

# BEATRICE: Bayesian Fine-mapping from Summary Data using Deep Variational Inference

Sayan Ghosal<sup>1</sup>, Michael C. Schatz<sup>2</sup>, Archana Venkataraman<sup>1,3\*</sup>

**1** Department of Electrical and Computer Engineering, Johns Hopkins University, Baltimore, MD, USA

**2** Department of Computer Science, Johns Hopkins University, Baltimore, MD, USA

**3** Department of Electrical and Computer Engineering, Boston University, Boston, MA, USA

\* Correspondences should be directed to [archanav@bu.edu](mailto:archanav@bu.edu)

## Abstract

We introduce a novel framework BEATRICE to identify putative causal variants from GWAS summary statistics (<https://github.com/sayangsep/Beatrice-Finemapping>). Identifying causal variants is challenging due to their sparsity and to highly correlated variants in the nearby regions. To account for these challenges, our approach relies on a hierarchical Bayesian model that imposes a binary concrete prior on the set of causal variants. We derive a variational algorithm for this fine-mapping problem by minimizing the KL divergence between an approximate density and the posterior probability distribution of the causal configurations. Correspondingly, we use a deep neural network as an inference machine to estimate the parameters of our proposal distribution. Our stochastic optimization procedure allows us to simultaneously sample from the space of causal configurations. We use these samples to compute the posterior inclusion probabilities and determine credible sets for each causal variant. We conduct a detailed simulation study to quantify the performance of our framework across different numbers of causal variants and different noise paradigms, as defined by the relative genetic contributions of causal and non-causal variants. Using this simulated data, we perform a comparative analysis against two state-of-the-art baseline methods for fine-mapping. We demonstrate that BEATRICE achieves uniformly better coverage with comparable power and set sizes, and that the performance gain increases with the number of causal variants. Thus, BEATRICE is a valuable tool to identify causal variants from eQTL and GWAS summary statistics across complex diseases and traits.

## Author summary

Fine-mapping provides a way to uncover genetic variants that causally affect some trait of interest. However, correct identification of the causal variants is challenging due to the correlation structure shared across variants. While current fine-mapping approaches take into account this correlation structure, they are often computationally intensive to run and cannot handle spurious effects from non-causal variants. In this paper, we introduce BEATRICE, a novel framework for Bayesian fine-mapping from summary data. Our strategy is to impose a binary concrete prior over the causal configurations that can handle non-zero spurious effects and to infer the posterior probabilities of the causal variant locations using deep variational inference. In a simulation study, we

demonstrate that BEATRICE achieves comparable or better performance to the current fine-mapping methods across increasing numbers of causal variants and increasing noise, as determined by the polygenicity of the trait.

## 1 Introduction

Genome-Wide Association Studies (GWAS) provide a natural way to quantify the contribution each genetic variant to the observed phenotype [1]. However, the univariate nature of GWAS does not take into account the correlation structure shared between the genetic variants, which arises from low genomic recombination of nearby DNA regions [2]. Consequently, strong correlations can inflate the effect size of a non-causal genetic variant, thus leading to false positive identifications [3]. Fine-mapping [4, 5] addresses this problem by analyzing the correlation structure of the data to identify small subsets of causal genetic variants [5, 6]. These subsets, known as credible sets, capture the uncertainty of finding the true causal variant within a highly correlated region [7]. Unlike p-values, the corresponding posterior inclusion probabilities (PIPs) computed during fine-mapping can be compared across studies of different sample sizes.

Traditional fine-mapping methods can be grouped into two general categories. The first category uses a penalized regression model to predict the output phenotype based on the collection of genetic variants [8, 9]. Popular regularizations like LASSO [10] and Elastic Net [9] simultaneously perform effect size estimation while slowly shrinking the smaller effect sizes to zero. The drawback of penalized regression models is that they optimize phenotypic prediction and, due to the correlation structure, do not always identify the true causal variants. The second category relies on Bayesian modeling. Here, the phenotype is modeled as a linear combination of the genetic variants, with sparsity incorporated into the prior distribution for the model weights. Approximate inference techniques, such as Markov Chain Monte Carlo (MCMC) [11] and variational methods [12] have been used to infer the effect sizes, PIPs, and credible sets. While these approaches represent valuable contributions to the field, they require the raw genotype and phenotype information, which raises privacy and regulatory concerns, particularly in the cases of publicly shared datasets. MCMC sampling also requires a burn-in period, which adds a substantial (100X) runtime overhead.

In response to these concerns, fine-mapping approaches have moved towards using summary statistics, which can be easily shared across sites. For example, the works of [13–15] use a stochastic or exhaustive search to identify the posterior probabilities of the causal configurations. However, exhaustive search based methods are restricted by the number of assumed causal variants, as this leads to an exponential increase in the dimensionality of the approximate posterior distribution. Stochastic search approaches [13] are computationally less expensive, but, by construction, they cannot handle nontrivial effects from spurious non-causal variants. The most recent contribution to fine-mapping is SuSiE [16, 17], which estimates the variant effect sizes as a sum of “single effects”. These “single effect” vectors contain one non-zero element representing a causal variant and are estimated using a Bayesian step-wise selection approach. SuSiE provides a simple framework to robustly estimate PIPs and credible sets; however, there is limited evidence for its performance given the presence of spurious genetic effects. Such scenarios can appear due to polygenicity of the trait, trans-interactions of variants, or varying correlation structure of the genomic region.

In this paper, we introduce BEATRICE, a novel framework for Bayesian fine-mapping from summary data using deep variational Inference. In contrast to sampling methods, we approximate the posterior distribution of the causal variants given the GWAS summary statistics as a binary concrete distribution [18, 19], whose parameters are estimated using a deep neural network. This unique formulation allows

BEATRICE to use computationally efficient gradient-based optimization to minimize the KL divergence between the proposal binary concrete distribution and the posterior distribution of the causal variants. In addition, our unique optimization strategy samples a representative set of causal configurations in the process of minimizing the empirical KL divergence; these configurations can be used to obtain the PIPs and the credible sets. We compare our model with two state-of-the-art fine-mapping approaches, SuSiE [16] and FINEMAP [13]. We perform an extensive simulation study and quantify the performance of each model across increasing numbers of causal variants and increasing noise, as determined by the degree to which non-causal variants explain the phenotype variance. The runtimes of both SuSiE and BEATRICE are less than one minute, in contrast FINEMAP requires significantly longer to converge. On average BEATRICE achieves 2.2 fold increase in coverage, 0.1 fold increase in AUPRC, and similar power in comparison to SuSiE and FINEMAP.

## 2 Generative Assumptions of Fine-mapping

BEATRICE is based on a generative additive effect model. Formally, let  $\mathbf{y} \in \mathbb{R}^{n \times 1}$  denote a vector of (scalar) quantitative traits across  $n$  subjects. The corresponding genotype data  $\mathbf{X} \in \mathbb{R}^{n \times m}$  is a matrix, where  $m$  represents the number of genetic variants in the analysis. Without loss of generality, we assume that the columns of  $\mathbf{X}$  have been normalized to have mean 0 and variance 1, i.e.,  $\frac{1}{n} \sum_i \mathbf{X}_{ij} = 0$  and  $\frac{1}{n} \sum_i \mathbf{X}_{ij}^2 = 1$  for  $j = 1, \dots, m$ . The quantitative trait is generated as follows:

$$\mathbf{y} = \mathbf{X}\beta + \eta \quad \eta \sim N\left(0, \frac{1}{\tau} \mathbf{I}_n\right), \quad (1)$$

where  $\beta \in \mathbb{R}^{m \times 1}$  is the effect size,  $\eta \in \mathbb{R}^{n \times 1}$  is additive white Gaussian noise with variance  $\tau$ , and  $\mathbf{I}_n$  is the  $n \times n$  identity matrix.

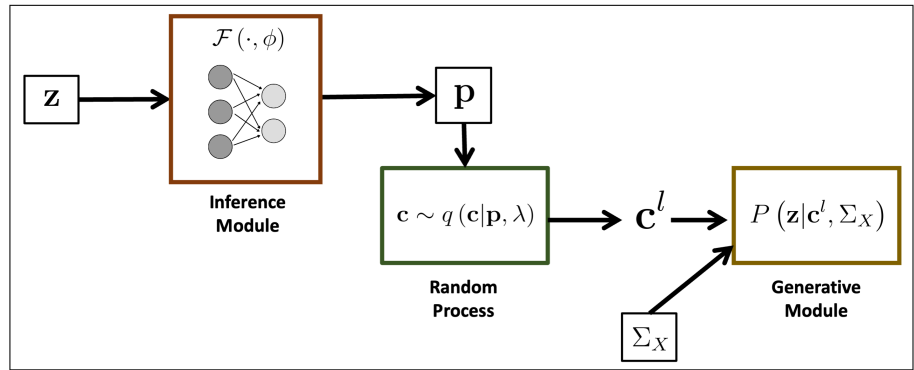
### 2.1 Genome Wide Association Studies (GWAS)

GWAS uses a collection of element-wise linear regression models to estimate the effect of each genetic variant. Mathematically, the GWAS effect sizes are computed as  $\hat{\beta} = \frac{1}{n} \mathbf{X}^T \mathbf{y}$ , with the corresponding vector of normalized z-scores equal to  $\mathbf{z} = \frac{1}{\sqrt{n\tau}} \mathbf{X}^T \mathbf{y}$  [1, 14]. The main drawback of GWAS is that non-causal genetic variants can have large effect sizes due to polygenicity of the quantitative trait [20], varying degrees of linkage disequilibrium (LD) with causal variants [3], and/or interactions of the variant with enriched genes [20]. One popular strategy to mitigate this drawback is to impose a sparse prior over  $\beta$  given the set of causal variants:

$$\beta \sim N\left(0, \frac{1}{\tau} \sigma^2 \Sigma_C\right) \quad (2)$$

$$\Sigma_C(i, j) = \begin{cases} 0, & i \neq j \\ 1, & i = j \text{ and } i \text{ is causal} \\ \epsilon, & i = j \text{ and } i \text{ is non-causal with non-zero effect} \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

Notice from Eq. (3) that the variance of  $\beta(i)$  for a causal variant is  $\frac{\sigma^2}{\tau}$  and the variance of  $\beta(i)$  for a non-causal variant with non-zero effect is  $\epsilon \frac{\sigma^2}{\tau}$ , where  $\epsilon$  is assumed to be small. This formulation handles residual influences from the non-causal variants,



**Fig 1.** Overview of BEATRICE. The inputs to our framework are the LD matrix  $\Sigma_X$  and the summary statistics  $\mathbf{z}$ . The inference module uses a neural network to estimate the underlying probability map  $\mathbf{p}$ . The random process generates random samples  $\mathbf{c}^l$  for the Monte Carlo integration in Eq. (12). Finally, the generative module calculates the likelihood of the summary statistics from the sample causal vectors  $\mathbf{c}^l$ .

which are often observed in real-world data. Under this assumed prior, we can show [14, 21] that the normalized GWAS effect sizes  $\mathbf{z}$  are distributed as:

$$p(\mathbf{z}|\Sigma_X, \Sigma_C) = N(\mathbf{z}; 0, \Sigma_X + \Sigma_X (n\sigma^2 \Sigma_C) \Sigma_X) \quad (4)$$

where  $\Sigma_X = \frac{1}{n} \mathbf{X}^T \mathbf{X}$  is the empirical correlation matrix of the genotype data, also known as the LD matrix. Broadly, the goal of fine-mapping is to identify the causal variants  $i$ , i.e., non-zero elements of  $\Sigma_C$  given the effect sizes  $\mathbf{z}$  and the LD matrix  $\Sigma_X$ . The derivation is provided in **Section S1** of the Supplement.

### 3 Materials and Methods

BEATRICE uses a variational inference framework for fine-mapping. For convenience, we represent the diagonal elements of  $\Sigma_C$  by the vector  $\mathbf{c} \in \mathbb{R}^{m \times 1}$ , and by construction,  $\mathbf{c}$  encodes the causal variant locations. Figure 1 provides an overview of BEATRICE. Our framework consists of three main components: an inference module, a random sampler, and a generative module. The inputs to BEATRICE are the summary statistics  $\mathbf{z}$  and the LD matrix  $\Sigma_X$ . The inference module estimates the parameters  $\mathbf{p}$  of our proposal distribution  $q(\cdot; \mathbf{p}, \lambda)$  using a neural network. The random process sampler uses the parameters  $\mathbf{p}$  to randomly sample potential causal vectors  $\mathbf{c}$  according to the given proposal distribution. Finally, the generative module calculates the likelihood of the observed summary statistics  $\mathbf{z}$  according to Eq. (4).

#### 3.1 Proposal Distribution

The goal of fine-mapping is to infer the posterior distribution  $p(\mathbf{c}|\{\mathbf{z}, \Sigma_X\})$ , where  $\mathbf{c}$  corresponds to the diagonal elements of  $\Sigma_C$ . Due to the prior formulation in Eqs. (2-3), solving for the true posterior distribution is computationally intractable, as it requires a combinatorial search over the possible causal configurations. Thus, we approximate the posterior distribution  $p(\mathbf{c}|\{\mathbf{z}, \Sigma_X\})$  with a binary concrete distribution  $q(\mathbf{c}; \mathbf{p}, \lambda)$  [18], where the parameters  $\mathbf{p}$  of the distribution are functions of the inputs  $\{\mathbf{z}, \Sigma_X\}$ . Samples  $\mathbf{c}$  generated under a binary concrete distribution can be viewed as continuous relaxations of independent Bernoulli random variables. This reparametrization [19] allows us to learn  $\mathbf{p}$  from the data using standard gradient descent.

Formally, let  $\mathbf{c}_i$  and  $\mathbf{p}_i$  denote the  $i$ th element of the vectors  $\mathbf{c}$  and  $\mathbf{p}$ , respectively. Each entry of  $\mathbf{c}$  is independent and is distributed as follows:

$$q(\mathbf{c}_i; \mathbf{p}_i, \lambda) = \frac{\lambda \mathbf{p}_i \mathbf{c}_i^{-\lambda-1} (1 - \mathbf{p}_i) (1 - \mathbf{c}_i)^{-\lambda-1}}{(\mathbf{p}_i \mathbf{c}_i^{-\lambda} + (1 - \mathbf{p}_i) (1 - \mathbf{c}_i)^{-\lambda})^2}, \quad (5)$$

where the parameter  $\lambda$  controls the extent of relaxation from a Bernoulli distribution.

We can easily sample from the binary concrete distribution in Eq. (5) via

$$\mathbf{c}_i = \xi \left( \frac{\log \left( \frac{U}{1-U} \right) + \log \left( \frac{\mathbf{p}_i}{1-\mathbf{p}_i} \right)}{\lambda} \right), \quad (6)$$

where  $\xi(\cdot)$  is the sigmoid function, and the random variable  $U$  is sampled from a uniform distribution over the interval  $[0, 1]$ . As seen,  $\mathbf{p}_i$  specifies the underlying probability map and  $U$  provides stochasticity for the sampling procedure in Eq. (6). We note that the gradient of Eq. (6) with respect to  $\mathbf{p}_i$  tends to have a low variance in practice, which helps to stabilize the optimization.

The two unique properties of binary concrete random variable are  $P(\mathbf{c}_i > \frac{1}{2}) = \mathbf{p}_i$  and  $\lim_{\lambda \rightarrow 0} P(\mathbf{c}_i = 1) = \mathbf{p}_i$ . The first property indicates that  $\mathbf{p}_i$  controls the degree to which  $\mathbf{c}_i$  assumes low values close to 0 and high values close to 1. This property also give BEATRICE flexibility to handle genetic variants with different levels of association, thus aligning with our generative process that assumes some non-causal variants may have small non-zero effects. The second property implies that a high probability  $\mathbf{p}_i$  at location  $i$  is highly indicative of a causal variant. Taken together, the binary concrete distribution has an easily-optimized parameterization with desirable properties.

### 3.2 Variational Inference

We select the variational parameters  $\{\mathbf{p}, \lambda\}$  to minimize the Kullback–Leibler (KL) divergence between the proposal distribution and the posterior distribution of the causal vector  $\mathbf{c}$  given the input data  $\{\mathbf{z}, \Sigma_{\mathbf{X}}\}$ , that is

$$\{\mathbf{p}^*, \lambda^*\} = \arg \min_{\{\mathbf{p}, \lambda\}} KL(q(\mathbf{c}; \mathbf{p}, \lambda) || p(\mathbf{c} | \{\mathbf{z}, \Sigma_{\mathbf{X}}\})) \quad (7)$$

Using Bayes' Rule, we can show that the optimization in Eq. (7) can be rewritten

$$\{\mathbf{p}^*, \lambda^*\} = \arg \min_{\{\mathbf{p}, \lambda\}} KL(q(\mathbf{c}; \mathbf{p}, \lambda) || p(\mathbf{c}; \mathbf{p}_0, \lambda_0)) - E_{q(\cdot; \mathbf{p}, \lambda)} [\log(p(\mathbf{z} | \Sigma_{\mathbf{X}}, \mathbf{c}))], \quad (8)$$

where we have assumed an element-wise binary concrete prior  $p(\mathbf{c}; \mathbf{p}_0, \lambda_0)$  over the vector  $\mathbf{c}$ . We fix the relaxation parameter to be small ( $\lambda = 0.01$ ) and the probability map to be uniform  $\mathbf{p}_0 = [\frac{1}{m}, \dots, \frac{1}{m}]^T$ . Thus, the first term of Eq. (8) can be viewed as a regularizer that encourages sparsity in causal vectors  $\mathbf{c}$ . The second term of Eq. (8) can be interpreted as the likelihood of the observed test statistics. The works of [22, 23] have demonstrated that under certain assumptions, the likelihood term of the summary statistics is the same as the original data likelihood  $p(\mathbf{y} | \mathbf{X}, \mathbf{c})$  derived from Eq. (1).

During optimization, the relaxation parameter  $\lambda$  is annealed [18, 19] to a small non-zero value (0.01) with fixed constant rate, and the underlying probability map  $\mathbf{p}$  is optimized using gradient descent. Specifically, we use a neural network to generate the vector  $\mathbf{p} = \mathcal{F}(\mathbf{z}; \phi)$ . The details of the neural network architecture are provided in **Section S3** of the Supplement. Optimizing  $\mathbf{p}^*$  now amounts to optimizing the learnable parameters of the neural network  $\phi$ . Given a fixed value of  $\lambda$ , the neural network loss function follows from Eq. (8) according to

$$\mathcal{L}(\phi) = KL(q(\mathbf{c}; \mathbf{p}(\phi), \lambda) || p(\mathbf{c}; \mathbf{p}_0, \lambda_0)) - E_{q(\cdot; \mathbf{p}(\phi), \lambda)} [\log(p(\mathbf{z} | \Sigma_{\mathbf{X}}, \mathbf{c}))], \quad (9)$$

where we have defined  $\mathbf{p}(\phi) \triangleq \mathcal{F}(\mathbf{z}; \phi)$  for notational convenience.

### 3.3 Optimization Strategy

The expectations in Eq. (9) do not have closed-form expressions. Therefore, we use Monte Carlo integration to accurately approximate  $\mathcal{L}(\phi)$  in the regime of small  $\lambda$ , i.e., when the binary concrete distribution behaves similar to a Bernoulli distribution.

Let  $\mathbf{c}^1(\phi), \dots, \mathbf{c}^L(\phi)$  be a collection of causal vectors sampled independently from  $q(\cdot|\mathbf{p}(\phi), \lambda)$  according to Eq. (6). The likelihood term of Eq. (9) is computed as

$$E_{q(\cdot|\mathbf{p}(\phi), \lambda)} [\log(p(\mathbf{z}|\boldsymbol{\Sigma}_{\mathbf{X}}, \mathbf{c}))] = \frac{1}{L} \sum_{l=1}^L \log \left( p \left( \mathbf{z} | \boldsymbol{\Sigma}_{\mathbf{X}}, \mathbf{c}^l(\phi) \right) \right), \quad (10)$$

where the right-hand side probability is computed according to Eq. (4) by substituting  $\mathbf{c}^l(\phi)$  for the diagonal entries of  $\boldsymbol{\Sigma}_C$  in each term of the summation. Once again, the continuous relaxation used to generate  $\mathbf{c}^l(\phi)$  in Eq. (6) allows us to directly optimize  $\phi$ .

We approximate the first term of Eq. (9) under the assumption of small  $\{\lambda, \lambda_0\}$  on the order of 0.01. In this case, the binary concrete distribution behaves like a  $\{0, 1\}$  Bernoulli distribution. Under these conditions, we can write the first term of Eq. (9) as

$$\begin{aligned} & KL(q(\mathbf{c}; \mathbf{p}(\phi), \lambda) \| p(\mathbf{c}; \mathbf{p}_0, \lambda_0)) \\ & \approx \sum_{i=1}^m \left[ \mathbf{p}_i(\phi) \log \left( \frac{\mathbf{p}_i(\phi)}{p_0} \right) + (1 - \mathbf{p}_i(\phi)) \log \left( \frac{1 - \mathbf{p}_i(\phi)}{1 - p_0} \right) \right], \end{aligned} \quad (11)$$

where  $p_0$  is a fixed scalar parameter used to construct the (constant) prior vector  $\mathbf{p}_0$ . We note that the criteria  $\{\lambda \rightarrow 0.01, \lambda_0 = 0.01\}$  is satisfied in practice, as  $\lambda$  is annealed during the optimization to progressively smaller values and  $\lambda_0$  is fixed *a priori*.

The above approximations allow us to rewrite the neural network loss as

$$\begin{aligned} \mathcal{L}(\phi) & \approx -\frac{1}{L} \sum_{l=1}^L \log N \left( \mathbf{z}; 0, \boldsymbol{\Sigma}_X + \boldsymbol{\Sigma}_X \left( n \sigma^2 \boldsymbol{\Sigma}_C^l(\phi) \right) \boldsymbol{\Sigma}_X \right) \\ & \quad + \sum_{i=1}^m \mathbf{p}_i(\phi) \log \left( \frac{\mathbf{p}_i(\phi)}{p_0} \right) + (1 - \mathbf{p}_i(\phi)) \log \left( \frac{1 - \mathbf{p}_i(\phi)}{1 - p_0} \right), \end{aligned} \quad (12)$$

where  $\boldsymbol{\Sigma}_C^l(\phi)$  corresponds to the diagonal matrix using the vector  $\mathbf{c}^l(\phi)$  as the diagonal entries. We use a stochastic gradient descent optimizer [24] to minimize the loss  $\mathcal{L}(\phi)$  with respect to the neural network weights  $\phi$ . This process is detailed in Algorithm 1.

---

#### Algorithm 1 Optimization scheme to minimize Eq. (12)

---

```

 $\mathcal{B}^R = \{\}$ 
Initialize  $\phi_0$ 
for  $t = [1 \dots T]$  do
    Generate  $\mathbf{p}(\phi_t) = \mathcal{F}(\mathbf{z}; \phi_t)$ 
    Randomly sample  $\mathbf{c}_t^l$  according to Eq. (6)
    Binarize  $\mathbf{c}_t^l$  to  $\mathbf{b}_t^l$  and add to  $\mathcal{B}^R$ 
     $\mathcal{S}_t^l = \{i\}$  s.t.  $\mathbf{c}_t^l(i) > 0.01$ 
    Prune set  $\mathcal{S}_t^l$  such that it consists of 50 indices.
     $\mathbf{c}_t^l(i) = 0$  if  $i \notin \mathcal{S}_t^l$ 
    Generate  $\mathcal{L}(\phi_t)$  according to Eq. (12)
     $\phi_{t+1} = \phi_t - \text{StepSize} \nabla \mathcal{L}(\phi_t)$ 
end for

```

---

### 3.4 Computational Complexity 153

Each iteration of stochastic gradient descent requires us to compute the likelihood term  $\left[ \log N \left( \mathbf{z}; 0, \boldsymbol{\Sigma}_X + \boldsymbol{\Sigma}_X \left( n \sigma^2 \boldsymbol{\Sigma}_C^l(\phi) \right) \boldsymbol{\Sigma}_X \right) \right]$ . This computation is expensive due to the covariance matrix inversion, whose run-time is on the order of  $O(m^3)$ , where  $m$  is the total number of variants. To mitigate this issue, the works of [25] show that if  $\boldsymbol{\Sigma}_C^l(\phi)$  is sparse, then the matrix inversion can be done with order  $O(k^3) + O(mk^2)$  run-time, where  $k$  is the number of non-zero diagonal elements of  $\boldsymbol{\Sigma}_C^l(\phi)$ . We leverage this result in the optimization by thresholding the elements of  $\mathbf{c}^l(\phi)$  to set small values exactly to zero. We specify this threshold such that at most 50 elements of  $\mathbf{c}^l(\phi)$  are non-zero at each iteration. This choice allows us to run BEATRICE in fixed time for all scenarios for a fixed  $m$ . We also regularize  $\boldsymbol{\Sigma}_X$  with a small diagonal load to ensure invertibility of the covariance matrix at each iteration. Finally, we run stochastic gradient descent with a batch size of one to further speed up BEATRICE. Effectively, this means that we sample a single  $\mathbf{c}^l(\phi)$  at each epoch rather than perform a true Monte Carlo integration. The authors of [26] have previously shown that a single random sample ( $L = 1$ ) is sufficient to guarantee convergence to a local minimum of Eq. (12). Algorithm 1 provides a detailed description of these optimization steps. 154  
155  
156  
157  
158  
159  
160  
161  
162  
163  
164  
165  
166  
167  
168  
169

### 3.5 Causal Configurations and Posterior Inclusion Probabilities 170

The desired outputs of each fine-mapping method are Posterior Inclusion Probabilities (PIPs) and credible sets. PIPs estimate how likely each variant is causal as a measure of its importance. Credible sets identify the subset of variants that are likely to contain a causal variant, which captures the uncertainty of finding the true variant. 171  
172  
173  
174

The main challenge to estimating the posterior probability of a given causal configuration (i.e., set of causal variant locations) is the exponentially large search space. Let  $\mathbf{b}$  denote a binary vector with a value of 1 at causal locations and a value of 0 at non-causal locations. At a high level,  $\mathbf{b}$  can be viewed as a binarized version of the causal vector  $\mathbf{c}$  in the previous sections. Using Bayes' Rule, the posterior probability of  $\mathbf{b}$  given the input data  $\{\mathbf{z}, \boldsymbol{\Sigma}_X\}$  can be written as follows: 175  
176  
177  
178  
179  
180

$$p(\mathbf{b}|\mathbf{z}, \boldsymbol{\Sigma}_X) = \frac{p(\mathbf{z}|\boldsymbol{\Sigma}_X, \mathbf{b}) p(\mathbf{b})}{\sum_{\mathbf{b}' \in \mathcal{B}} p(\mathbf{z}|\boldsymbol{\Sigma}_X, \mathbf{b}') p(\mathbf{b}')} \quad (13)$$

where  $\mathcal{B}$  is the set of all  $2^m$  possible causal configurations. Once again,  $\mathbf{z}$  captures the summary statistics and  $\boldsymbol{\Sigma}_X$  is the LD matrix. Even though  $\mathcal{B}$  is exponentially large, it has been argued [27] that the majority of these configurations have negligible probability and do not contribute to the denominator of Eq. (13). 181  
182  
183  
184

Our stochastic optimization provides a natural means to track causal configurations with non-negligible probability to compute  $p(\mathbf{b}|\mathbf{z}, \boldsymbol{\Sigma}_X)$ . Namely, at each iteration of stochastic gradient descent, we randomly generate a sample causal vector  $\mathbf{c}^l$  to minimize Eq. (12). In parallel, we binarize the vector  $\mathbf{c}^l$  via

$$\mathbf{b}_i^l = \begin{cases} 1, & \mathbf{c}_i^l > \gamma, \\ 0, & \text{otherwise} \end{cases}$$

and add the resulting vector  $\mathbf{b}^l$  to a reduced set of causal configurations  $\mathcal{B}^R$ . The variational objective ensures that our proposal distribution converges to the true posterior distribution of the causal vectors. Thus, the samples  $\mathbf{c}^l$  lie near modes of the posterior distribution which is the neighborhood of non-negligible probability. 185  
186  
187  
188

In this work, we use a threshold  $\gamma = 0.1$  to binarize the vectors  $\mathbf{c}^l$ . Empirically, we find that this threshold value preserves the main interactions between variants. However, the user of BEATRICE can adjust this threshold as needed. 189  
190  
191

After obtaining the sampled vectors, we replace the exhaustive set  $\mathbf{B}$  in Eq. (13) with the reduced set  $\mathcal{B}^R$  for tractable computation of  $p(\mathbf{b}|\mathbf{z}, \Sigma_X)$ . We then compute the posterior inclusion probability (PIP) of each variant by summing the probabilities over the subset of  $\mathcal{B}^R$  with a value of 1 at that variant location. Mathematically,

$$P(\mathbf{b}_i = 1|\mathbf{z}, \Sigma_X) \approx \sum_{\mathbf{b} \in \mathcal{S}} p(\mathbf{b}|\mathbf{z}, \Sigma_X) \quad (14)$$

$$\text{s.t. } \mathcal{S} \subset \mathcal{B}^R \text{ and } \mathcal{S} = \{\mathbf{b} | \mathbf{b}_i = 1\} \quad (15)$$

where  $\mathcal{S}$  is a subset of  $\mathcal{B}^R$  that contains binary configurations with 1 at location  $i$ . 192

Finally, we identify the credible sets in two steps. First, in a conditional step-wise 193 fashion, we identify the variants with the highest conditional probability given the 194 previously selected variants. This strategy identifies the set of “key” variants with a high 195 probability of being causal. Second, we determine the credible set for each key variant, 196 by computing the conditional inclusion probabilities of each variant given the key 197 variants and adding variants to the credible set. A detailed description of this process 198 can be found in the Supplementary Methods document (**Section S2 in S1 text**). 199

### 3.6 Baselines 200

We compare our approach with the state-of-the-art methods, FINEMAP and SuSiE. 201

**FINEMAP:** This approach uses a stochastic shotgun search to identify causal 202 configurations with non-negligible posterior probability. FINEMAP defines the 203 neighborhood of a configuration at every step by deleting, changing or adding a causal 204 variant from the current configuration. The next iteration samples from this 205 neighborhood, thus reducing the exponential search space to a smaller high-probability 206 region. Finally, the identified causal configurations are used to determine the posterior 207 inclusion probabilities for each variant. The computationally efficient shotgun approach 208 makes FINEMAP a viable tool for finemapping from multiple GWAS summary data 209 in [28, 29]. However, the FINEMAP algorithm [13] does not provide definitive credible 210 sets, thus we rely on the same approach used in [16] for this task. Details of this 211 procedure are provided in **Section S4** of the Supplement. 212

**SuSiE:** The recent works of [16, 17] introduced an iterative Bayesian selection 213 approach for fine-mapping that represents the variant effect sizes as a sum of 214 “single-effect” vectors. Each vector contains only one non-zero element, which represents 215 the causal signal. In addition to finding causal variants, SuSiE provides a way to 216 quantify the uncertainty of the causal variants locations via credible sets. SuSiE has also 217 been used widely to find putative causal variants GWAS summary statistics [30, 31]. 218

### 3.7 Evaluation Strategy 219

We evaluate several metrics of performance in our simulation study. 220

**Area Under Precision Recall Curve (AUPRC):** We threshold the PIPs and 221 compute the precision and recall for identifying the ground-truth causal configuration. 222 High precision indicates a low false positive rate, while high recall relates to a low false 223 negative rate. Thus, the AUPRC, which is computed by sweeping the PIP threshold, 224 can be viewed as a holistic measure of performance across both classes. AUPRC is also 225 robust to severe class imbalance [32], which is the case in fine-mapping. 226



**Coverage, Power and Size of the Credible Sets:** We assess the quality of the credible sets using three metrics [16,17]. Coverage is the percentage of credible sets that contain at least one true causal variant, and power is the percentage of ground-truth causal variants that appear in at least one credible set. Higher coverage indicates that the method is more confident about its prediction of causal variants, whereas higher power indicates the method can accurately identify all the causal variants.

One caveat is that we can generally achieve both higher coverage and higher power simply by adding variants to the credible sets. To counter this trend, we report the average size of the credible sets identified by each method. Ideally, we would like the credible sets to be as small as possible while retaining high coverage and high power.

## 4 Experimental Setup

**Genotype Simulations:** We use the method of [33] to simulate genotypes  $\mathbf{X}$  based on data from the 1000 Genomes Project. We select an arbitrary sub-region (39.9Mb – 40.9Mb) from Chromosome 2 as the base. After filtering for rare variants (MAF < 0.02), the remaining 3.5K variants are used to simulate pairs of haplotypes to generate 10,000 unrelated individuals. In each experiment below, we randomly select  $m = 1k$  variants and  $n = 5000$  individuals to generate the phenotype data.

**Phenotype Generation:** We generate the phenotype  $\mathbf{y}$  from a standard mixed linear model [22], where the influences of the causal variants are modeled as fixed effects, and the influences of other non-causal variants are modeled as random effects. In this case, the genetic risk for a trait is spread over the entire dataset, with each variant having small individual effects, as per the polygenicity assumption of a complex trait.

Given a set of  $d$  causal variants  $C$ , let  $\mathbf{X}_C \in \mathbf{R}^{n \times d}$  denote the corresponding subset of the genotype data and  $\mathbf{X}_{NC} \in \mathbf{R}^{n \times m-d}$  denote the remaining non-causal variants. From here, we generate the phenotype data  $\mathbf{y}$  as follows:

$$\begin{aligned} \mathbf{y} &= \mathbf{X}_C \beta + \mathbf{g}_{NC} + \epsilon \triangleq \mathbf{g}_C + \mathbf{g}_{NC} + \epsilon \\ \mathbf{g}_{NC} &\sim N\left(0, \frac{1}{m-d} \mathbf{X}_{NC} \mathbf{X}_{NC}^T\right) \\ \beta &\sim N(0, \mathbf{I}_d) \\ \epsilon &\sim N(0, \alpha^2 \mathbf{I}_n) \end{aligned}$$

where  $\beta$  is the  $d$ -dimensional effect sizes sampled from a Gaussian, and  $\epsilon$  is an zero-mean Gaussian noise with variance  $\alpha^2$ . The random variable  $\mathbf{g}_{NC}$  models the effect of the non-causal variants as a multivariate Gaussian vector with mean 0 and covariance  $\frac{1}{m-d} \mathbf{X}_{NC} \mathbf{X}_{NC}^T$ . Likewise,  $\mathbf{g}_C = \mathbf{X}_C \beta$  captures the effect of the causal variants.

In our experiments, we define  $\omega^2$  as the total phenotypic variance attributed to the genotype (e.g., both  $\mathbf{g}_C$  and  $\mathbf{g}_{NC}$ ) and  $p$  as the proportion of this variance associated with the causal variants in  $\mathbf{g}_C$ . Using the strategy described in [34], we enforce these conditions by normalizing the phenotype  $\mathbf{y}$  as follows:

$$\begin{aligned} \tilde{\mathbf{y}} &= \sqrt{\frac{p\omega^2}{\text{var}(\mathbf{g}_C)}} \mathbf{g}_C + \sqrt{\frac{(1-p)\omega^2}{\text{var}(\mathbf{g}_{NC})}} \mathbf{g}_{NC} + \tilde{\epsilon} \\ \tilde{\epsilon} &\sim N(0, (1-\omega^2)\mathbf{1}_n) \end{aligned} \tag{16}$$

where  $\text{var}(\mathbf{g}_C)$  and  $\text{var}(\mathbf{g}_{NC})$  are the empirical variances of  $\mathbf{g}_C$  and  $\mathbf{g}_{NC}$ , respectively.

After generating the genotype  $\mathbf{X}$  and the normalized phenotype  $\tilde{\mathbf{y}}$ , we run a GWAS to estimate the effect size  $\hat{\beta}_i$  of each variant  $i$ . From here, we convert the estimated effect sizes to z-scores via  $\mathbf{z}_i = \frac{\hat{\beta}_i}{se(\hat{\beta}_i)}$ , where  $se(\cdot)$  denotes the standard error. The LD matrix is computed from the genotype data as  $\Sigma_X = \frac{1}{n} \mathbf{X}^T \mathbf{X}$ . The z-scores and LD matrix are input to each of the fine-mapping methods above.

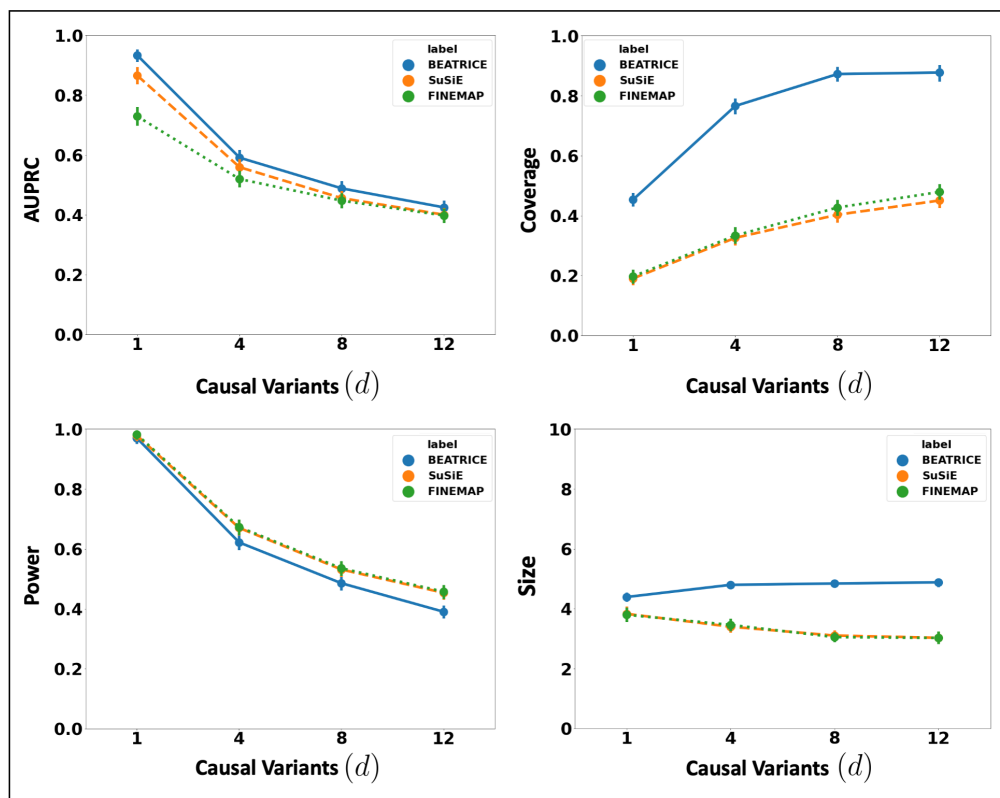
**Noise Configurations:** We evaluate the performance of each method while varying the number of causal variants  $d$ , the total genotype variance  $\omega^2$ , and the proportion of this variance associated with the causal variants  $p$ . Formally, we sweep over one order of magnitude for  $d = [1, 4, 8, 12]$ ,  $\omega^2 = [0.1, 0.2, 0.4, 0.5, 0.7, 0.8]$ , and  $p = [0.1, 0.3, 0.5, 0.7, 0.9]$ . For each noise setting, we randomly generate 20 datasets by independently re-sampling the causal variant locations, the effect sizes  $\{\beta_i\}$ , the non-causal component  $\mathbf{g}_{NC}$ , and the noise  $\tilde{\epsilon}$ . We run all three fine-mapping methods over a total of  $4 \times 6 \times 5 \times 20 = 2400$  configurations for a comprehensive evaluation.

## 5 Results

**Varying the Number of Causal Variants** Figure 2 illustrates the performance of each method (BEATRICE, FINEMAP, and SuSiE) while increasing the number of causal variants from  $d = 1$  to  $d = 12$ . The points denote the mean performance across all noise configurations  $(\omega^2, p)$  for fixed  $d$ , and the error bars represent the 95% confidence interval across these configurations. We note that BEATRICE achieves a uniformly higher AUPRC than both baseline method, which suggests that BEATRICE can better estimate the PIPs than FINEMAP or SuSiE. BEATRICE also provides 0.9 – 1.4 fold increase in coverage than the baselines with similar power, which indicates that the credible sets generated by BEATRICE are more likely to contain a causal variant as compared to SuSiE and FINEMAP. Finally, we note that although FINEMAP and SuSiE identify smaller credible sets, the difference in set size between them and BEATRICE is  $< 2$  variants. Taken together, as the number of causal variants increases, BEATRICE gives us a better estimate of the PIPs and arguably better credible sets. Compared to the baselines BEATRICE does not impose any prior assumptions over the total number of causal variants, which may lead to its improved performance.

**Increasing the Genotype Contribution:** Figure 3 shows the performance of each method while increasing the genetically-explained variance from  $\omega^2 = 0.1$  to  $\omega^2 = 0.8$ . Similar to above, the points denote the mean performance across all configurations  $(d, p)$  for fixed  $\omega^2$ , and the error bars represent the 95% confidence interval across these configurations. We note that BEATRICE achieves a significantly higher AUPRC than FINEMAP and a slightly higher AUPRC than SuSiE. When evaluating the credible sets, we observe similar trends in coverage (BEATRICE is 0.25 – 2.34 folds higher) and power (similar performance across methods). While the FINEMAP and SuSiE identify slightly smaller credible, the difference to BEATRICE is only a few variants. Taken together, we submit that BEATRICE achieves the best trade-off across the four performance metrics.

**Varying the Contributions of Causal and Non-Causal Variants:** Figure 4 illustrates the performance of each method while increasing the contribution of the causal variants from  $p = 0.1$  to  $p = 0.9$ . Once again, the points denote the mean performance across all configurations configurations  $(d, \omega^2)$  for fixed  $p$ , and the error bars represent the 95% confidence interval across these configurations. From an application standpoint, the presence of non-causal variants with small non-zero effects makes it difficult to detect the true causal variants. Accordingly, we observe a

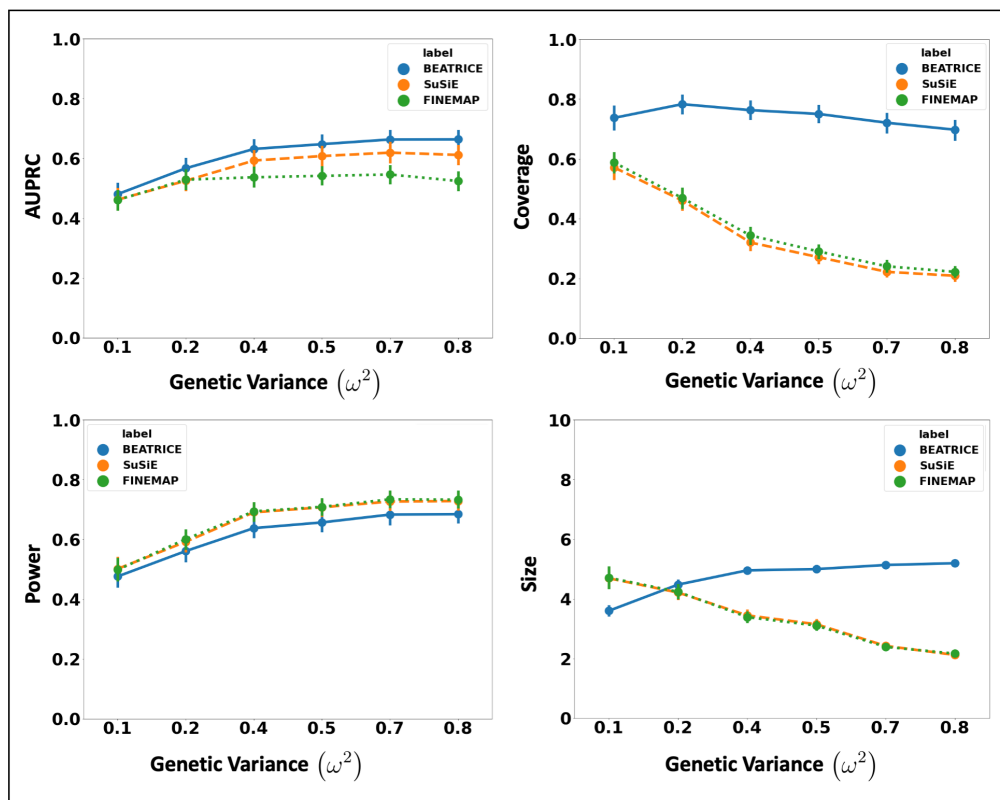


**Fig 2.** The performance metrics for the three methods across varying numbers of causal variants. Along the x-axis, we plot the number of causal variants, and across the y-axis, we plot the mean and confidence interval (95%) of each metric. We calculate the mean by fixing  $d$  to a specific value  $d = d^*$  and sweep over all the noise settings where  $d = d^*$ .

performance boost across all methods when  $p$  is larger. Similar to our previous experiments, BEATRICE provides the best AUPRC, with converging performance as  $p \rightarrow 1$ . In addition, BEATRICE identifies better credible sets with significantly higher coverage while maintaining power. Thus, we conclude that BEATRICE is the most robust of the three methods to the presence of noise from non-causal variants. This performance gain may arise from our binary concrete proposal distribution for the causal vector  $\mathbf{c}$ , which provides flexibility to accommodate varying degrees of association.

## 6 Discussion

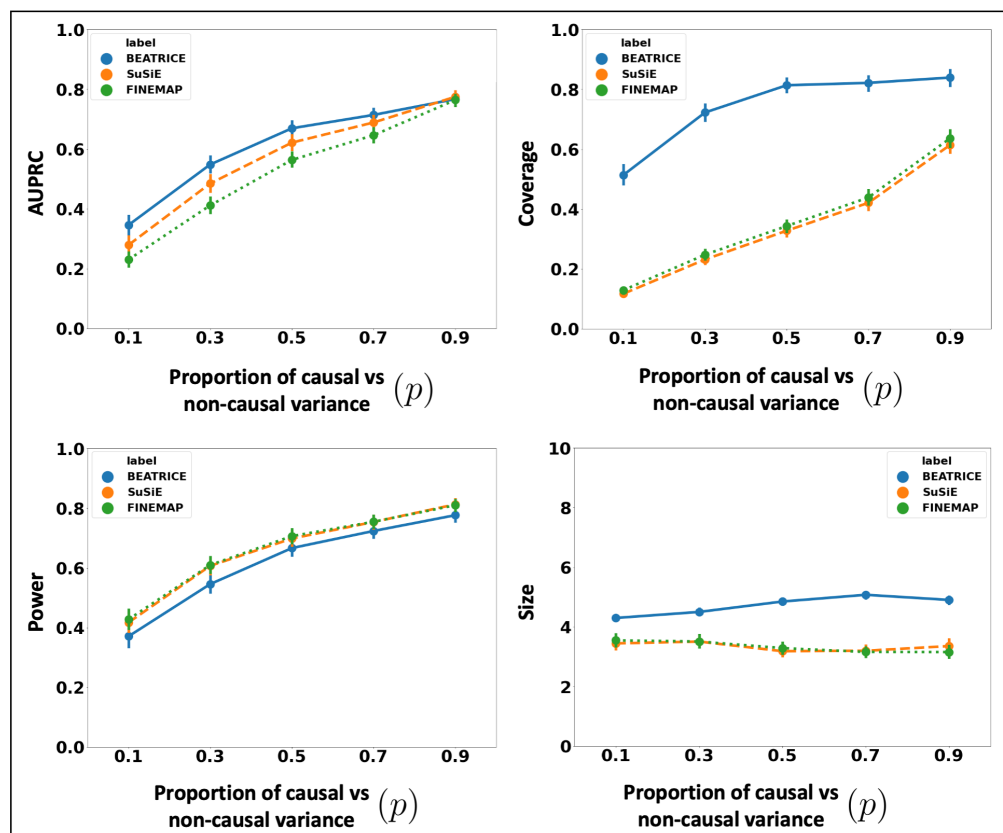
BEATRICE is a novel, robust, and general purpose tool for fine-mapping that can be used across a variety of studies. One key contribution of BEATRICE over methods like FINEMAP and SuSiE is its ability to discern spurious effects from non-causal variants, including non-causal variants in high LD with true causal variants. Our simulated experiments capture this improved performance by sweeping the proportion of the observed variance attributed to causal (fixed effects) and non-causal (random effects) genetic variants. This parameter  $p \in [0, 1]$  is swept over its natural domain, such that  $p = 1$  implies that the only link between the genotype and phenotype comes from the causal variants. At this extreme, Figure 4 shows that all methods achieve comparable performance. However, as  $p$  decreases, meaning that the effects of non-causal variants increase, BEATRICE outperforms both baselines.



**Fig 3.** The performance metric for increasing phenotype variance explained by genetics. Along the x-axis, we plot the variance explained by genetics ( $\omega^2$ ), and across the y-axis, we plot each metric’s mean and confidence interval (95%). We calculate the mean by fixing  $\omega^2$  to a specific value  $\omega = \omega^*$  and sweep over all the noise settings where  $\omega = \omega^*$ .

We further probe this behavior by illustrating the element-wise PIPs and the credible sets identified by all three methods under two simulation settings:  $\{d = 1, \omega^2 = 0.2, p = 0.9\}$  (Figure 5) and  $\{d = 1, \omega^2 = 0.2, p = 0.1\}$  (Figure 6). As seen in Figure 5, the variance explained by the non-causal variants is small, so the causal variant is easy to distinguish and has been correctly identified by all three approaches. In contrast, we see in Figure 6 that when the non-causal variants play a larger role, the causal variant no longer has the maximum GWAS  $z$ -score. Here, only BEATRICE correctly identifies the causal variant and assigns it the highest PIP. Both FINEMAP and SuSiE give uncertain predictions, as captured by the large credible sets and multiple high PIPs. We conjecture that BEATRICE takes advantage of the binary concrete distribution to model non-causal variants with non-zero effects, while using the sparsity term of  $\mathcal{L}(\cdot)$  to prioritize potentially causal variants.

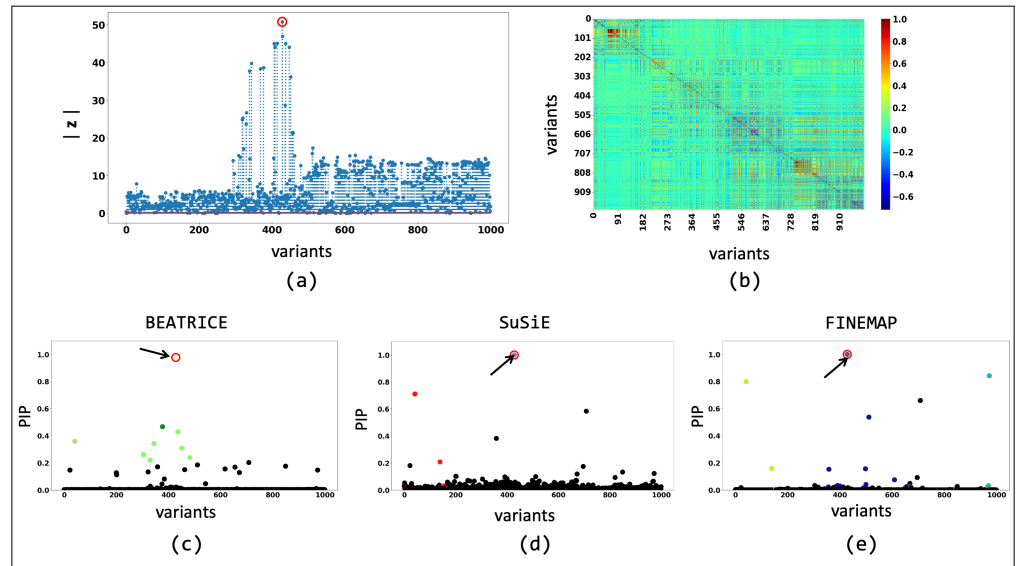
A second contribution of BEATRICE is our strategic integration of neural networks within a larger statistical framework. Specifically, we use the neural network in Figure 1 as an inference engine to estimate the parameters  $\mathbf{p}$  of our proposal distribution. In this case, the standard over-parameterization in the neural network helps BEATRICE to manage the complexity of the data while providing a buffer against overfitting. BEATRICE leverage the continuous representation of the causal vectors  $\mathbf{c}^l$  to backpropagate the gradients through the random sampler and train the network. Additionally, the continuous representation of  $\mathbf{c}^l$  results in low-variance gradients with respect to the underlying probability map, thus leading to a stable optimization.



**Fig 4.** The performance metric for multiple levels of noise introduced by non-causal variants. The noise level ( $p$ ) is explained by the variance ratio of non-causal variants vs. causal variants. Along the x-axis, we plot the noise level ( $p$ ); across the y-axis, we plot each metric's mean and confidence interval (95%). We calculate the mean by fixing  $p$  to a specific value  $p = p^*$  and sweep over all the noise settings where  $p = p^*$ .

Related to the above point, a third contribution of BEATRICE is its ability to efficiently build and evaluate a representative set of causal configurations during the optimization process. This set identifies key regions of the exponential search space to compute the PIPs and credible sets. In particular, we keep track of the sampled vectors at every iteration of the optimization, as described in Section 3.5. By minimizing the KL divergence between the proposal distribution and the true posterior distribution, we ensure that the randomly sampled causal vectors slowly converge to the causal configurations that have non-negligible posterior probability. Our strategy lies in stark contrast with traditional mean-field approaches, where independence assumptions between elements of the proposal distribution do not allow for joint inference of the causal configurations. Furthermore, this strategy allows us to efficiently estimate the PIPs in finite run-time. Figure 7 compares the average run-time of each method across all parameter settings. We observe that the run-time of BEATRICE and SuSiE are less than one minute. In contrast, FINEMAP requires significantly more time to converge.

The final contribution of BEATRICE is its simple and flexible design. Importantly, BEATRICE can easily incorporate priors based on the functional annotations of the variants. Formally, in the current setup, the prior over  $\mathbf{c}$  is effectively constant, as captured by  $p_0 = \frac{1}{m}$ . We can integrate prior knowledge simply by modifying the distribution of  $p_0$  across the variants. Thus, BEATRICE is a general-purpose tool for



**Fig 5.** The fine-mapping performance of BEATRICE , SuSiE, and FINEMAP at a noise setting of  $\{d = 1, \omega^2 = 0.2, p = 0.9\}$ . (a) The absolute z-score of each variant as obtained from GWAS. (b) Pairwise correlation between the variants. (c), (d), and (e) are the posterior inclusion probabilities of each variant as identified by BEATRICE , SuSiE, and FINEMAP, respectively. The red circle marked by an arrow shows the location of the causal variant. We have further color-coded the variants based on their assignment to credible sets. The non-black markers represent the variants assigned to a credible set. Additionally, the variants in a credible set are marked by the same color.

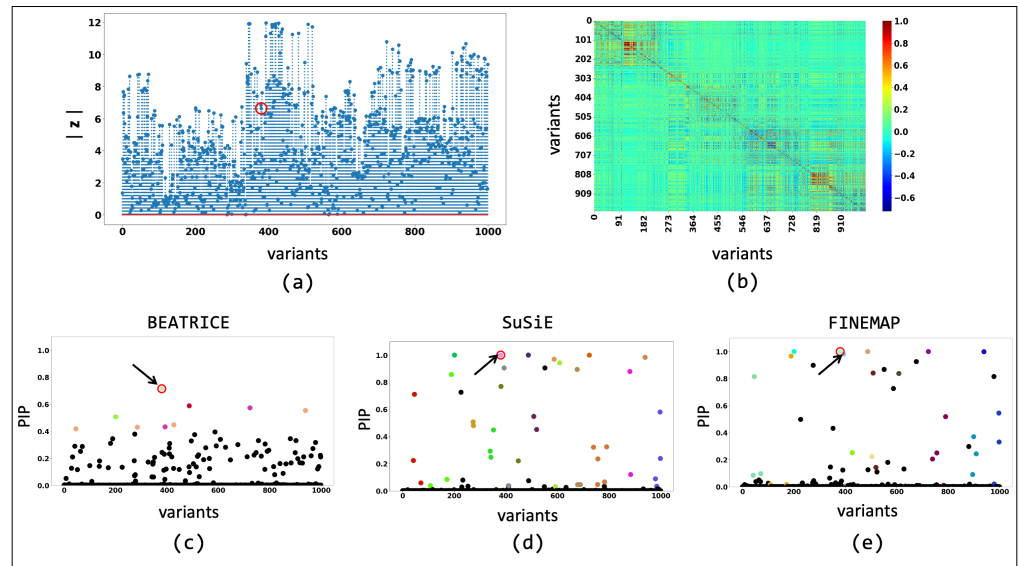
fine-mapping. Going one step further, a recent direction in fine-mapping is to aggregate data across multiple studies to identify causal variants [25]. Here, different LD matrices across studies helps to refine the fine-mapping results. BEATRICE can be applied in this context as well simply by modifying Eq. (12) as follows:

$$\begin{aligned} \mathcal{L}(\phi) = & -\frac{1}{SL} \sum_{s=1}^S \sum_{l=1}^L \log \left( N \left( \mathbf{z}_s; 0, \boldsymbol{\Sigma}_{X_s} + \boldsymbol{\Sigma}_{X_s} \left( n\sigma^2 \boldsymbol{\Sigma}_C^l(\phi) \right) \boldsymbol{\Sigma}_{X_s} \right) \right) \\ & + \sum_i \mathbf{p}_i \log \left( \frac{\mathbf{p}_i}{p_0} \right) + (1 - \mathbf{p}_i) \log \left( \frac{1 - \mathbf{p}_i}{1 - p_0} \right) \end{aligned} \quad (17)$$

where  $s$  denotes each separate study,  $S$  is the total number of studies in the analysis, and  $\mathbf{z}_s, \boldsymbol{\Sigma}_{X_s}$  are the summary statistics for each study.

## 7 Code Availability

We have compiled the code for BEATRICE and its dependencies into a docker image, which can be found at <https://github.com/sayangsep/Beatrice-Finemapping>. We have also provided installation instructions and a detailed description of the usage. The compact packaging will allow any user to directly download and run BEATRICE on their data. Namely, all the user must specify are a directory path to the summary statistics (i.e., z-scores), the LD matrix, and the number of subjects. Fig. 8 shows the outputs generated by BEATRICE. The results are output in (1) a PDF document that displays the PIPs and corresponding credible sets, (2) a table with PIPs, (3) a text file with credible sets, and (4) a text file with the conditional inclusion probability of the



**Fig 6.** The fine-mapping performance of BEATRICE , SuSiE, and FINEMAP at a noise setting of  $\{d = 1, \omega^2 = 0.2, p = 0.1\}$ . (a) The absolute z-score of each variant as obtained from GWAS. (b) Pairwise correlation between the variants. (c), (d), and (e) are the posterior inclusion probabilities of each variant as identified by BEATRICE , SuSiE, and FINEMAP, respectively. The red circle marked by an arrow shows the location of the causal variant. We have further color-coded the variants based on their assignment to credible sets. The non-black markers represent the variants assigned to a credible set. Additionally, the variants in a credible set are marked by the same color.

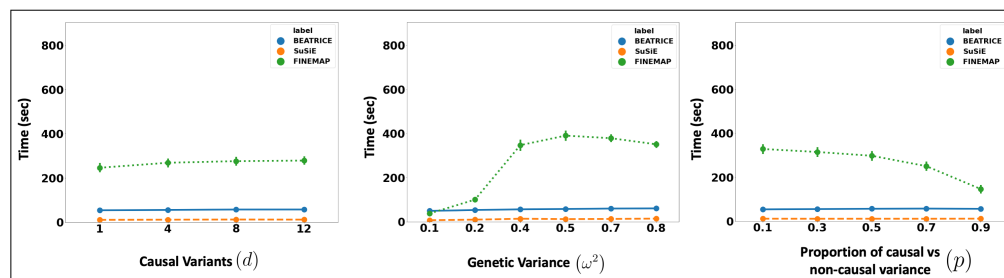
variants within the credible sets. The user can also generate the neural network losses describe in Eq. (12) by adding a flag to the run command.

## 8 Conclusion

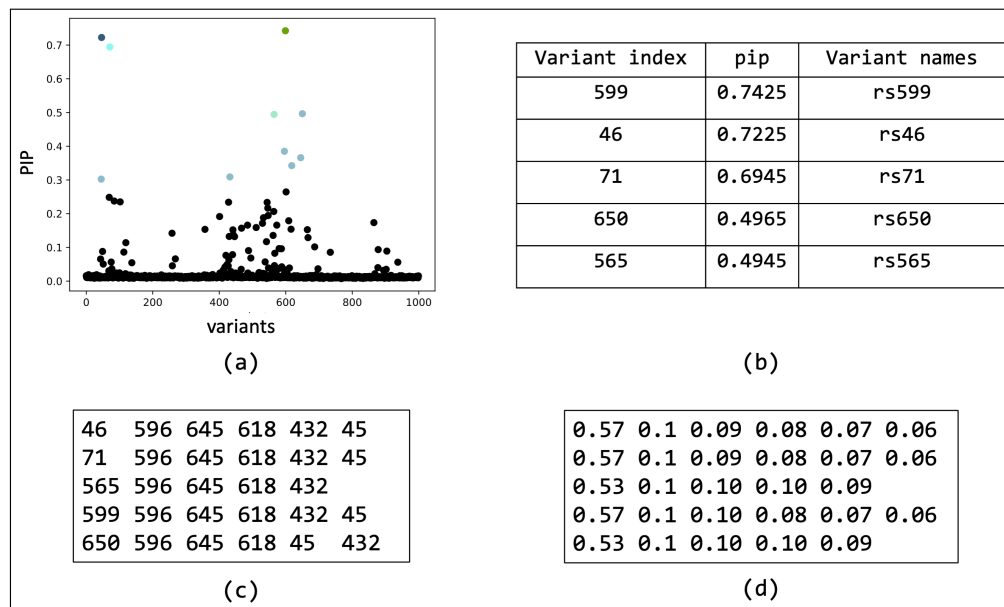
We present BEATRICE, a novel Bayesian framework for fine-mapping that identifies potentially causal variants within GWAS risk loci through the shared LD structure. Using a variational approach, we approximate the posterior probability of the causal location(s) via a binary concrete distribution. We leverage the unique properties of binary concrete random variables to build an optimization algorithm that can successfully model variants with differing levels of association. Moreover, we introduce a new strategy to build a reduced set of causal configurations within the exponential search space that can be neatly folded into our optimization routine. This reduced set is used to approximate the PIPs and identify credible sets. In a detailed simulation study, we compared BEATRICE with two state-of-the-art baselines and demonstrated the advantages of BEATRICE under different noise settings. Finally, our model does not have any prior on the causal variants and is agnostic to the original GWAS study. Hence, BEATRICE is a powerful tool to refine the results of a GWAS or eQTL analysis. It is also flexible enough to accommodate a variety of experimental settings.

## Acknowledgments

This work was supported by the National Science Foundation CAREER Award 1845430 (PI: Venkataraman), the National Institutes of Health Awards R01-HD108790 (PI:



**Fig 7.** The runtime comparison of BEATRICE, SuSiE, and FINEMAP across all the simulation settings.



**Fig 8.** Overview of the outputs generated by BEATRICE. (a) The PIPs are displayed and color coded by their assignment to credible sets. (b) A table with the PIPs and the corresponding name of the variants. (c) A text file with the credible sets. Here each row represent a credible set and the entries are indices of the variants present in the credible set. The first column of each row represents the key index. (d) The conditional inclusion probability of each of the credible variants given the key variant. The calculations can be found in **Section S2** of the Supplements.

Venkataraman), U24-HG010263 (PI: Schatz) and U41-HG006620 (PI: Schatz).

386

## References

1. Uffelmann E, Huang QQ, Munung NS, de Vries J, Okada Y, Martin AR, et al. Genome-wide association studies. *Nature Reviews Methods Primers* 2021 1:1. 2021;1(1):1–21.
2. Visscher PM, Brown MA, McCarthy MI, Yang J. Five Years of GWAS Discovery. *The American Journal of Human Genetics*. 2012;90(1):7–24.



3. Brzyski D, Peterson CB, Sobczyk P, Candès EJ, Bogdan M, Sabatti C. Controlling the Rate of GWAS False Discoveries. *Genetics*. 2017;205:61–75. doi:10.1534/GENETICS.116.193987.
4. Maller JB, McVean G, Byrnes J, Vukcevic D, Palin K, Su Z, et al. Bayesian refinement of association signals for 14 loci in 3 common diseases. *Nature Genetics* 2012 44:12. 2012;44:1294–1301. doi:10.1038/ng.2435.
5. Spain SL, Barrett JC. Strategies for fine-mapping complex traits. *Human Molecular Genetics*. 2015;24(R1):R111–R119.
6. Schaid DJ, et al. From genome-wide associations to candidate causal variants by statistical fine-mapping. *Nature Reviews Genetics* 2018 19:8. 2018;19(8).
7. Hutchinson A, Watson H, Wallace C. Improving the coverage of credible sets in Bayesian genetic fine-mapping. *PLOS Computational Biology*. 2020;16:e1007829. doi:10.1371/JOURNAL.PCBI.1007829.
8. Cho S, et al. Elastic-net regularization approaches for genome-wide association studies of rheumatoid arthritis. *BMC proceedings*. 2009;3(Suppl 7).
9. Sabourin J, et al. Fine-Mapping Additive and Dominant SNP Effects Using Group-LASSO and Fractional Resample Model Averaging. *Genetic Epidemiology*. 2015;39(2).
10. Tibshirani R. Regression Shrinkage and Selection Via the Lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*. 1996;58:267–288. doi:10.1111/J.2517-6161.1996.TB02080.X.
11. Guan Y, et al. BAYESIAN VARIABLE SELECTION REGRESSION FOR GENOME-WIDE ASSOCIATION STUDIES AND OTHER LARGE-SCALE PROBLEMS. *The Annals of Applied Statistics*. 2011;5(3).
12. Carbonetto P, Stephens M. Scalable Variational Inference for Bayesian Variable Selection in Regression, and Its Accuracy in Genetic Association Studies. *Bayesian Analysis*. 2012;7(1):73–108.
13. Benner C, Spencer CCA, Havulinna AS, Salomaa V, Ripatti S, Pirinen M. FINEMAP: efficient variable selection using summary data from genome-wide association studies. *Bioinformatics*. 2016;32(10):1493–1501.
14. Chen W, Larrabee BR, Ovsyannikova IG, Kennedy RB, Haralambieva IH, Poland GA, et al. Fine mapping causal variants with an approximate bayesian method using marginal test statistics. *Genetics*. 2015;200(3):719–736.
15. Hormozdiari F, Kostem E, Kang EY, Pasaniuc B, Eskin E. Identifying causal variants at loci with multiple signals of association. *Genetics*. 2014;198(2):497–508.
16. Zou Y, Carbonetto P, Wang G, Stephens M. Fine-mapping from summary data with the “Sum of Single Effects” model. *PLOS Genetics*. 2022;18:e1010299. doi:10.1371/JOURNAL.PGEN.1010299.
17. Wang G, Sarkar A, Carbonetto P, Stephens M. A simple new approach to variable selection in regression, with application to genetic fine mapping. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2020;82(5):1273–1300.

18. Maddison CJ, et al. The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables. 5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings. 2016;.
19. Jang E, Gu S, Poole B. Categorical Reparameterization with Gumbel-Softmax. 5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings. 2016;.
20. Cheng W, Ramachandran S, Crawford L. Estimation of non-null SNP effect size distributions enables the detection of enriched genes underlying complex traits. *PLOS Genetics*. 2020;16:e1008855. doi:10.1371/JOURNAL.PGEN.1008855.
21. Servin B, Stephens M. Imputation-Based Analysis of Association Studies: Candidate Regions and Quantitative Traits. *PLOS Genetics*. 2007;3:e114. doi:10.1371/JOURNAL.PGEN.0030114.
22. Pirinen M, Donnelly P, Spencer CCA. EFFICIENT COMPUTATION WITH A LINEAR MIXED MODEL ON LARGE-SCALE DATA SETS WITH APPLICATIONS TO GENETIC STUDIES. *The Annals of Applied Statistics*. 2013;7(1):369–390.
23. Wakefield J. Bayes factors for Genome-wide association studies: Comparison with P-values. *Genetic Epidemiology*. 2009;33(1):79–86.
24. Kingma DP, Ba JL. Adam: A Method for Stochastic Optimization. 3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings. 2014;.
25. LaPierre N, Taraszka K, Huang H, He R, Hormozdiani F, Eskin E. Identifying causal variants by fine mapping across multiple studies. *PLOS Genetics*. 2021;17(9):e1009733.
26. Kingma DP, Welling M. Auto-Encoding Variational Bayes; 2013. Available from: <https://arxiv.org/abs/1312.6114>.
27. Hans C, Dobra A, West M. Shotgun Stochastic Search for “Large p” Regression. <https://doi.org/10.1198/016214507000000121>. 2012;102:507–516. doi:10.1198/016214507000000121.
28. Wang Z, et al. An autoimmune pleiotropic SNP modulates IRF5 alternative promoter usage through ZBTB3-mediated chromatin looping. *Nature Communications* 2023 14:1. 2023;14(1).
29. Trubetskoy V, et al. Mapping genomic loci implicates genes and synaptic biology in schizophrenia. *Nature* 2022 604:7906. 2022;604(7906).
30. Albiñana C, Zhu Z, Borbye-Lorenzen N, Boelt SG, Cohen AS, Skogstrand K, et al. Genetic correlates of vitamin D-binding protein and 25-hydroxyvitamin D in neonatal dried blood spots. *Nature Communications* 2023 14:1. 2023;14(1):1–16. doi:10.1038/s41467-023-36392-5.
31. Li Y, Dang X, Chen R, Wang J, Li S, Mitchell BL, et al. Cross-ancestry genome-wide association study and systems-level integrative analyses implicate new risk genes and therapeutic targets for depression. *medRxiv*. 2023; p. 2023.02.24.23286411. doi:10.1101/2023.02.24.23286411.

32. Davis J, Goadrich M. The relationship between precision-recall and ROC curves. *ACM International Conference Proceeding Series*. 2006;148:233–240. doi:10.1145/1143844.1143874.
33. Dimitromanolakis A, Xu J, Krol A, Briollais L. sim1000G: A user-friendly genetic variant simulator in R for unrelated individuals and family-based designs. *BMC Bioinformatics*. 2019;20:1–9. doi:10.1186/S12859-019-2611-1/TABLES/4.
34. Meyer HV, Birney E. PhenotypeSimulator: A comprehensive framework for simulating multi-trait, multi-locus genotype to phenotype relationships. *Bioinformatics*. 2018;34:2951–2956. doi:10.1093/BIOINFORMATICS/BTY197.