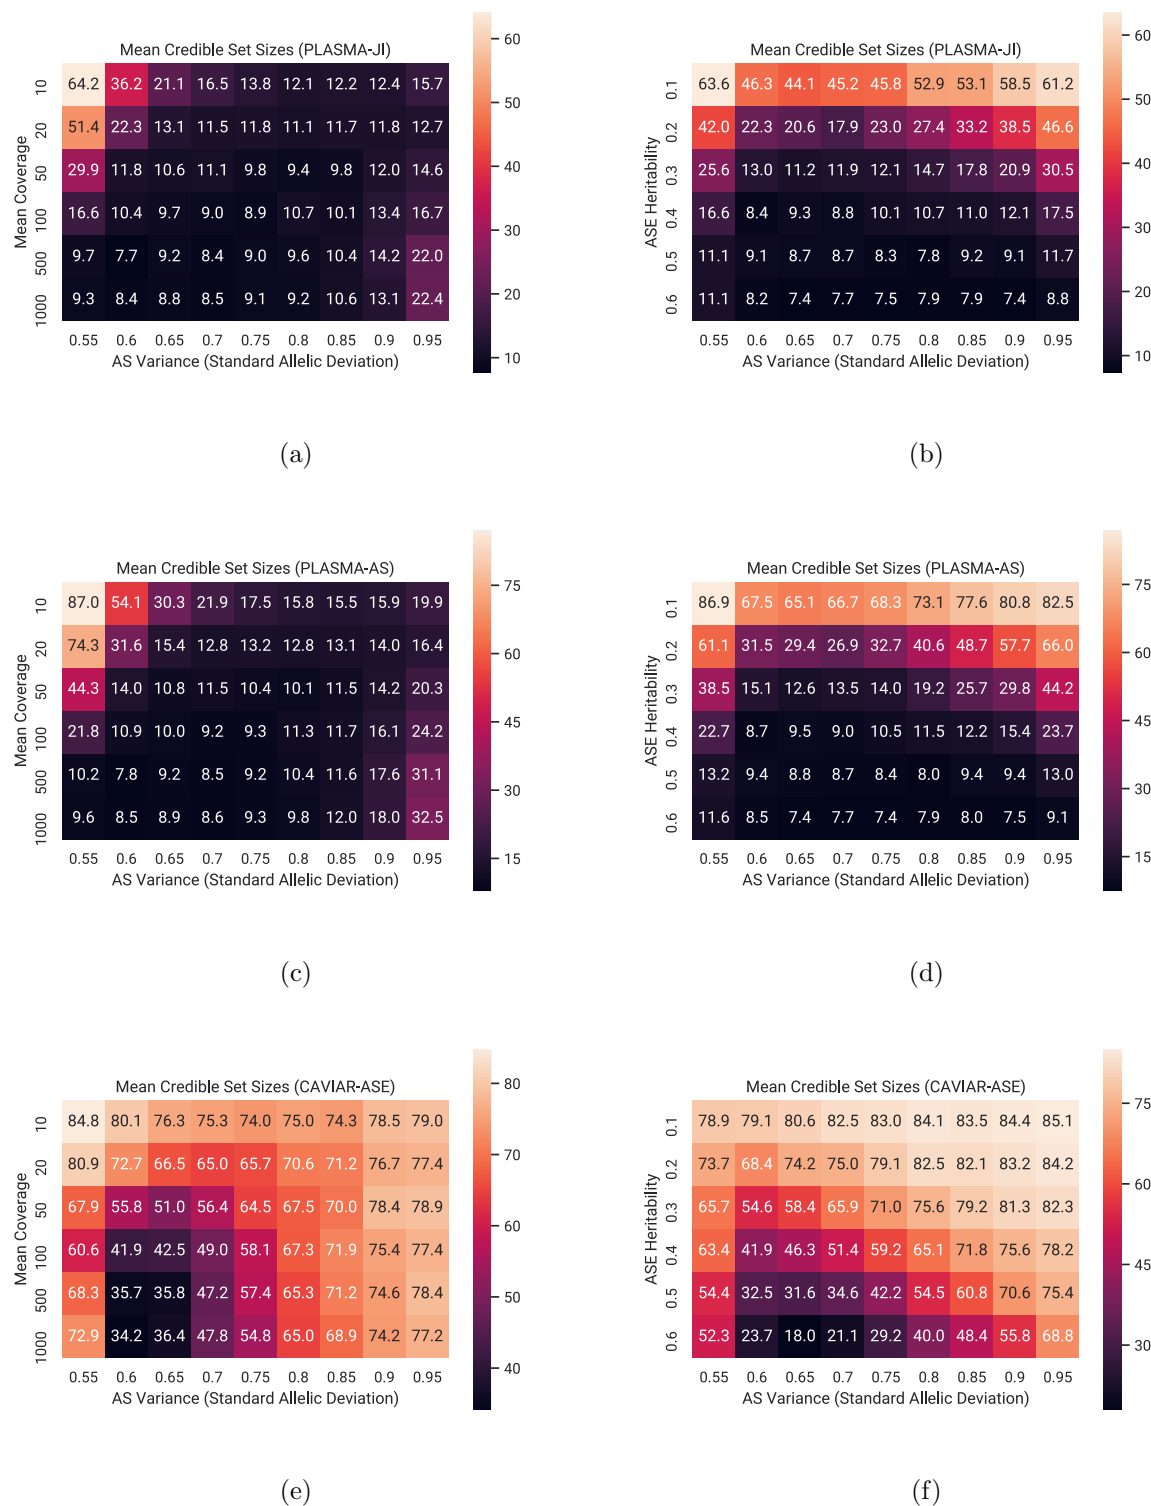


Figure 3: Comparison of the mean 95% confidence credible set sizes across different AS fine-mapping methods. Each square is the mean set size calculated over 500 simulated loci of 100 markers and 100 samples each, with one causal variant. (a, c, e) Mean credible set sizes as a function of standard allelic deviation and mean read coverage, with AS heritability set to 0.5. (b, d, 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 varying numbers of causal variants. Each locus is of 100 markers and 95 samples each, with AS heritability set to 0.4, QTL heritability set to 0.05, and mean read coverage set to 100 reads. (a) Inclusion curve for 2 causal variants. (b) Distribution of 95% confidence credible set sizes for 2 causal variants.

2.4 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 flavor of PLASMA. Figure 4 shows the inclusion curve and the distribution of 95% confidence credible set sizes for each flavor. For both curves, a success is defined as the inclusion of both causal variants. For PLASMA-JI, the median credible set size increases from a mean of 9.5 for one causal variant to 87.0 for two causal variants. This apparent decrease in power 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. Nevertheless, PLASMA-JI and PLASMA-AS deliver an improvement over QTL-Only fine-mapping, with mean credible set sizes of 87.0, 93.0, 95.1, for PLASMA-JI, PLASMA-AS, and QTL-Only (respectively). Due to the difficulty of fine-mapping multiple causal variants [10], along with the estimate that over 75% of loci do not display allelic heterogeneity, further analyses on experimental data were performed under the one-causal-variant assumption.

Unlike the single causal variant case, where all model hyperparameters were inferred from simulation parameters, the causal variance hyperparameters in this case were manually calibrated. We believe that this need for calibration is due to linkage disequilibrium obfuscating the relationship between causal effect sizes and total heritability at a locus. (See Supplemental Methods for information about hyperparameter estimation.) The results shown in this section is calibrated such that the recall rates for 95% confidence credible sets are 0.95, 0.95, and 0.978 for the PLASMA-JI, QTL-Only, and PLASMA-AS flavors, respectively.

2.5 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 [25]. The data was processed through a rigorous QC pipeline to account for mapping biases based on established best practices [19, 22]. Figure 5 shows credible set size distribution plots for tumor and normal data under a

1 causal variant assumption. Among the tumor samples ($N=524$), 22.4% of loci are fine-mapped within 10 variants with PLASMA-JI, while 3.1% of loci are fine-mapped within 10 variants with QTL-Only fine-mapping (Table S1a). Furthermore, PLASMA-JI achieves a median credible set size for 45 variants, whereas QTL-Only achieves a median credible set size of 289 variants (Table S2a). We also see a significant improvement over CAVIAR-ASE, which has 4.7% of loci fine-mapped within 10 causal variants, and a median credible set size of 292. Results for normal samples ($N=70$) have a similar trend, with 7.0%, 0.2%, and 0.5% of loci fine-mapped within 10 causal variants, for PLASMA, QTL-Only, and CAVIAR-ASE respectively (Table S1b). Median credible set sizes of 67, 374, and 348 variants, for PLASMA, QTL-Only, and CAVIAR-ASE respectively (Table S2b). The 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, we also ran 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 S3 and Table S3.

To investigate how the methods perform at lower sample sizes, we randomly subsample individuals prior to fine-mapping. Figure 6 plots the credible set size distributions for PLASMA-JI, QTL-Only, and CAVIAR-ASE at various sample sizes. In terms of loci fine-mapped to credible set sizes within 10 variants in tumor (Table S4a), PLASMA with 50 samples has approximately the same power as QTL-Only fine-mapping with 500 samples. In terms of median credible set size, PLASMA with 10 samples has about the same power as QTL-Only fine-mapping with 500 samples (Table S4c). All methods increase in credible set size as the sample size is restricted but the relative gain of PLASMA over the other methods decreases across sample sizes. PLASMA yields a 6.4-fold decrease in median credible set size over QTL-Only fine-mapping at 524 samples, but a 1.3-fold decrease at 10 samples (Table S4c). This implies 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 287 in tumor, compared to a median set size of 381 with QTL-Only fine-mapping. We furthermore see in Figure 6b that at a given sample size, PLASMA has higher power for normal samples than for tumor samples, which we believe is 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 we do not know the true causal variants in each locus, as a proxy we use markers with a posterior probability of at least 0.1 when fine-mapped with the QTL-Only method on all samples. We note that this will strongly bias the credible set in favor of the QTL model and thus do not compare to the QTL-Only model. In Figure S4, we again see that PLASMA is more effective than CAVIAR-ASE at each sample size. In terms of loci fine-mapped to within 10 variants in tumor (Table S4a), PLASMA with 100 samples has greater power than CAVIAR-ASE fine-mapping with 500 samples. At a given sample size, PLASMA is thus better able to prioritize variants that will be ranked highly in larger studies.

Lastly, we look at how PLASMA prioritizes experimentally-verified causal variants at GWAS risk loci. Figure 7 shows the strength AS and QTL associations for DPF3 and SCARB1, genes in two kidney GWAS loci that have verified causal variants [23,26]. 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, we see that 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 accurate and confident prioritization of causal variants, PLASMA can substantially reduce the difficulty of experimentally validating causal variants.

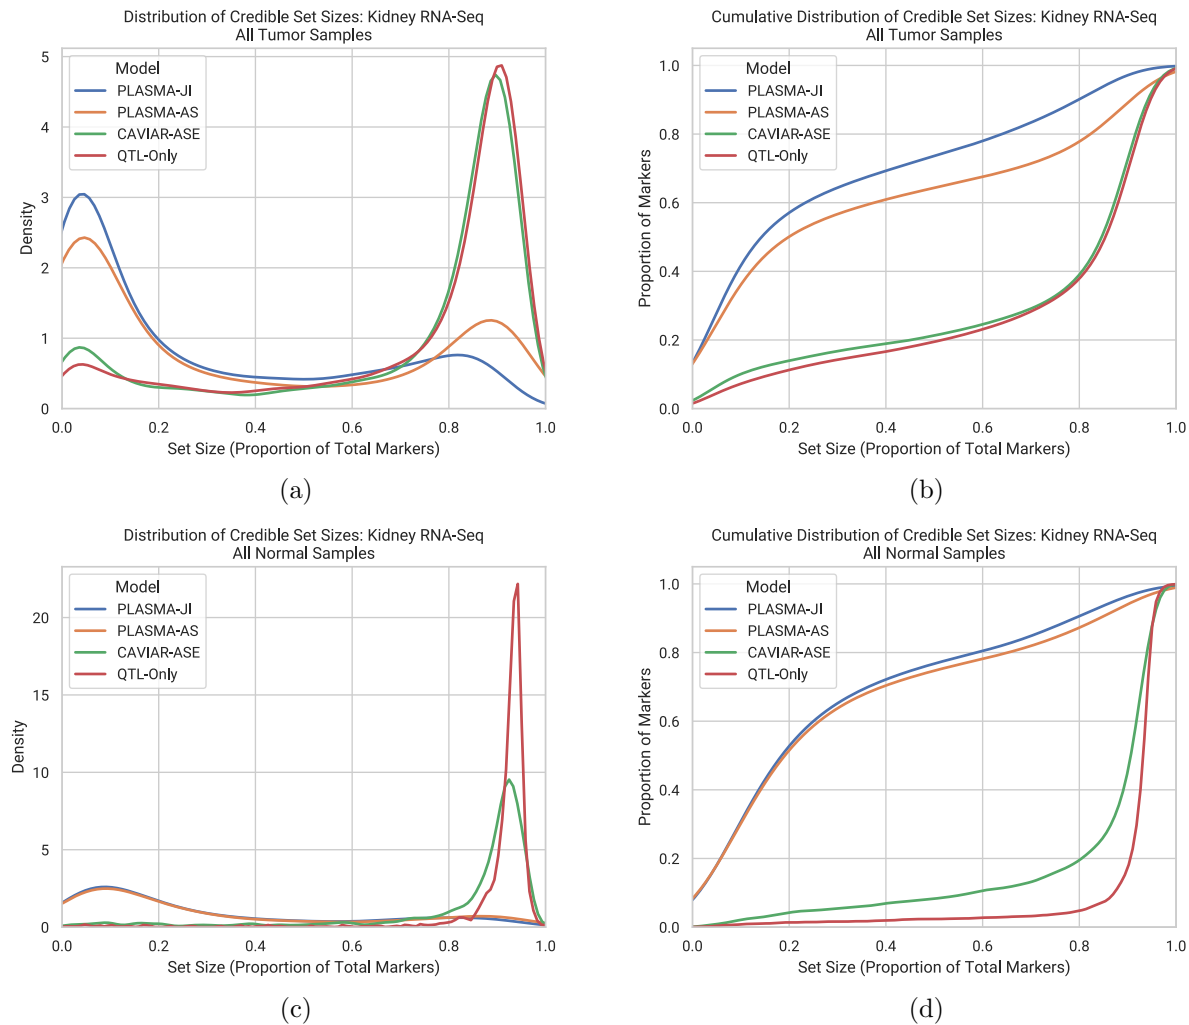


Figure 5: Comparison of the distribution of 95% confidence credible set sizes across loci in kidney tumor and normal samples, with an allelic imbalance false discovery rate of 0.05. Expression was measured as RNA-Seq read counts mapped to phased genotypes. Fine-mapping was conducted with 500 tumor samples and 90 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) Cumulative distribution of credible set sizes for tumor and samples, respectively.

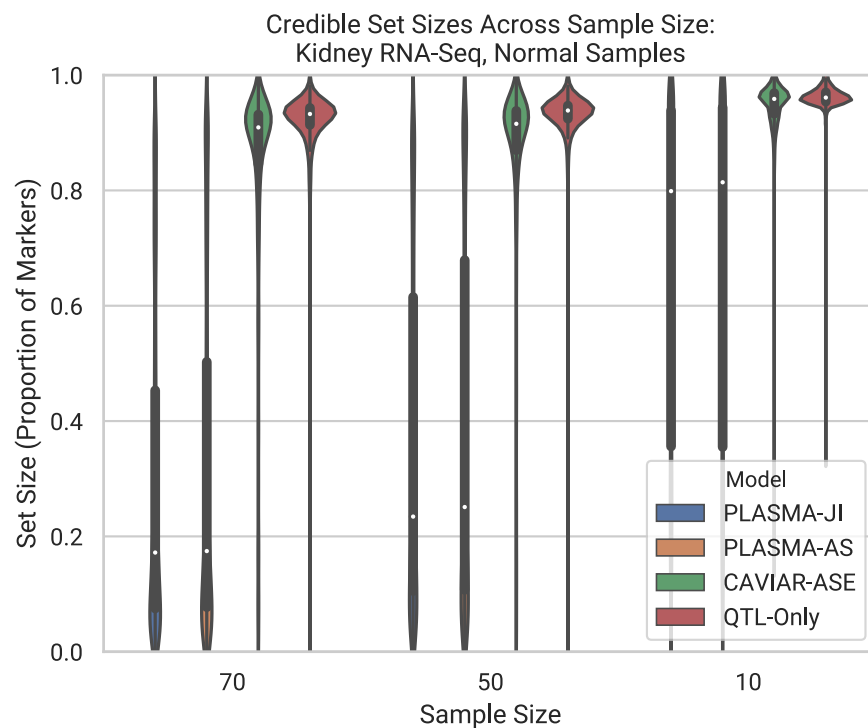
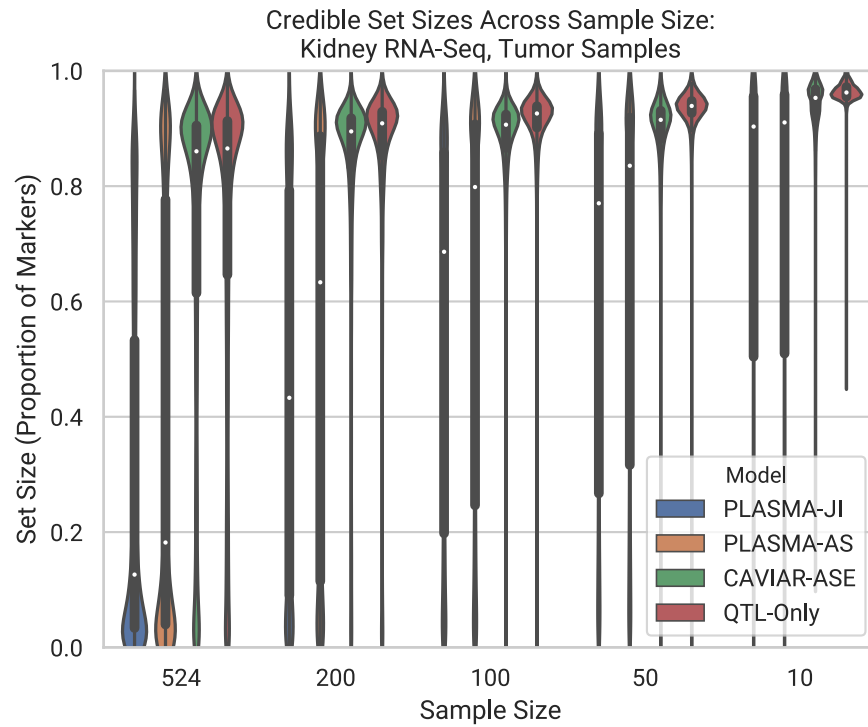


Figure 6: Performance of PLASMA as a function of sample size. (a) Credible set size distributions with varying down-sampled sample sizes, in tumor samples. During down-sampling, sample subsets were constructed by selecting randomly from total samples. (b) Credible set size distributions with varying down-sampled sample sizes, in normal samples.

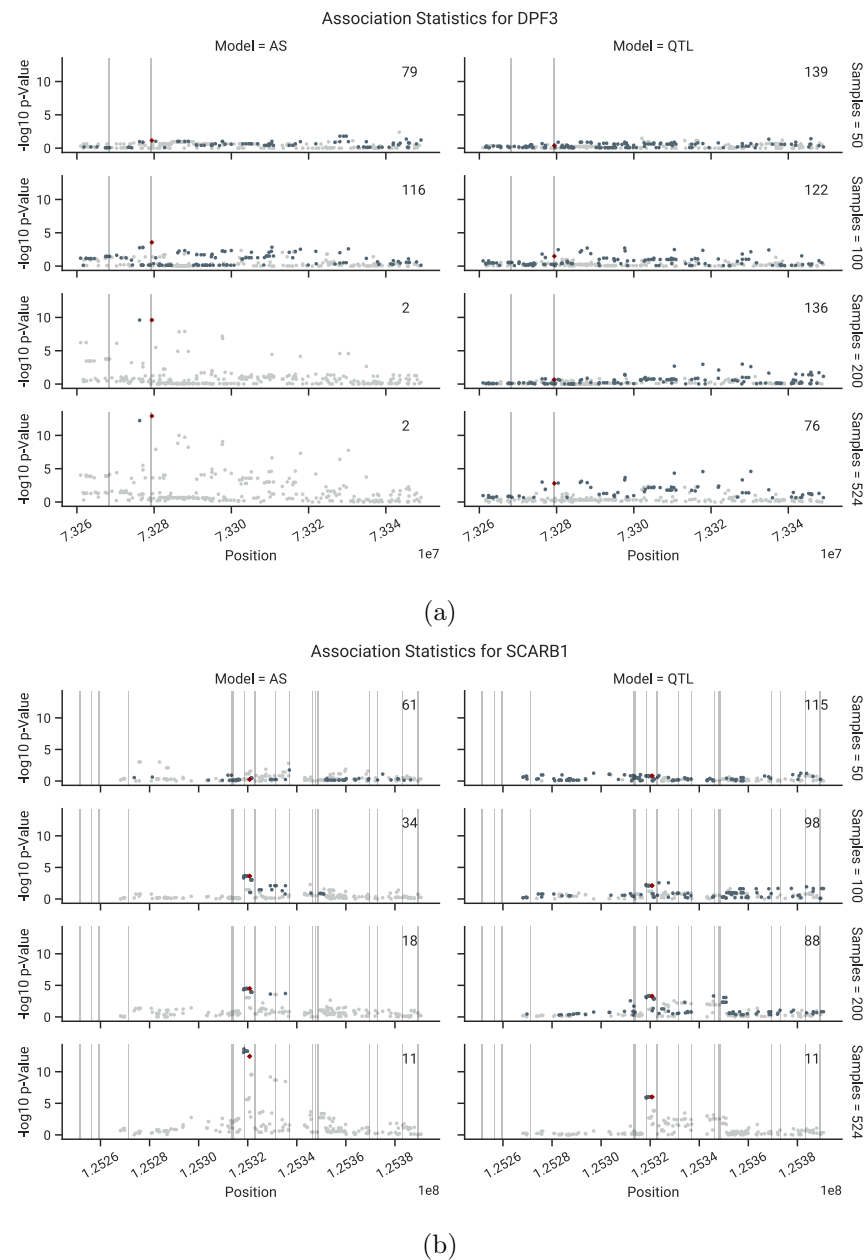


Figure 7: Comparison of AS and QTL associations in the experimentally-verified loci, as a function of sample size. Regions of 70 kb are shown around the causal marker. AS and QTL p-Values were calculated using from z_ϕ and z_β , respectively. Fine-mapping was conducted with the PLASMA-AS and QTL-Only 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.

2.6 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. Figures 8 shows distribution plots for tumor and normal data under a 1 causal variant assumption. Among the tumor loci, 15.2% of loci are fine-mapped within 50 variants with PLASMA-JI, while 1.9% of loci are fine-mapped within 50 variants with QTL-Only (Table S5a). Furthermore, PLASMA achieves a median credible set size of 226, compared to QTL-Only fine-mapping achieving a size of 322 (Table S6a). PLASMA also outperforms CAVIAR-ASE, with 1.9% of loci fine-mapped within 50 causal variants (no gain over QTL-Only), and a median credible set size of 310. Results from normal samples are similar, with 10.2%, 2.4%, and 2.6% of loci fine-mapped within 50 causal variants, for PLASMA, QTL-Only, and CAVIAR-ASE respectively (Table S5b). These methods achieve a median credible size of 232, 321, and 313 variants, respectively (Table S6b). Overall, these ChIP fine-mapping results are roughly in line with those from RNA-Seq fine-mapping.

2.7 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 [27]. Since chromatin accessibility is an indicator of transcription factor binding and regulation [28], 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 7a). We 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 QTL-Only analysis. We define the null distribution as the credible set markers being located independently of open chromatin and use Fisher's exact test to calculate enrichment as a function of minimum causal variant probability. Figures 9a and b, and Tables S7 and S8 show the p-values and odds ratios, respectively, (computed by Fishers exact test) as a function of posterior probability threshold from each fine-mapping method. In terms of both p-values and odds ratios, we see that the credible set markers produced by PLASMA, for the most part, display a significantly stronger enrichment with open chromatin compared to existing methods in terms of both p-values and odds ratios. For instance, at the $p = 0.1$ threshold for tumor samples, PLASMA's credible set markers achieve a p-value of 5.61×10^{-57} and an odds ratio of 2.30, respectively. In comparison, credible sets from QTL-Only fine-mapping at that threshold achieve a p-value of 1.50×10^{-5} and odds ratio of 1.58, respectively. 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. We believe that this is due to PLASMA 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

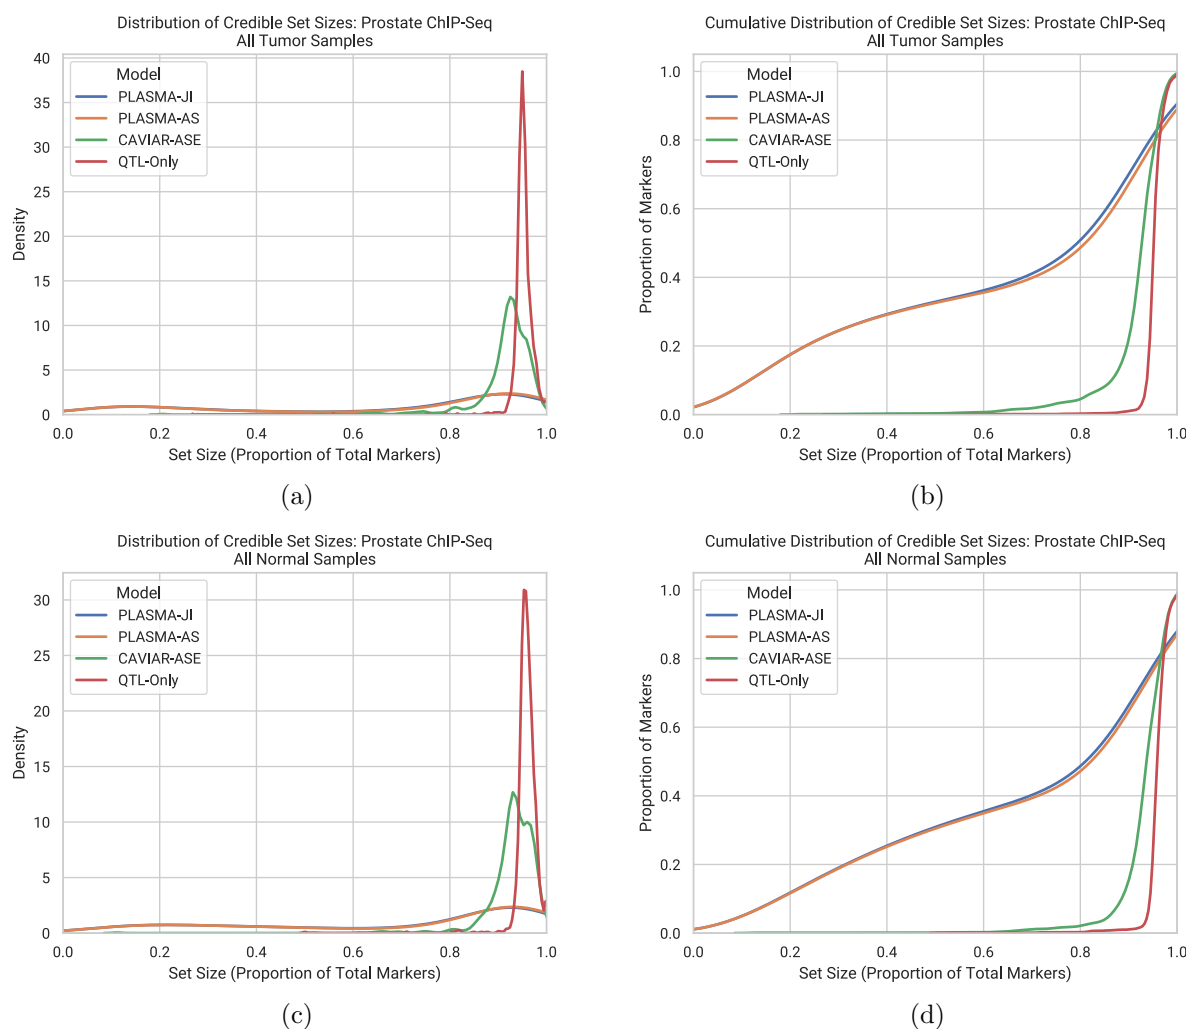


Figure 8: Comparison of 95% confidence credible set sizes for peaks in prostate tumor and normal cells, with an allelic imbalance false discovery rate of 0.05. Presence of H3k27ac histone marks was quantified as ChIP-Seq read counts mapped to phased genotypes. (a) Distribution of credible set sizes for tumor samples, under a 1-causal-variant assumption. (b) Cumulative distribution of credible set sizes for tumor samples, under a 1-causal-variant assumption. (c) Distribution for normal samples, under a 1-causal-variant assumption. (d) Cumulative distribution for normal samples, under a 1-causal-variant assumption.

we reasoned that they should therefore be enriched for true causal cQTLs [29, 30]. 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 9c and Table S9 show the p-values and odds ratios, respectively, across models as a function of posterior probability threshold (computed by Fishers exact test). Comparing the methods, we see that the credible set markers produced by PLASMA display a significantly stronger enrichment with looping anchors compared to the other methods in terms of both p-values and odds ratios. For instance, at the $p = 0.1$ threshold, PLASMA’s credible sets achieve a p-value of 9.13×10^{-7} and an odds ratio of 1.76, respectively. In contrast, credible set markers from QTL-Only fine-mapping at that threshold achieve a p-value of 0.37 and odds ratio of 0.34, respectively.

3 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 [31]. 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 [32], and direct allele-specific measurement of expression using single-cell RNA-Seq [33]. 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 identified as a QTL association, but would not exhibit AS association. This would be a model violation

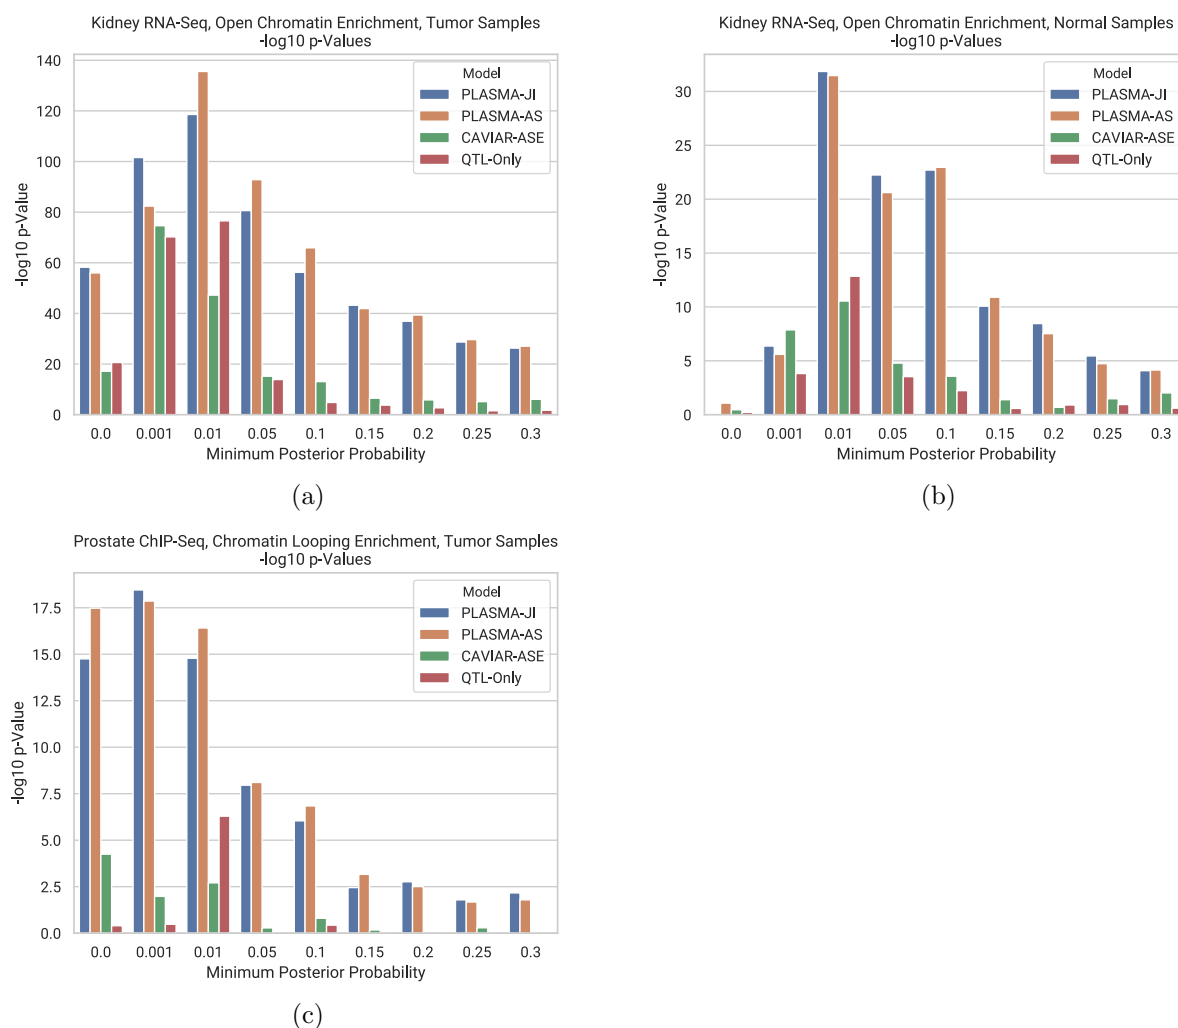


Figure 9: (a, b) Enrichment of fine-mapped credible sets with open chromatin for kidney tumor RNA-seq data in terms of negative-log p-values for tumor and normal samples, respectively. We took the 95% credible sets across loci, and looked at open chromatin regions that these markers intersect. We repeat this process on subsets of markers in the credible sets, thresholded by posterior probability. p-Values were calculated with Fisher's exact test. (c) Enrichment of fine-mapped credible sets with chromatin looping for prostate H3K27ac ChIP-Seq data.

for PLASMA and produce larger credible set sizes. Although PLASMA can consider correlations between causal AS and QTL effect sizes, this parameter is hard to estimate, and we find in real data that the model with correlation set to zero (PLASMA-JI) 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 [34]. 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, 35] 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) [36], 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.

4 Methods

4.1 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. We define β_i , the genetic effect of marker i on total gene expression as follows:

$$\mathbf{y} = \mathbf{x}_i \beta_i + \boldsymbol{\epsilon}_i \quad (3)$$

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

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

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

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

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

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

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

423 Experimentally-derived AS data, such as RNA-Seq data, yield reads that are mapped to a
 424 particular haplotype. For a given individual j , we define $c_{A,j}$ as the allele-specific read count from
 425 haplotype A . We model the allele-specific read count as drawn a beta-binomial distribution, given
 426 the total mapped read count c_j :

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

427 We use this beta binomial model to estimate the variance of the sampling error τ_j :

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

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

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

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

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

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

435 With this estimator, we define the AS association statistic for marker i as the AS effect size
 436 divided by the estimated variance of the effect size (full derivation in Supplemental Methods):

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

4.2 Inference of credible sets and posterior probabilities

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

$$\hat{\mathbf{z}} = \begin{bmatrix} \hat{\mathbf{z}}_\phi \\ \hat{\mathbf{z}}_\beta \end{bmatrix} \quad (12)$$

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. We define the combined correlation matrix \mathbf{R} as:

$$\mathbf{R} = \begin{bmatrix} \mathbf{R}_z & r_{\beta\phi}\mathbf{R}_z \\ r_{\beta\phi}\mathbf{R}_z & \mathbf{R}_z \end{bmatrix} \quad (13)$$

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

$$\hat{\mathbf{z}} \sim \mathcal{N}_{2m}(\mathbf{z}, \mathbf{R}) \quad (14)$$

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. We assume that the causal configuration is the same for the QTL and AS signals.

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. See Supplementary methods for a mathematical relationship between these two hyperparameters.) We show that these three hyperparameters are closely related to the heritability of gene expression (Supplemental Methods). We define $\Sigma_{\mathbf{c}}$, the covariance matrix of causal effect sizes, given a causal configuration:

$$\Sigma_{\mathbf{c}} = \begin{bmatrix} \text{diag}(\mathbf{c})\sigma_{c,\phi}^2 & \text{diag}(\mathbf{c})r_{c,\beta\phi}\sigma_{c,\phi}\sigma_{c,\beta} \\ \text{diag}(\mathbf{c})r_{c,\beta\phi}\sigma_{c,\phi}\sigma_{c,\beta} & \text{diag}(\mathbf{c})\sigma_{c,\beta}^2 \end{bmatrix} \quad (15)$$

We define our likelihood for a causal configuration as:

$$\mathcal{L}(\mathbf{c}; \hat{\mathbf{z}}) = \mathcal{N}_{2m}(\mathbf{0}, \mathbf{R} + \mathbf{R}\Sigma_{\mathbf{c}}\mathbf{R}) \quad (16)$$

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

$$P(\mathbf{c}) = \prod_{i=1}^m \gamma^{c_i} (1 - \gamma)^{1-c_i} \quad (17)$$

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

$$P(\mathbf{c}|\hat{\mathbf{z}}) = \frac{P(\hat{\mathbf{z}}|\mathbf{c})P(\mathbf{c})}{\sum_{\mathbf{c}^* \in \mathbb{C}} P(\hat{\mathbf{z}}|\mathbf{c}^*)P(\mathbf{c}^*)} \quad (18)$$

We define the ρ -level credible set \mathbb{K} as the smallest set of markers with a ρ_c probability of including 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. We calculate ρ_c as the sum of the probabilities of the configurations in $\mathbb{C}_{\mathbb{K}}$:

$$\rho_c = \sum_{\mathbf{c} \in \mathbb{C}_{\mathbb{K}}} \Pr(\mathbf{c}|\hat{\mathbf{z}}) \quad (19)$$

Additionally, we define a marker’s posterior inclusion probability (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.

To reduce the number of configurations to evaluate in the case of multiple causal variants, we use the heuristic that configurations with significant probabilities tend to be similar to each other. We use 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. Upon termination, we assume 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, we use a greedy approximation algorithm. At each step, before ρ_c is reached, the algorithm adds the marker that increases the confidence the most.

4.3 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 in Chromosome 22 is randomly chosen from the genotypes. Next, a random selection of samples are randomly selected from the section. The genotypes corresponding to the chosen samples yield two haplotype matrices, which we denote \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, we calculate \mathbf{q}_a and \mathbf{q}_b , the ideal un-scaled gene expression for each haplotype, by multiplying the relative effect sizes with each haplotype matrix.

With this haplotype-specific expression, we simulate read count data. 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, we calculate the allelic imbalance and the total read count (QTL) separately.

To calculate total read count data, we model total ideal un-scaled expression \mathbf{q}_t as $\mathbf{q}_a + \mathbf{q}_b$, the sum of the haplotype-specific un-scaled gene expression. We then add Gaussian-distributed noise so that the variance of \mathbf{q}_t is consistent with the total variance across samples as specified by the QTL heritability. Finally, we scale this final expression so that the total expression across samples is of unit variance. We do not explicitly generate total read counts, 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, we take into account heritability, mean read coverage, and the total variance of the AS phenotype. We model the ideal allelic imbalance phenotype as $\text{logit}\left(\frac{\mathbf{q}_a}{\mathbf{q}_b}\right)$ (calculated element-wise). We then add Gaussian-distributed noise 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.

4.4 Quality control of genotype data

For TCGA data, germline genotype calls were downloaded from the Genomic Data Commons. For PrCa ChIP samples, germline genotypes were called from blood as described in Ref. [37]. Genotypes were imputed to the Haplotype Reference Consortium [38] using the Michigan Imputation Server [39] and restricted to variants with INFO greater than 0.9 and MAF greater than 0.01. Variants were further restricted to QC-passing SNPs from Ref. [38] which represent common, well-mapped variants from the 1000 Genomes project.

4.5 Quality control of RNA-seq data

Raw RNA-seq BAM files were downloaded from the Genomic Data Commons. Initial RNA-seq mapping and alignment was performed following TCGA parameters for the STAR aligner [40]. Mapping bias was accounted for by re-mapping using the WASP pipeline [19] and the STAR aligner with the same parameters. Reads were randomly de-duplicated as recommended by the WASP pipeline.

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

4.6 Quality control of ChIP-seq data

ChIP-seq experiments were performed as described in Ref. [37]. Reads were aligned using bwa and default parameters [41], and peaks were called using MACS2 and default parameters (with DNA-seq input provided as control) [42]. Peaks were then unified across all samples. Mapping bias was accounted for by re-mapping using the WASP pipeline and the bwa aligner with the same parameters. Reads were randomly de-duplicated as recommended by the WASP pipeline. Beta binomial overdispersion parameters were estimated globally for each sample as somatic copy number was expected to be minimal.

4.7 Allele-specific quantification

The StratAS algorithm was 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 were 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 was then tested for an allele-specific association with the feature and features that were significant at a genome-wide false discovery rate (FDR) of 5% were retained for fine-mapping.

4.8 Functional enrichment analysis

For QTLs fine-mapped from RNA-seq we selected 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 selected 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 was then estimated by computing the proportion of markers in credible sets that intersect with the functional feature. Controls were calculated as the intersection between all tested

markers and the functional feature. Odds ratios and p-values were computed with Fisher’s exact test.

Acknowledgments

We thank F. Hormozdiari for guidance on statistical fine-mapping and C. Kalita for guidance on model validation. We also thank B. Pasanuic and C. Giambartolomei 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.
- [24] Zou J, Hormozdiari F, Jew B, Ernst J, *et al.* Leveraging allele-specific expression to refine fine-mapping for eQTL studies. *bioRxiv*, 2018.
- [25] 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.
- [26] 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.
- [27] Kundaje A, Meuleman W, Ernst J, Bilenky M, *et al.* Integrative analysis of 111 reference human epigenomes. *Nature*, 518(7539):317–330, 2015.
- [28] 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.
- [29] 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.

- [30] 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.
- [31] 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.
- [32] Herzel L, Straube K, Neugebauer KM. Long-read sequencing of nascent RNA reveals coupling among RNA processing events. *Genome Res*, 28(7):1008–1019, 2018.
- [33] van der Wijst MGP, Brugge H, de Vries DH, Deelen P, *et al.* Single-cell RNA sequencing identifies celltype-specific cis-eQTLs and co-expression QTLs. *Nat Genet*, 50(4):493–497, 2018.
- [34] 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.
- [35] 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.
- [36] 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.
- [37] Pomerantz M, Qiu X, Zhy Y, Takeda D, *et al.* Charting the prostate epigenomic landscape from transformation to progression. (In preparation).
- [38] 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.
- [39] Das S, Forer L, Schonherr S, Sidore C, *et al.* Next-generation genotype imputation service and methods. *Nat Genet*, 48(10):1284–1287, 2016.
- [40] Dobin A, Davis CA, Schlesinger F, Drenkow J, *et al.* STAR: ultrafast universal RNA-seq aligner. *Bioinformatics*, 29(1):15–21, 2013.
- [41] Li H, Durbin R. Fast and accurate short read alignment with Burrows-Wheeler transform. *Bioinformatics*, 25(14):1754–1760, 2009.
- [42] Zhang Y, Liu T, Meyer CA, Eeckhoutte J, *et al.* Model-based analysis of ChIP-Seq (MACS). *Genome Biol*, 9(9):R137, 2008.

646

54




Figure S1: 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.

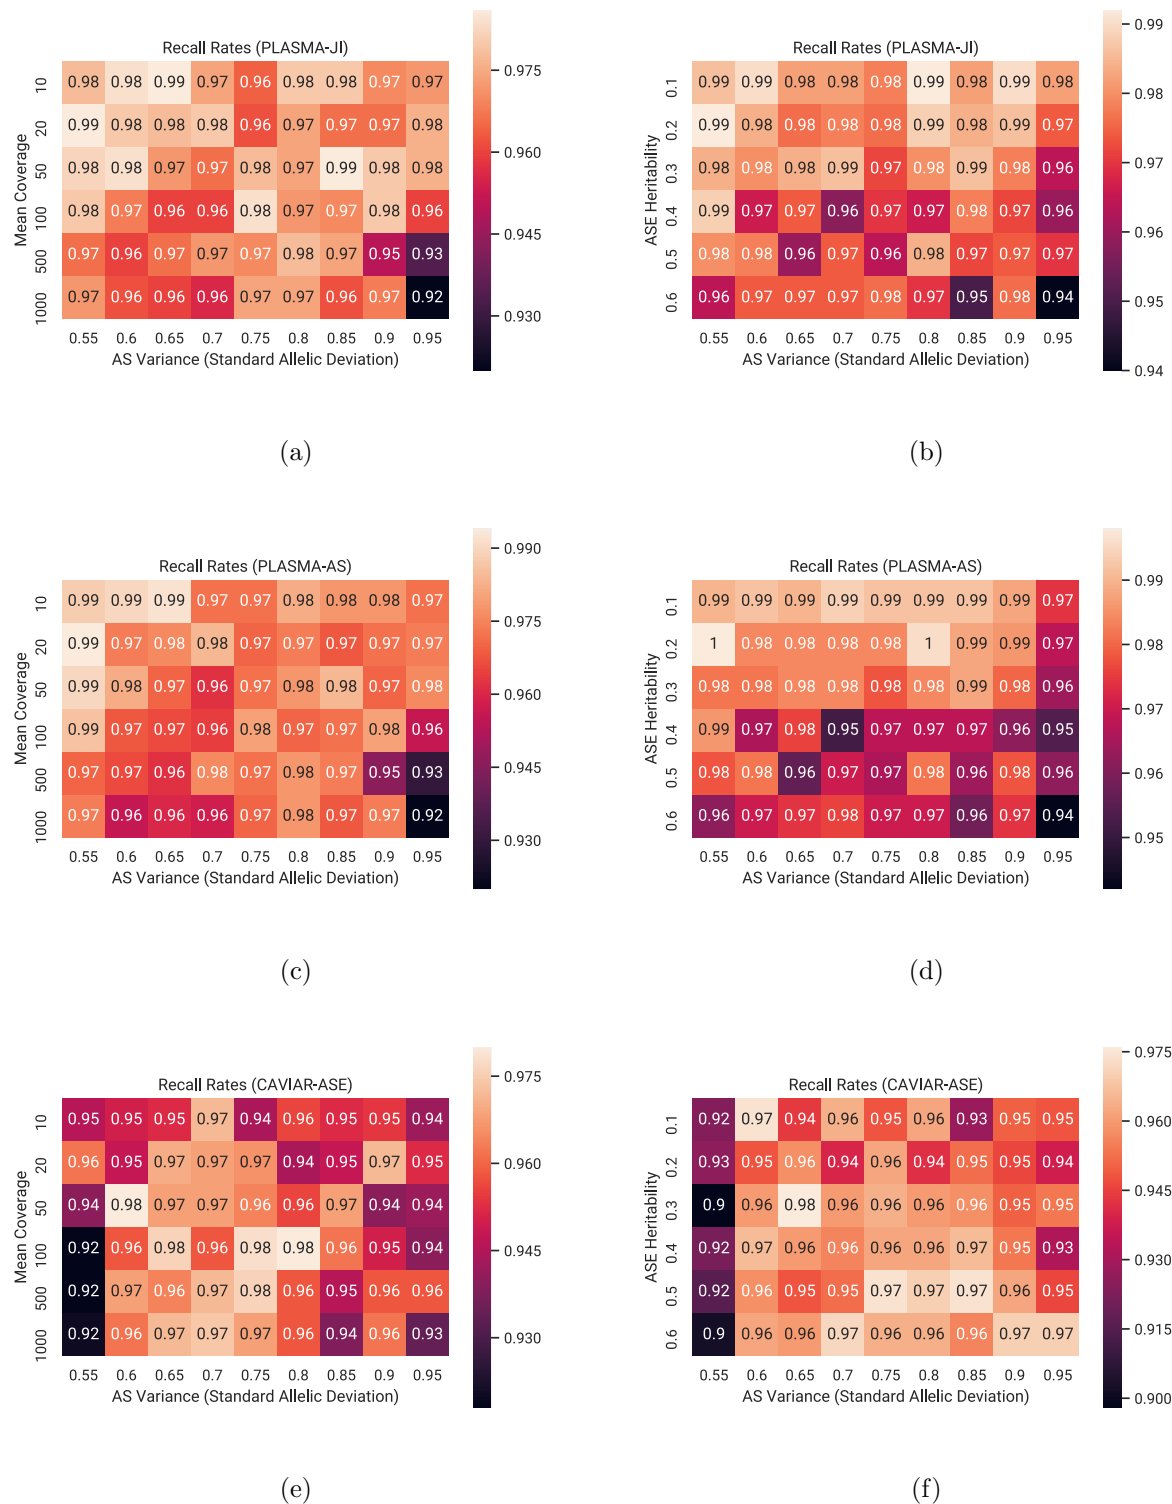


Figure S2: Recall rates for 95% credible sets under various conditions of mean read coverage, heritability, and standard allelic deviation (intensity of AS variance). Recall rate is defined as the expected proportion of causal markers included in the credible set. Each square is the mean statistic over 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.

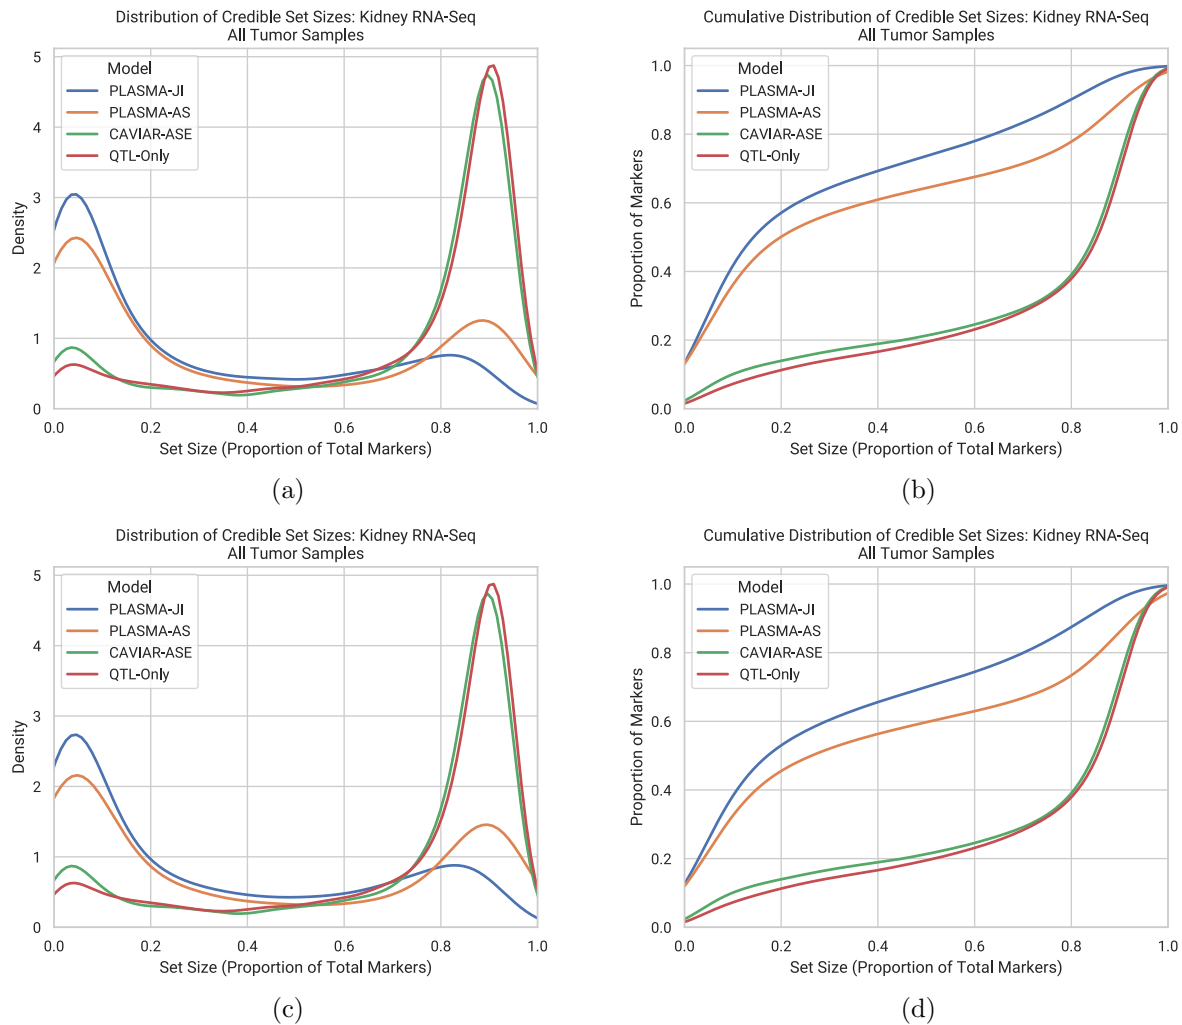


Figure S3: Effect of the AS heritability hyperparameter on 95% credible set sizes with tumor RNA-Seq data. (a) Set size distribution with the heritability hyperparameter at 50% (b) Cumulative set size distribution with the heritability hyperparameter at 50% (c) Set size distribution with the heritability hyperparameter at 5% (d) Cumulative set size distribution with the heritability hyperparameter at 5%

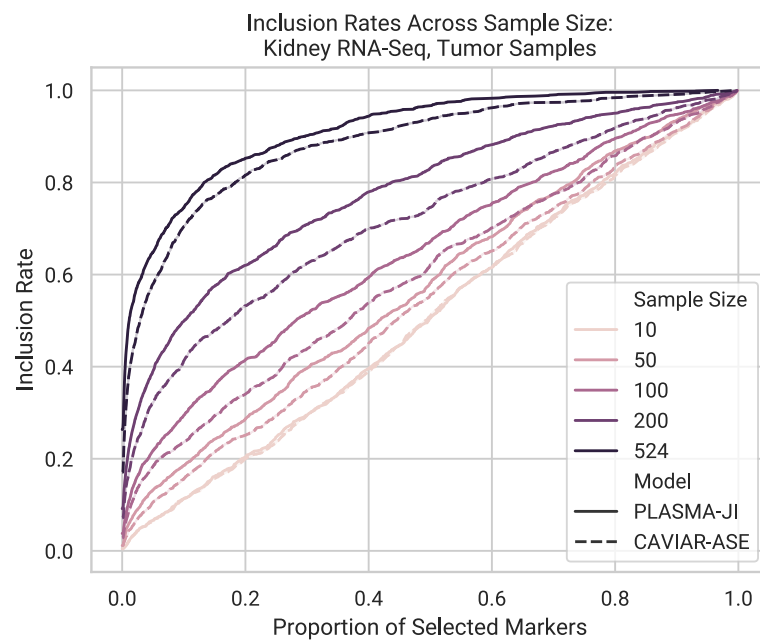


Figure S4: Inclusion curve for fine-mapped kidney tumor RNA-Seq data for PLASMA and CAVIAR-ASE across sample sizes. Inclusion was evaluated against a gold-standard set of markers, defined as all markers with at least 0.1 posterior probability under QTL-Only fine-mapping with all samples.

Table S1: Proportions of kidney RNA-Seq loci fine-mapped within various credible set size thresholds. (a) Threshold proportions for tumor samples (b) Threshold proportions for normal samples

(a) Tumor Samples				
Threshold	PLASMA-JI	QTL-Only	PLASMA-AS	CAVIAR-ASE
5	0.1356	0.0185	0.1202	0.0284
10	0.2237	0.0314	0.1973	0.0473
20	0.3433	0.0499	0.3037	0.0740
50	0.5216	0.0914	0.4665	0.1312
100	0.6507	0.1481	0.5800	0.1785

(b) Normal Samples				
Threshold	PLASMA-JI	QTL-Only	PLASMA-AS	CAVIAR-ASE
5	0.0277	0.0010	0.0272	0.0024
10	0.0704	0.0019	0.0709	0.0053
20	0.1661	0.0044	0.1598	0.0107
50	0.4031	0.0107	0.3929	0.0287
100	0.6333	0.0223	0.6207	0.0602

Table S2: Statistics for credible sets sizes of fine-mapped kidney RNA-Seq loci. (a) Statistics for tumor samples (b) Statistics for normal samples

(a) Tumor Samples				
	PLASMA-JI	QTL-Only	PLASMA-AS	CAVIAR-ASE
Mean	110.0	288.6	144.4	285.2
Variance	18155.	29337.	27310.	33197.
25th Percentile	12.0	176.0	15.0	163.0
Median	45.0	289.0	62.0	292.0
75th Percentile	176.0	392.0	260.0	394.0

(b) Normal Samples				
	PLASMA-JI	QTL-Only	PLASMA-AS	CAVIAR-ASE
Mean	119.1	384.5	126.9	285.2
Variance	16454.	29666.	19223.	352.7
25th Percentile	30.0	284.0	31.0	30546.
Median	67.0	374.0	68.0	249.0
75th Percentile	166.5	468.0	176.0	348.0

Table S3: Statistics for credible sets sizes of fine-mapped kidney RNA-Seq loci under different AS heritability hyperparameters. Credible sets were generated with PLASMA-JI.

	0.5	0.05
Mean	110.0	121.9
Variance	18155.	20064.
25th Percentile	12.0	14.0
Median	45.0	56.0
75th Percentile	176.0	199.0

Table S4: Proportions of kidney RNA-Seq loci fine-mapped to a 95% credible set size of within 10 markers with varying sample sizes. (a) Proportions of tumor sample loci across sample sizes at a credible set size threshold of 10 (b) Proportions of normal sample loci across sample sizes at a credible set size threshold of 10 (C) Median credible set sizes for tumor sample loci across sample sizes (b) Median credible set sizes for normal sample loci across sample sizes threshold

(a) Tumor Samples, Credible Set Size Threshold of 10

Samples	PLASMA-JI	QTL-Only	PLASMA-AS	CAVIAR-ASE
524	0.2237	0.0314	0.1973	0.0473
200	0.1081	0.0150	0.0950	0.0108
100	0.0589	0.0105	0.0501	0.0068
50	0.0335	0.0030	0.0300	0.0019
10	0.0016	0.0000	0.0014	0.0000

(b) Normal Samples, Credible Set Size Threshold of 10

Samples	PLASMA-JI	QTL-Only	PLASMA-AS	CAVIAR-ASE
70	0.0704	0.0019	0.0709	0.0053
50	0.0452	0.0024	0.0423	0.0034
10	0.0016	0.00052	0.0016	0.0005

(c) Tumor Samples, Median Credible Set Size

Samples	PLASMA-JI	QTL-Only	PLASMA-AS	CAVIAR-ASE
524	45.0	289.0	62.0	292.0
200	143.0	333.0	188.0	328.0
100	212.0	351.0	244.0	341.0
50	239.0	362.0	257.0	347.0
10	287.0	381.0	290.0	371.0

(d) Normal Samples, Median Credible Set Size

Samples	PLASMA-JI	QTL-Only	PLASMA-AS	CAVIAR-ASE
70	67.0	374.0	68.0	348.0
50	89.0	377.0	92.0	357.0
10	259.0	398.0	259.5	390.0

Table S5: Proportions of prostate ChIP-Seq loci fine-mapped within various credible set size thresholds. (a) Threshold proportions for tumor samples (b) Threshold proportions for normal samples

(a) Tumor Samples				
Threshold	PLASMA-JI	QTL-Only	PLASMA-AS	CAVIAR-ASE
5	0.0036	0.0000	0.0015	0.0007
10	0.0138	0.0029	0.0138	0.0029
20	0.0386	0.0080	0.0378	0.0080
50	0.1521	0.0189	0.1507	0.0189
100	0.2766	0.0255	0.2751	0.0284

(b) Normal Samples				
Threshold	PLASMA-JI	QTL-Only	PLASMA-AS	CAVIAR-ASE
5	0.0022	0.0022	0.0022	0.0022
10	0.0088	0.0055	0.0088	0.0055
20	0.0297	0.0143	0.0308	0.0143
50	0.1023	0.0242	0.1001	0.0264
100	0.2332	0.0330	0.2230	0.0341

Table S6: Statistics for credible sets sizes of fine-mapped prostate ChIP-Seq loci. (a) Statistics for tumor samples (b) Statistics for normal samples

(a) Tumor Samples				
	PLASMA-JI	QTL-Only	PLASMA-AS	CAVIAR-ASE
Mean	207.3	311.7	209.9	299.8
Variance	14982.	7393.	15448.	7337.
25th Percentile	85.25	267.0	86.0	256.0
Median	226.0	322.0	228.5	310.0
75th Percentile	307.75	366.0	312.0	355.0

(b) Normal Samples				
	PLASMA-JI	QTL-Only	PLASMA-AS	CAVIAR-ASE
Mean	215.9	311.5	218.2	302.2
Variance	13809.	8266.	13977.	8125.
25th Percentile	108.0	265.0	109.0	257.0
Median	232.0	321.0	237.0	313.0
75th Percentile	312.0	369.0	314.0	358.0

Table S7: Enrichment of fine-mapped credible sets with chromatin looping for kidney RNA-Seq tumor data. We took the 95% credible sets across peaks, and looked at regions of open chromatin that these markers intersect. We repeat this process on subsets of markers in the credible sets, thresholded by posterior probability. P-values, odds ratios and confidence intervals were calculated with Fisher’s exact test. (a-d) Enrichment statistics with PLASMA-JI, QTL-Only, PLASMA-AS, and CAVIAR-ASE fine-mapped credible sets, respectively.

(a) PLASMA-JI

Min Posterior Prob.	Odds Ratio	95% C.I. Lower	95% C.I. Upper	$-\log_{10}$ p-Value
0.0	1.174	1.152	1.197	58.24
0.001	1.256	1.231	1.282	101.6
0.01	1.545	1.492	1.599	118.6
0.05	1.998	1.871	2.132	80.61
0.1	2.300	2.094	2.521	56.25
0.15	2.492	2.217	2.793	43.23
0.2	2.694	2.350	3.077	36.92
0.25	2.754	2.347	3.215	28.70
0.3	2.906	2.437	3.443	26.29

(b) QTL-Only

Min Posterior Prob.	Odds Ratio	95% C.I. Lower	95% C.I. Upper	$-\log_{10}$ p-Value
0.0	1.119	1.093	1.146	20.61
0.001	1.198	1.174	1.222	70.22
0.01	1.473	1.417	1.532	76.56
0.05	1.603	1.429	1.794	13.88
0.1	1.582	1.289	1.924	4.822
0.15	1.707	1.294	2.214	3.768
0.2	1.724	1.215	2.380	2.700
0.25	1.604	1.028	2.396	1.540
0.3	1.792	1.096	2.780	1.759

(c) PLASMA-AS

Min Posterior Prob.	Odds Ratio	95% C.I. Lower	95% C.I. Upper	$-\log_{10}$ p-Value
0.0	1.166	1.144	1.188	55.99
0.001	1.211	1.188	1.235	82.38
0.01	1.604	1.549	1.662	135.54
0.05	2.157	2.017	2.305	92.83
0.1	2.537	2.306	2.786	65.91
0.15	2.574	2.276	2.901	41.90
0.2	2.910	2.525	3.340	39.36
0.25	2.948	2.495	3.463	29.65
0.3	3.135	2.607	3.744	27.04

(d) CAVIAR-ASE

Min Posterior Prob.	Odds Ratio	95% C.I. Lower	95% C.I. Upper	$-\log_{10}$ p-Value
0.0	1.108	1.082	1.134	17.16
0.001	1.198	1.175	1.221	74.64
0.01	1.388	1.330	1.448	47.22
0.05	1.610	1.441	1.794	15.19
0.1	1.971	1.665	2.318	13.06
0.15	1.864	1.477	2.324	6.511
0.2	2.060	1.547	2.695	5.843
0.25	2.197	1.574	2.995	5.151
0.3	2.551	1.783	3.554	6.081

Table S8: Enrichment of fine-mapped credible sets with chromatin looping for kidney RNA-Seq normal data. We took the 95% credible sets across peaks, and looked at regions of open chromatin that these markers intersect. We repeat this process on subsets of markers in the credible sets, thresholded by posterior probability. P-values, odds ratios and confidence intervals were calculated with Fisher’s exact test. (a-d) Enrichment statistics with PLASMA-JI, QTL-Only, PLASMA-AS, and CAVIAR-ASE fine-mapped credible sets, respectively.

(a) PLASMA-JI

Min Posterior Prob.	Odds Ratio	95% C.I. Lower	95% C.I. Upper	$-\log_{10}$ p-Value
0.0	1.005	.9731	1.037	.1098
0.001	1.095	1.057	1.134	6.373
0.01	1.445	1.363	1.531	31.85
0.05	1.776	1.594	1.974	22.24
0.1	2.265	1.952	2.617	22.70
0.15	2.047	1.668	2.489	10.07
0.2	2.154	1.689	2.713	8.455
0.25	2.044	1.521	2.695	5.453
0.3	1.942	1.398	2.637	4.083

(b) QTL-Only

Min Posterior Prob.	Odds Ratio	95% C.I. Lower	95% C.I. Upper	$-\log_{10}$ p-Value
0.0	1.010	.9729	1.048	.2094
0.001	1.065	1.031	1.1013	3.819
0.01	1.450	1.318	1.592	12.85
0.05	1.955	1.362	2.729	3.532
0.1	2.363	1.266	4.071	2.231
0.15	1.573	0.499	3.801	.5835
0.2	2.158	0.568	5.840	.8992
0.25	2.665	0.524	8.476	.9411
0.3	2.158	0.249	8.557	.6024

(c) PLASMA-AS

Min Posterior Prob.	Odds Ratio	95% C.I. Lower	95% C.I. Upper	$-\log_{10}$ p-Value
0.0	1.027	.9963	1.060	1.072
0.001	1.086	1.049	1.123	5.598
0.01	1.442	1.360	1.528	31.47
0.05	1.746	1.565	1.943	20.61
0.1	2.291	1.972	2.649	22.954
0.15	2.116	1.722	2.577	10.90
0.2	2.087	1.625	2.644	7.520
0.25	1.965	1.449	2.610	4.730
0.3	1.965	1.414	2.668	4.144

(d) CAVIAR-ASE

Min Posterior Prob.	Odds Ratio	95% C.I. Lower	95% C.I. Upper	$-\log_{10}$ p-Value
0.0	1.016	.9831	1.05	.4543
0.001	1.084	1.054	1.115	7.869
0.01	1.332	1.226	1.444	10.55
0.05	1.748	1.361	2.215	4.787
0.1	2.119	1.412	3.070	3.570
0.15	1.855	.9981	3.178	1.390
0.2	1.692	.6663	3.590	.6872
0.25	2.403	.9386	5.162	1.478
0.3	3.156	1.22	6.871	2.018

Table S9: Enrichment of fine-mapped credible sets with chromatin looping for prostate H3K27ac ChIP-Seq tumor data. We took the 95% credible sets across peaks, and looked at regions of chromatin looping that these markers intersect. We repeat this process on subsets of markers in the credible sets, thresholded by posterior probability. P-values, odds ratios and confidence intervals were calculated with Fisher’s exact test. (a-d) Enrichment statistics with PLASMA-JI, QTL-Only, PLASMA-AS, and CAVIAR-ASE fine-mapped credible sets, respectively

(a) PLASMA-JI

Min Posterior Prob.	Odds Ratio	95% C.I. Lower	95% C.I. Upper	$-\log_{10}$ p-Value
0.0	1.162	1.119	1.206	14.75
0.001	1.182	1.140	1.227	18.46
0.01	1.353	1.258	1.454	14.79
0.05	1.592	1.364	1.850	7.958
0.1	1.762	1.413	2.177	6.040
0.15	1.568	1.158	2.082	2.450
0.2	1.759	1.234	2.444	2.7616
0.25	1.666	1.081	2.471	1.790
0.3	1.834	1.177	2.745	2.164

(b) QTL-Only

Min Posterior Prob.	Odds Ratio	95% C.I. Lower	95% C.I. Upper	$-\log_{10}$ p-Value
0.0	1.020	0.975	1.067	.3985
0.001	1.022	0.978	1.070	.4764
0.01	1.392	1.226	1.575	6.291
0.05	.9287	.4183	1.800	.0000
0.1	.3370	.0084	1.951	.4306
0.15	.0000	.0000	5.245	.0000
0.2	.0000	.0000	14.03	.0000
0.25	.0000	.0000	14.03	.0000
0.3	.0000	.0000	30.64	.0000

(c) PLASMA-AS

Min Posterior Prob.	Odds Ratio	95% C.I. Lower	95% C.I. Upper	$-\log_{10}$ p-Value
0.0	1.179	1.135	1.223	17.48
0.001	1.179	1.136	1.223	17.86
0.01	1.382	1.28	1.49	16.41
0.05	1.598	1.369	1.857	8.105
0.1	1.823	1.466	2.246	6.842
0.15	1.668	1.239	2.205	3.161
0.2	1.715	1.191	2.403	2.505
0.25	1.645	1.058	2.456	1.675
0.3	1.721	1.096	2.593	1.793

(d) CAVIAR-ASE

Min Posterior Prob.	Odds Ratio	95% C.I. Lower	95% C.I. Upper	$-\log_{10}$ p-Value
0.0	1.094	1.047	1.143	4.248
0.001	1.056	1.013	1.102	1.978
0.01	1.183	1.063	1.314	2.705
0.05	.7941	.4182	1.378	.2815
0.1	1.686	.7075	3.458	.7999
0.15	1.264	.1468	4.956	.1712
0.2	.9194	.0223	5.689	.0000
0.25	1.445	.0342	9.499	.2881
0.3	.0000	.0000	9.026	.0000

5 Supplemental Methods

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:

$$\mathbf{y} = \mathbf{x}_i \beta_i + \boldsymbol{\epsilon}_i \quad (20)$$

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:

$$\hat{\beta}_i = \left(\mathbf{x}_i^\top \mathbf{x}_i \right)^{-1} \mathbf{x}_i^\top \mathbf{y} \quad (21)$$

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:

$$\begin{aligned} \text{Var}(\hat{\beta}_i) &= \left(\mathbf{x}_i^\top \mathbf{x}_i \right)^{-2} \text{Var}(\mathbf{x}_i^\top \mathbf{y}) \\ &= \left(\mathbf{x}_i^\top \mathbf{x}_i \right)^{-2} \left(\mathbf{x}_i^\top \mathbf{x}_i \right) \text{Var}(y) \\ &= \left(\mathbf{x}_i^\top \mathbf{x}_i \right)^{-1} \text{Var}(y) \\ &= \left(\mathbf{x}_i^\top \mathbf{x}_i \right)^{-1} \sigma_{y,i}^2 \end{aligned} \quad (22)$$

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

$$\hat{\sigma}_{\beta,i}^2 = \frac{\boldsymbol{\epsilon}_i^\top \boldsymbol{\epsilon}_i}{n - 1} \quad (23)$$

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

$$\hat{z}_{\beta,i} = \frac{\hat{\beta}_i}{\sqrt{\left(\mathbf{x}_i^\top \mathbf{x}_i \right)^{-1} \hat{\sigma}_{y,i}^2}} \quad (24)$$

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

$$\hat{z}_{\beta,i} \sim \mathcal{N}(z_{\beta,i}, 1) \quad (25)$$

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

$$\hat{z}_{\beta,i} = \frac{\hat{\beta}_i \sqrt{n}}{\sqrt{\hat{\sigma}_{y,i}^2}} \quad (26)$$

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:

$$\frac{e_1}{e_0} = \delta_i \frac{e'_1}{e'_0} \quad (27)$$

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

$$\frac{e_A}{e_B} = (\delta_i)^{v_i} \frac{e'_A}{e'_B} \quad (28)$$

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 28 as a linear model. Let w be the log expression ratio between haplotypes A and B :

$$w = \log \left(\frac{e_A}{e_B} \right) \quad (29)$$

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

$$\phi_i = \log(\delta_i) \quad (30)$$

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

$$\zeta_i = \log \left(\frac{e'_A}{e'_B} \right) \quad (31)$$

With these parameters we rewrite Equation 28 as:

$$w = v_i \phi_i + \zeta_i \quad (32)$$

Given n individuals, this expression becomes:

$$\mathbf{w} = \mathbf{v}_i \phi_i + \boldsymbol{\zeta}_i \quad (33)$$

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:

$$\hat{w} = \log \left(\frac{c_A}{c_B} \right) \quad (34)$$

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 :

$$c_{A,j} \sim \text{BB}(\alpha_j, \beta_j, c_j) \quad (35)$$

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

$$\pi_j = \frac{\mathbb{E}[c_{A,j}]}{c_j} = \frac{\alpha_j}{\alpha_j + \beta_j} \quad (36)$$

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

$$\rho_e = \frac{1}{\alpha_j + \beta_j + 1} \quad (37)$$

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

$$\mathbb{E}[c_{A,j}] = c_j \pi_j \quad (38)$$

$$\text{Var}(c_{A,j}) = c_j \pi_j (1 - \pi_j) (1 + \rho_e (c_j - 1)) \quad (39)$$

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:

$$\mathbb{E} \left[\frac{c_{A,j}}{c_j} \right] = \pi_j \quad (40)$$

$$\text{Var} \left(\frac{c_{A,j}}{c_j} \right) = \frac{1}{c_j} \pi_j (1 - \pi_j) (1 + \rho_e (c_j - 1)) \quad (41)$$

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

$$w_j^* = \text{logit}(\pi_j) = \log \frac{\pi_j}{1 - \pi_j} \quad (42)$$

$$\frac{dw_j^*}{d\pi_j} = \frac{1}{\pi_j(1 - \pi_j)} \quad (43)$$

$$\frac{d^2 w_j^*}{d\pi_j^2} = \frac{2\pi_j - 1}{\pi_j^2(1 - \pi_j)^2} \quad (44)$$

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

$$\begin{aligned}\mathbb{E}[\hat{w}_j] &= \mathbb{E}\left[\text{logit}\left(\frac{c_{A,j}}{c_j}\right)\right] \\ &\approx \text{logit}\left(\mathbb{E}\left[\frac{c_{A,j}}{c_j}\right]\right) + \frac{1}{2}\text{Var}\left(\frac{c_{A,j}}{c_j}\right)\frac{d^2}{d\pi_j^2}\text{logit}(\pi_j) \\ &\approx \text{logit}(\pi_j) + \frac{1}{2}\left(\frac{1}{c_j}\pi_j(1-\pi_j)(1+\rho_e(c_j-1))\right)\left(\frac{2\pi_j-1}{\pi_j^2(1-\pi_j)^2}\right) \\ &\approx \text{logit}(\pi_j) + \frac{2\pi_j-1}{2c_j\pi_j(1-\pi_j)}(1+\rho_e(c_j-1)) \\ &\approx w_j^* + \frac{1}{c_j}\sinh(w_j^*)(1+\rho_e(c_j-1))\end{aligned}\quad (45)$$

709

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

710 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 45
711 implies that \hat{w} is a biased estimator of w^* , especially at low read counts and/or high overdispersion.
712 To get an estimator of w^* with reduced bias, we take the approximation that $\sinh(w^*) \approx w^*$ around
713 zero:

$$\hat{w}_j^* = \frac{\hat{w}_j}{1 + \frac{1}{c_j}(1 + \rho_e(c_j - 1))}\quad (47)$$

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

$$\hat{\sigma}_{c,j}^2 = \frac{2}{c_j}(1 + \cosh(\hat{w}_j^*))(1 + \rho_e(c_j - 1))\quad (48)$$

715 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
716 $\text{Var}(\tau_j) = \sigma_{c,j}^2$. Thus, across individuals:

$$\hat{\mathbf{w}} = \mathbf{v}_i\phi_i + \boldsymbol{\zeta}_i + \boldsymbol{\tau}\quad (49)$$

717 5.2.2 Calculation of AS summary statistics

718 Due to heteroscedasticity among individuals, we estimate the AS effect size ϕ_i in a weighted manner,
719 giving larger weights to individuals with lower expected sampling error. Given individual j , we
720 define the weight for j as the inverse of the estimated read count variance:

$$\omega_j = \frac{1}{\hat{\sigma}_{c,j}^2}\quad (50)$$

721 We define our weight matrix Ω as a diagonal matrix with $\Omega_{j,j} = \omega_j$.

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

$$\hat{\phi}_i = \left(\mathbf{v}_i^\top \Omega \mathbf{v}_i \right)^{-1} \mathbf{v}_i^\top \Omega \hat{\mathbf{w}} \quad (51)$$

723 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
724 $\sigma_{w,i}^2 + \sigma_{c,j}^2$. Thus, under the null:

$$\begin{aligned} \text{Var}(\hat{\phi}_i) &= \left(\mathbf{v}_i^\top \Omega \mathbf{v}_i \right)^{-2} \text{Var}(\mathbf{v}_i^\top \Omega \hat{\mathbf{w}}) \\ &= \left(\mathbf{v}_i^\top \Omega \mathbf{v}_i \right)^{-2} \left(\text{Var}(\mathbf{v}_i^\top \Omega \mathbf{w}) + \text{Var}(\mathbf{v}_i^\top \Omega \boldsymbol{\tau}) \right) \\ &= \left(\mathbf{v}_i^\top \Omega \mathbf{v}_i \right)^{-2} \left(\text{Var}(\mathbf{v}_i^\top \Omega \mathbf{w}) + \text{Var}(\mathbf{v}_i^\top \Omega^{1/2} (\Omega^{1/2} \boldsymbol{\tau})) \right) \\ &= \left(\mathbf{v}_i^\top \Omega \mathbf{v}_i \right)^{-2} \left((\mathbf{v}_i^\top \Omega^2 \mathbf{v}_i) \text{Var}(w) + \mathbf{v}_i^\top \Omega \mathbf{v}_i \right) \\ &= \left(\mathbf{v}_i^\top \Omega \mathbf{v}_i \right)^{-2} \left((\mathbf{v}_i^\top \Omega^2 \mathbf{v}_i) \sigma_{w,i}^2 + \mathbf{v}_i^\top \Omega \mathbf{v}_i \right) \end{aligned} \quad (52)$$

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

$$\hat{\sigma}_{w,i}^2 = \frac{\sum_{j=1}^n \left(\omega_j (\zeta_{i,j} + \tau_j)^2 - 1 \right)}{\sum_{j=1}^n \omega_j} \quad (53)$$

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

$$\begin{aligned} \mathbb{E}[\hat{\sigma}_{w,i}^2] &= \frac{\sum_{j=1}^n \left(\omega_j \mathbb{E}[(\zeta_{i,j} + \tau_j)^2] - 1 \right)}{\sum_{j=1}^n \omega_j} \\ &= \frac{\sum_{j=1}^n (\omega_j \text{Var}(\zeta_{i,j} + \tau_j) - 1)}{\sum_{j=1}^n \omega_j} \\ &= \frac{\sum_{j=1}^n (\omega_j \text{Var}(\zeta_{i,j}) + \omega_j \text{Var}(\tau_j) - 1)}{\sum_{j=1}^n \omega_j} \\ &= \frac{\sum_{j=1}^n \omega_j \text{Var}(\zeta_{i,j})}{\sum_{j=1}^n \omega_j} \\ &= \text{Var}(\zeta_i) \\ &= \sigma_{w,i}^2 \end{aligned} \quad (54)$$

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

$$\hat{z}_{\phi,i} = \frac{\hat{\phi}_i}{\sqrt{\left(\mathbf{v}_i^\top \Omega \mathbf{v}_i \right)^{-2} \left((\mathbf{v}_i^\top \Omega^2 \mathbf{v}_i) \hat{\sigma}_{w,i}^2 + \mathbf{v}_i^\top \Omega \mathbf{v}_i \right)}} \quad (55)$$

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

$$\hat{z}_{\phi,i} \sim \mathcal{N}(z_{\phi,i}, 1) \quad (56)$$

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 $\Omega = \frac{c}{k}\mathbf{I}$ for coverage c and some constant k . Equation 55 simplifies to:

$$\hat{z}_{\phi,i} = \frac{\hat{\phi}_i \sqrt{n}}{\sqrt{\hat{\sigma}_{w,i}^2 + \frac{k}{c}}} \quad (57)$$

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.

5.3 Inference of causal variants with QTL and AS statistics

5.3.1 Modeling the correlation of summary statistics among markers

Due to linkage disequilibrium, there exist significant correlations of genotypes among markers. This correlation is reflected in the correlations in the association statistics. Given a set of m markers, we model a set of association statistics $\hat{\mathbf{z}}_{\alpha}$ for a locus as following a multivariate normal distribution with covariance $\mathbf{R}_{\mathbf{z}}$:

$$\hat{\mathbf{z}}_{\alpha} \sim \mathcal{N}_m(\mathbf{z}_{\alpha}, \mathbf{R}_{\mathbf{z}}) \quad (58)$$

Note that because the statistics are all of unit variance, $\mathbf{R}_{\mathbf{z}}$ is also the correlation matrix.

We now show that the correlation matrices for QTL and AS association statistics are both equivalent to the correlation matrix of the marker genotypes. Let $\mathbf{u}_{j,h}$ be the haploid 0/1 genotypes of markers on haplotype h of individual j . We assume that the genotypes are well-approximated by a multivariate normal distribution:

$$\mathbf{u}_{j,h} \sim \mathcal{N}_m(\boldsymbol{\mu}_{\mathbf{u}}, \boldsymbol{\Sigma}_{\mathbf{u}}) \quad (59)$$

The uncentered diploid 0/1/2 genotypes \mathbf{x}'_j can thus be expressed as the sum of two independent haploid genotypes of haplotypes A and B :

$$\mathbf{x}'_j = \mathbf{u}_{j,A} + \mathbf{u}_{j,B} \sim \mathcal{N}_m(2\boldsymbol{\mu}_{\mathbf{u}}, 2\boldsymbol{\Sigma}_{\mathbf{u}}) \quad (60)$$

Likewise, the -1/0/1 marker phases \mathbf{v}_j can be expressed as the difference of haploid genotypes:

$$\mathbf{v}_j = \mathbf{u}_{j,A} - \mathbf{u}_{j,B} \sim \mathcal{N}_m(\mathbf{0}, 2\boldsymbol{\Sigma}_{\mathbf{u}}) \quad (61)$$

Note that \mathbf{x}'_j and \mathbf{v}_j refer to the (uncentered) genotypes and phases across markers for a particular individual j . This is in contrast to \mathbf{x}_i and \mathbf{v}_i used earlier, which refer to the (centered) genotypes and phases across individuals for a particular marker i .

Examining Equations 21 and 24, we see that the QTL association statistic can be expressed as:

$$\hat{z}_{\beta,i} = f(\mathbf{x}_i) \left(\mathbf{x}_i^{\top} \mathbf{y} \right) \quad (62)$$

where $f(\mathbf{x}_i)$ is a vector-to-scalar function that ensures unit variance. Let \mathbf{X} be an n by m matrix of genotypes for all individuals and markers, and let \mathbf{F} be an m by m diagonal matrix such that $\mathbf{F}_{i,i} = f(\mathbf{x}_i)$. Across markers, the expression becomes:

$$\hat{\mathbf{z}}_{\beta} = \mathbf{F}\mathbf{X}^{\top} \mathbf{y} \quad (63)$$

Given that each row of \mathbf{X} is an independent realization of a multivariate-normally-distributed variable \mathbf{x}_j , the distribution of $\hat{\mathbf{z}}_\beta$ can be expressed as an affine transformation of the distribution of $\mathbf{x}_j = \mathbf{x}'_j - 2\boldsymbol{\mu}_u$.

$$\begin{aligned}\hat{\mathbf{z}}_\beta &\sim \mathcal{N}_m\left(\mathbf{z}_\beta, 2\left(\mathbf{y}^\top \mathbf{y}\right) \mathbf{F} \boldsymbol{\Sigma}_u \mathbf{F}\right) \\ &\sim \mathcal{N}_m\left(\mathbf{z}_\beta, \left(\sqrt{2\left(\mathbf{y}^\top \mathbf{y}\right) \mathbf{F}}\right) \boldsymbol{\Sigma}_u \left(\sqrt{2\left(\mathbf{y}^\top \mathbf{y}\right) \mathbf{F}}\right)^\top\right)\end{aligned}\quad (64)$$

Since $\sqrt{2\left(\mathbf{y}^\top \mathbf{y}\right) \mathbf{F}}$ is a diagonal matrix, the correlation matrix of $\hat{\mathbf{z}}_\beta$ is the same as the correlation matrix calculated from $\boldsymbol{\Sigma}_u$.

We examine the AS association statistic in a similar manner. Looking at Equations 51 and 55, the AS association statistic can be expressed as:

$$\hat{z}_{\phi,i} = g(\mathbf{v}_i, \boldsymbol{\Omega}) \left(\mathbf{v}_i^\top \boldsymbol{\Omega} \mathbf{w}\right) \quad (65)$$

where $g(\mathbf{v}_i, \boldsymbol{\Omega})$ is a vector-to-scalar function that ensures unit variance. Let \mathbf{V} be an n by m matrix of genotypes for all individuals and markers, and let \mathbf{G} be an m by m diagonal matrix such that $\mathbf{G}_{i,i} = g(\mathbf{v}_i, \boldsymbol{\Omega})$. Across markers, the expression becomes:

$$\hat{\mathbf{z}}_\phi = \mathbf{G} \mathbf{V}^\top \boldsymbol{\Omega} \mathbf{w} \quad (66)$$

The distribution of $\hat{\mathbf{z}}_\phi$ can thus be expressed as:

$$\begin{aligned}\hat{\mathbf{z}}_\phi &\sim \mathcal{N}_m\left(\mathbf{z}_\phi, 2\left((\boldsymbol{\Omega} \mathbf{w})^\top ((\boldsymbol{\Omega} \mathbf{w}))\right) \mathbf{G} \boldsymbol{\Sigma}_u \mathbf{G}\right) \\ &\sim \mathcal{N}_m\left(\mathbf{z}_\phi, \left(\sqrt{2\left((\boldsymbol{\Omega} \mathbf{w})^\top ((\boldsymbol{\Omega} \mathbf{w}))\right) \mathbf{G}}\right) \boldsymbol{\Sigma}_u \left(\sqrt{2\left((\boldsymbol{\Omega} \mathbf{w})^\top ((\boldsymbol{\Omega} \mathbf{w}))\right) \mathbf{G}}\right)^\top\right)\end{aligned}\quad (67)$$

Since $\boldsymbol{\Sigma}_u$ is transformed by a diagonal matrix, the correlation matrix of $\hat{\mathbf{z}}_\phi$ is the same as the correlation matrix calculated from $\boldsymbol{\Sigma}_u$.

Thus, both sets of summary statistics have the same correlation matrix \mathbf{R}_z , which is also the genotype correlation matrix

5.3.2 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} :

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

5.3.3 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:

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

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 $\lambda = -\log(\gamma)$. 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:

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

5.3.4 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.

5.3.5 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 to find all configurations with a significant probability. 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. 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. 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:

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

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

$$\hat{\mathbf{z}} = \begin{bmatrix} \hat{\mathbf{z}}_{\mathbf{C}} \\ \hat{\mathbf{z}}_{\mathbf{N}} \end{bmatrix} \quad (84)$$

We likewise permute the rows and columns of \mathbf{R} and $\Sigma_{\mathbf{c}}$ such that:

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

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

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

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

Blockwise inversion yields:

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

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

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

We can also simplify the determinant in Equation 83:

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

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

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

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:

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

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:

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

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:

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

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:

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

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:

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

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

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

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

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

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

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

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

$$h_{\hat{\phi}}^2 = h_{\phi}^2 \frac{\sigma_w^2}{\hat{\sigma}_w^2} \quad (100)$$

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:

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

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

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

We model the AS statistic in a similar manner:

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

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

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

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

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

893 The correlation is thus:

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