Accounting for uncertainty in residual variances improves calibration of the Sum of Single Effects model for small sample sizes

William R.P. Denault[1], Peter Carbonetto[1], Ruixi Li[2], The Alzheimer's Disease Functional Genomics Consortium, Gao Wang[2,3], and Matthew Stephens[1,4]

[1]Department of Human Genetics, University of Chicago, Chicago, IL, USA
[2]Center for Statistical Genetics, The Gertrude H. Sergievsky Center, Columbia University, New York, NY, USA
[3]Department of Neurology, Columbia University, New York, NY, USA
[4]Department of Statistics, University of Chicago, Chicago, IL, USA
[*]Correspondence: William R.P. Denault, wdenault@uchicago.edu, Matthew Stephens, mstephens@uchicago.edu

Abstract

The Sum of Single Effects (SuSiE) model, and a fitting procedure based on variational approximation, were recently introduced as a way to capture uncertainty in variable selection in multiple linear regression. This approach is particularly well-suited to cases where variables are highly correlated, and it has become quickly adopted for genetic fine-mapping studies. Here, we show that in small studies (< 50 samples) the original fitting procedure can produce substantially higher rates of false positive findings than it does in larger studies. We show that a simple alternative fitting procedure, which takes account of uncertainty in the residual variance, notably improves performance in small sample studies.

1. Introduction

[1] introduced the Sum of Single Effects (SuSiE) model as a way to capture uncertainty in variable selection in multiple linear regression. This work was motivated by settings like genetic fine-mapping, which often involves very highly correlated variables, where it is difficult to confidently identify the "relevant" variables (those with a non-zero regression coefficient). [1] point out that, even in such settings, it may be possible, and useful, to narrow down the potentially relevant variables. They suggest doing this by identifying sets of (highly correlated) variables that, with high probability, contain at least one variable with non-zero coefficient. [1] refer to such sets of variables as "Credible Sets" (CSs), and the ability of the SuSiE method to quickly produce reliable CSs is a key attraction. Since its introduction several papers have extended the SuSiE method in various ways [2–6].

Here we highlight a limitation of the original SuSiE method from [1]: it can produce badly miscalibrated CSs when applied to data sets with small sample sizes ($n \leq 50$). We show that this miscalibration occurs even in the simplest case, where exactly one variable is assumed to have a non-zero coefficient. While many fine-mapping studies are large, this miscalibration issue may nonetheless be important in smaller studies such as pilot experiments , complex experimental designs (e.g., dynamic eQTL from [7], $n = 19$, and more recently [8], $n = 53$), costly multi-modal omics experiments (e.g., multi-modal QTL mapping in influenza-infected cells [9], $n = 35$), or studies of rare cell types (e.g., eQTL mapping in microglia from the subventricular zone [10], $n = 47$).

We show that this miscalibration can be greatly improved by a simple modification of the procedures in [1] to account for uncertainty in the residual variance parameter. Specifically, whereas [1] estimate the residual variance by maximum likelihood, here we introduce methods to integrate it out under a diffuse prior, an idea that has its roots in the classical Bayesian literature [11–14]. This simple modification substantially improves CS calibration in simulations, and provides more realistic result when applying the SuSiE model to small-sample settings (Figure 1).

Accounting for uncertainty in the residual variance is not a new idea in the Bayesian statistics literature [11–14]. In this report we showcase that bring back this to extend the SuSiE model

2. The single effect regression, accounting for uncertainty in residual variance estimates

The SuSiE model is built on the simpler "Single Effect Regression" (SER) model, which assumes that *exactly one* of the variables in the regression has a non-zero effect on the outcome. The SER model can be written as:

$$\boldsymbol{y} = \boldsymbol{X}\boldsymbol{\beta} + \epsilon \tag{1}$$
$$\boldsymbol{\beta} = \boldsymbol{\gamma}\beta \tag{2}$$
$$\boldsymbol{\gamma} = Mult(1, \boldsymbol{\pi}) \tag{3}$$
$$\epsilon \sim N(0, \sigma^2) \tag{4}$$
$$\beta \sim N(0, \sigma_0^2\sigma^2). \tag{5}$$

Here $\boldsymbol{y}$ is the vector of observed outcomes, $\boldsymbol{X}$ is a matrix of explanatory variables, $\boldsymbol{\beta}$ is the vector containing the effects of these variables on the phenotype, and ϵ is Gaussian noise with variance σ^2. Note that $\boldsymbol{\gamma}$ is a vector in which exactly one entry is different from zero, the prior on which entry is non-zero being captured by $\boldsymbol{\pi}$ (a vector with non-negative entries that sum to one). The hyperparameter σ_0^2 determines the prior on the effect size relative to the residual variance. [1] estimate the residual variance σ^2 by maximum likelihood (or by maximizing the evidence lower bound when extending the SER to the SuSiE model). Given σ^2, σ_0^2 it is straightforward to compute the posterior distribution on $\boldsymbol{\beta}$ analytically, essentially by simply summing over all the possible values for γ [1, 12].

Here we augment the SER model above by placing a (conjugate) Inverse-Gamma prior on the residual variance:

$$\sigma^2 \sim \text{IG}(\kappa/2, \lambda/2) \tag{6}$$

This augmented SER is, in essence, the model considered by [12], and so we call it the SS-SER. Again, given $\kappa, \lambda, \sigma_0^2$ it is straightforward to compute the posterior distribution on $\boldsymbol{\beta}$ analytically. [12] show that this posterior distribution has a sensible limit as $\kappa, \lambda \to 0$, and we use that limiting posterior distribution here.

2.1. Extension to the SuSiE model

The SER model is very limited because it allows for only a single non-zero effect. [1] suggest allowing for multiple effects by summing multiple copies of the SER model, to yield the "sum of single effects" (SuSiE) model. They develop a fitting procedure, Iterative Bayesian Stepwise Selection (IBSS), which fits the SuSiE model by *iteratively* fitting SER models. They show that this IBSS algorithm, when used with their SER, corresponds to optimizing a variational approximation to the posterior distribution. In brief, here we use the same IBSS algorithm as [1], but replace their SER (which fixes σ at a point estimate) with the SS-SER (which integrates out σ); see online material for details. Unlike the original IBSS algorithm, this modified IBSS is not guaranteed to be exactly optimizing a variational approximation to the posterior distribution; however, our empirical results below show that in small sample studies it can produce better CS calibration, while maintaining the simple computations of the original algorithm.

Compares of the number of genes in which SuSiE with default SER reports

the size of each point reflects the number the number of gene regression (SER) components selected by the default SuSiE model (x-axis) versus the sum-of-squares (SS) optimized variant (y-axis), across multiple genes. Each point represents a gene, with size indicating the number of genes sharing that SER combination. The diagonal line indicates agreement between the two methods. Larger bubbles below the diagonal suggest that the SS-based method tends to select fewer components than the default.

3. Results

3.1. Numerical Simulations

We compared our new approach (*SS-SER*) with the original approach (*default SER*) on simulated data. For each simulation, we sampled genotype data by randomly selecting loci from the 1000

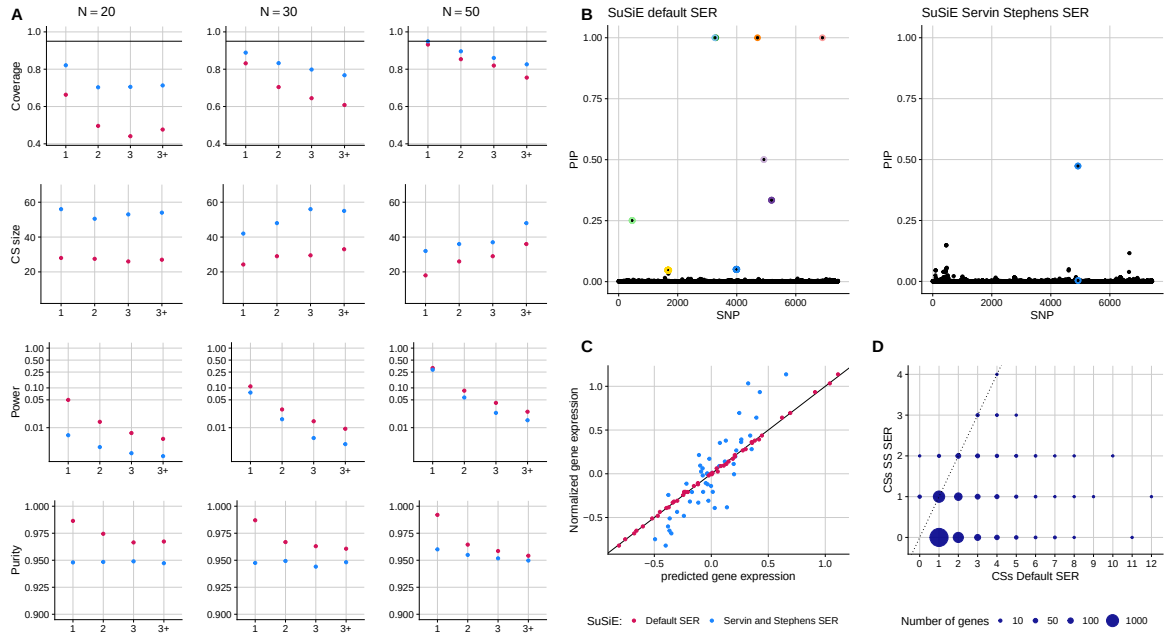

Fig. 1. Panel A: Comparison between SuSiE with the default SER (red) and SuSiE with the SS-SER (blue) in a simulation where the genotype accounts for 25% of the trait variance. The top row displays the observed empirical coverage of the two methods for N=20,30, and 50. The black line represents the target 95% coverage. Coverage was calculated as$1 - \frac{FP}{TP+FP}$ where FP and TP denote the number of false positives and true positives, respectively. The second row presents the observed mean CS size in the simulation. The third row presents the methods' power in the simulation. Note that the y-axis is displayed on a logarithmic scale. The fourth row presents the median CS purity in the simulation. **Panel B-C:** Illustration of severe miscalibration of SuSiE with default SER in fine-mapping *ADAM17* expression. These figures showcase a case of severe miscalibration in SuSiE when using the default SER, as observed in real data analysis, and contrast it with the more realistic fit of SuSiE with SS-SER using data from microglia in the subventricular zone from [10] , $n = 47$. **Panel D:** summary of the genome wide fine mapping of the RNA seq data of microglia in the subventricular zone from [10]. **Panel B** shows the estimated posterior inclusion probabilities (PIP) and credible sets (CS) for both SuSiE with the default SER and SS-SER. The x-axis represents the SNP index, while the y-axis denotes the corresponding PIP estimates. Colored circles indicate inferred CS (one color per CS). SuSiE with the default SER identifies nine CSs, whereas SuSiE with SS-SER reports only a single CS (also reported by SuSiE with default SER). **Panel C** compares the predicted ADAM17 gene expression under both models (red: default SER, blue: SS-SER) against normalized gene expression. The default SER exhibits severe overfitting, explaining more than 99.9% of the variance, while the SS-SER model provides a more realistic fit, capturing 68% of the observed variance. **Panel D:** comparison of the number of 95% credible sets (CSs) reported by SuSiE using the default SER prior (x-axis) versus the SS-optimized SER prior (y-axis), in genome-wide fine-mapping of RNA-seq data from microglia in the subventricular zone, as reported in [10]. Each point represents a distinct configuration of CS counts across genes; the size of the point reflects the number of genes exhibiting that configuration. For example, the large point at (1, 0) indicates that SuSiE with the default SER reports one CS in 2,634 genes, while SuSiE with SS SER reports none in those same genes. For visualization purposes, the point at (0,0) is removed; in 12,232 genes, both methods report zero CSs.

Genomes Project Phase 3 [15], including only unrelated individuals. Within the selected loci, we retained biallelic SNPs with observed minor allele frequencies $\geq 5\%$. From the remaining SNPs, we randomly selected 1 to 10 causal variants that collectively explain a fixed proportion of the phenotypic variance ($h^2 = 25\%, 30\%, 50\%, 75\%$). This operation was repeated using different sample sizes ($N = 20, 30, 50, 75, 100$). We ran both methods using an upper bound of $L = 10$ causal variants, and with the purity filtering step from [1], which removes "low purity" CSs (CSs that contain pairs of variables that are not highly correlated, correlation < 0.5). See supplementary material for details.

Our results (Figures 1A and 2) show that, for small sample size studies ($n \leq 50$), SuSiE with default SER has consistently poor CS coverage, substantially below the nominal 95% level. By comparison,

using the SS-SER substantially improves coverage (although not attaining 95%). For example, with $N = 20$ and $h^2 = 0.25$, the default SER CS coverage can be as low as 44.1%, increasing to 82.1% using the SS-SER (Figure 1A). The reason for this improved coverage is that SS-SER tends to produce larger CSs (Figure 1A; Figure 3) with lower purity (Figure 1A; Figure 4). Consequently more CSs are discarded by the purity filter, and so it also displays a corresponding drop in "power" (Figure 1A; Figure 5). It seems likely that this loss of power and localization are an inevitable consequence of improving the calibration.

To gauge the impact of estimating the number of causal variants, and the purity filtering step, we conducted an additional set of simpler simulations where a single variant affects the trait, and assessed CS coverage when (i) allowing SuSiE to estimate the number of causal effects (setting $L = 10$ as above), and (ii) running SuSiE with the correct number of causal effects ($L = 1$). We ran (i) and (ii) both with and without purity filtering. Note that fitting SuSiE with $L = 1$ does not involve any variational approximation, so (ii) removes any potential error in this approximation. We summarize the results of these simulations in Figures 6 and 7.

In these simpler simulations, the mis-calibration issue appears only with the smallest sample sizes ($n = 20, 30$) and smallest effect size $h^2 = 0.25, 0.3$, and only in the cases with purity filtering. Having SuSiE estimate the number of causal effects has only a minor impact on the calibration in all cases. These results may seem to suggest avoiding the purity filter altogether; however in practice this leads to very large numbers of "noninformative CSs" that contain very large numbers of variables. These noninformative CSs technically improve calibration because, by including almost every variable, they usually include a relevant variable by chance. However, outputting large numbers of uninformative CSs to improve calibration in this way seems undesirable in practice. Our recommendation for small studies is therefore to maintain the purity filter, use the SS-SER, and keep in mind that the resulting CSs may still be somewhat less reliable than for large studies.

3.2. Fine mapping Aracena et al. and Lopes et al. data

To compare results from the different methods on real data we fine-mapped the normalized gene expression data from two studies. The first dataset comes from [9], which examines the transcriptional response in primary macrophages (n = 35) before and after infection with influenza. The second dataset is from [10], which investigates expression QTL in microglia, a key cell type involved in brain aging.

For each gene we selected the SNPs within ±200kb of the gene body with MAF> 0.05 and fine-mapped these SNPs using both approaches. For the [9] study, across 13,680 fine-mapped genes, default SuSiE reported 1020 unique CSs, whereas SuSiE with SS-SER reported 389 CSs. Among the genes where SuSiE with SS-SER identified a single CS (381 out of 389), this CS always overlapped with a CS reported by default SuSiE. For the [10] study, across 17,132 fine-mapped genes, default SuSiE identified 6,918 CSs, while SuSiE with SS-SER reported 1,556 CSs. Again, nearly all genes with a single CS reported by SuSiE with SS-SER (844 out of 858) had overlapping CSs in default SuSiE. In the fine-mapping analysis of [10], SuSiE exhibited clear signs of miscalibration (see Figure 1), with 146 genes having four or more CSs, including up to 12 CSs for *GAL3ST2*. In contrast, SuSiE with SS-SER reported only one gene with four CSs.

4. Conclusion

We have highlighted a limitation of the SuSiE method: it often provides poorly calibrated CSs in small sample-size studies. We also provided a simple modification that uses the SS-SER to better account for uncertainty in the residual variance parameter, and improves credible set calibration. Our method retains the computational advantages of SuSiE and can be applied to summary statistics [16], making it practical for a wide range of fine-mapping studies. While our approach improves credible set coverage, some miscalibration remains in some settings, and further improving performance for small sample studies could be an area for future work.

References

[1] Wang, G., A. Sarkar and P. Carbonetto (2020). A simple new approach to variable selection in regression, with application to genetic fine mapping. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) 82*(5), 1273–1300. _eprint: https://rss.onlinelibrary.wiley.com/doi/pdf/10.1111/rssb.12388.

[2] Gao, B. and X. Zhou (2024, January). MESuSiE enables scalable and powerful multi-ancestry fine-mapping of causal variants in genome-wide association studies. *Nature Genetics 56*(1), 170–179. Publisher: Nature Publishing Group.

[3] Rossen, J., H. Shi and B. J. Strober (2024, May). MultiSuSiE improves multi-ancestry fine-mapping in All of Us whole-genome sequencing data. Pages: 2024.05.13.24307291.

[4] Yuan, K., R. J. Longchamps and A. F. Pardiñas (2023, January). Fine-mapping across diverse ancestries drives the discovery of putative causal variants underlying human complex traits and diseases. ISSN: 2328-4293 Pages: 2023.01.07.23284293.

[5] Zou, Y., P. Carbonetto and D. Xie (2023, June). Fast and flexible joint fine-mapping of multiple traits via the Sum of Single Effects model. Pages: 2023.04.14.536893 Section: New Results.

[6] Cui, R., R. A. Elzur and M. Kanai (2024, January). Improving fine-mapping by modeling infinitesimal effects. *Nature Genetics 56*(1), 162–169. Publisher: Nature Publishing Group.

[7] Strober, B., R. Elorbany and K. Rhodes (2019, June). Dynamic genetic regulation of gene expression during cellular differentiation. *Science (New York, N.Y.) 364*(6447), 1287–1290.

[8] Popp, J. M., K. Rhodes and R. Jangi (2024, December). Cell type and dynamic state govern genetic regulation of gene expression in heterogeneous differentiating cultures. *Cell Genomics 4*(12). Publisher: Elsevier.

[9] Aracena, K. A., Y.-L. Lin and K. Luo (2024, March). Epigenetic variation impacts individual differences in the transcriptional response to influenza infection. *Nature Genetics 56*(3), 408–419. Publisher: Nature Publishing Group.

[10] Lopes, K. d. P., G. J. L. Snijders and J. Humphrey (2022, January). Genetic analysis of the human microglial transcriptome across brain regions, aging and disease pathologies. *Nature Genetics 54*(1), 4–17. Publisher: Nature Publishing Group.

[11] Gönen, M., W. O. Johnson and Y. Lu (2005). The Bayesian Two-Sample t Test. *The American Statistician 59*(3), 252–257. Publisher: [American Statistical Association, Taylor & Francis, Ltd.].

[12] Servin, B. and M. Stephens (2007, July). Imputation-Based Analysis of Association Studies: Candidate Regions and Quantitative Traits. *PLOS Genetics 3*(7), e114.

[13] Liang, F., R. Paulo and G. Molina (2008, March). Mixtures of g Priors for Bayesian Variable Selection. *Journal of the American Statistical Association 103*(481), 410–423. Publisher: ASA Website _eprint: https://doi.org/10.1198/016214507000001337.

[14] Rouder, J. N., P. L. Speckman and D. Sun (2009, April). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review 16*(2), 225–237.

[15] Auton, A., G. R. Abecasis and D. M. Altshuler (2015, October). A global reference for human genetic variation. *Nature 526*(7571), 68–74. Publisher: Nature Publishing Group.

[16] Zou, Y., P. Carbonetto and G. Wang (2021, November). Fine-mapping from summary data with the "Sum of Single Effects" model. preprint, Genetics.

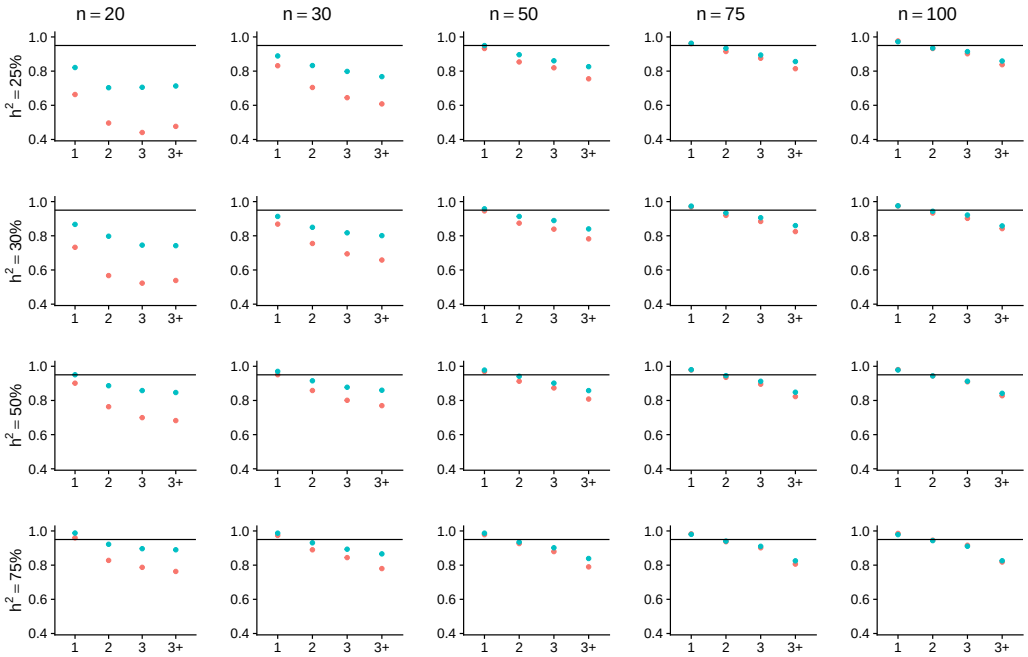

Fig. 2. Comparison of coverage for SuSiE with the default SER (orange) and SuSiE with the SS-SER (blue). The black line represents the target 95% coverage. Coverage was calculated as $1 - \frac{FP}{TP+FP}$ where FP and TP denote, respectively, the number of false positives, true positives, false negatives, and true negatives.

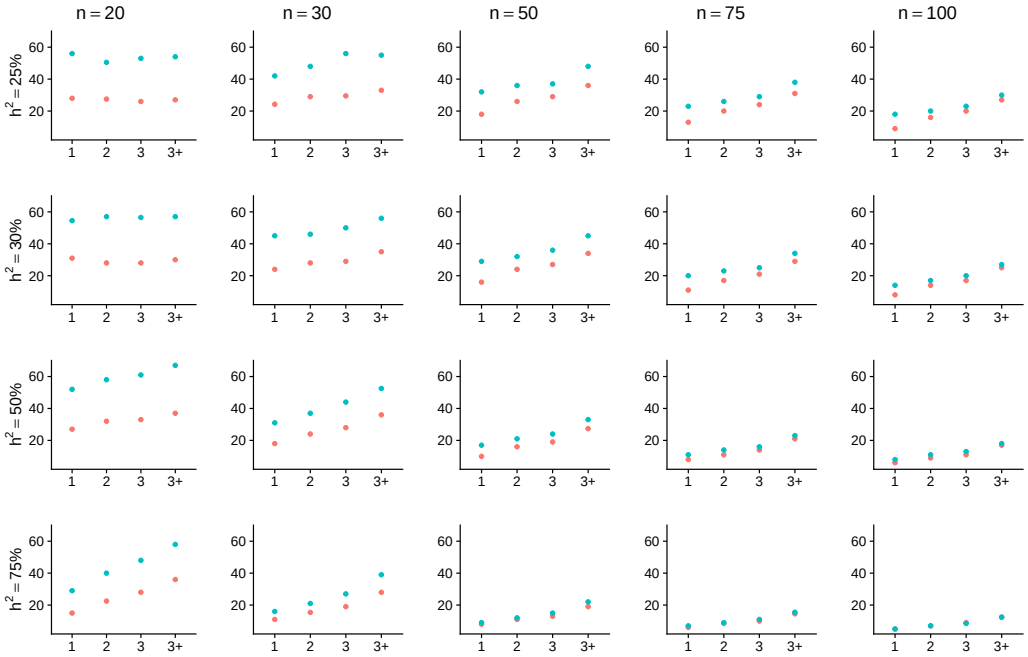

Fig. 3. Comparison of mean CS size for SuSiE with the default SER (orange) and SuSiE with the SS-SER (blue).

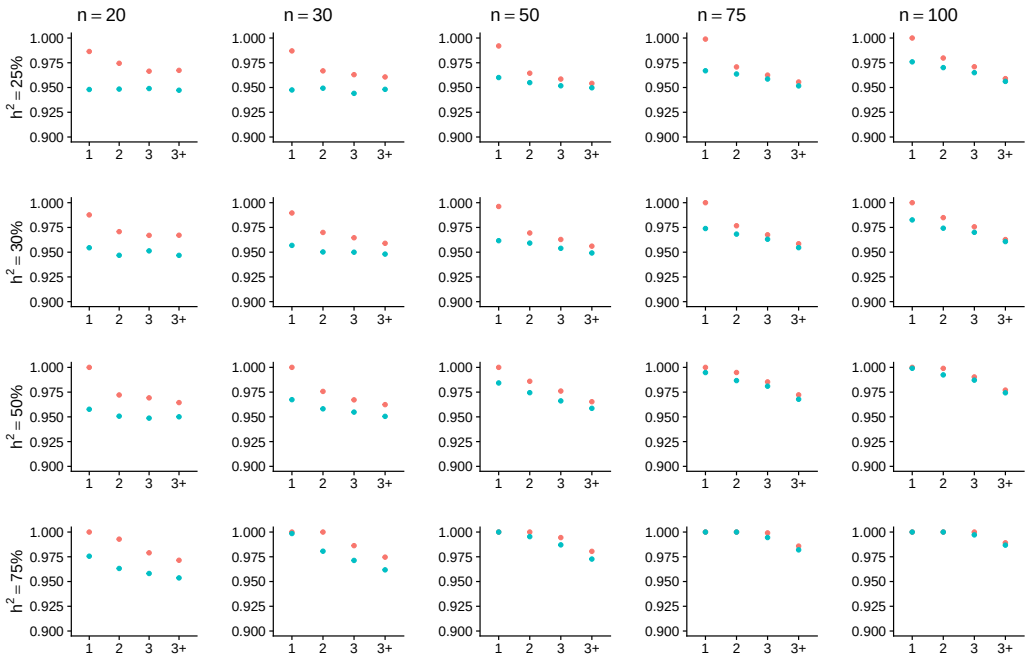

Fig. 4. Comparison of CS median purity for SuSiE with the default SER (orange) and SuSiE with the SS-SER (blue).

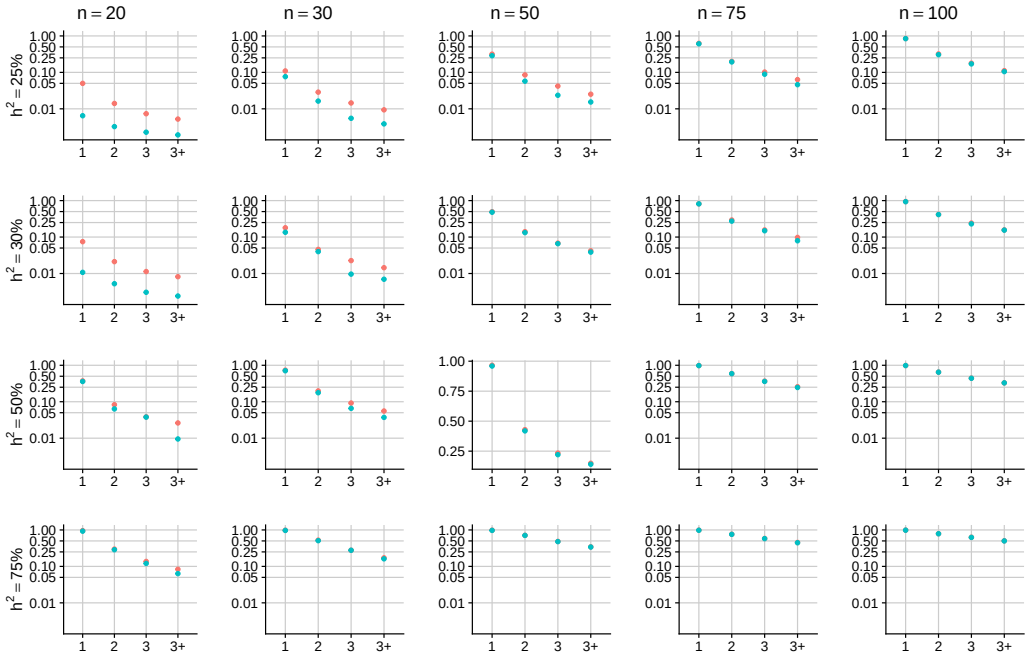

Fig. 5. Comparison of SuSiE power using the default Gaussian prior (orange) and the Servin and Stephens prior (blue). The y-axis is displayed on a logarithmic scale.

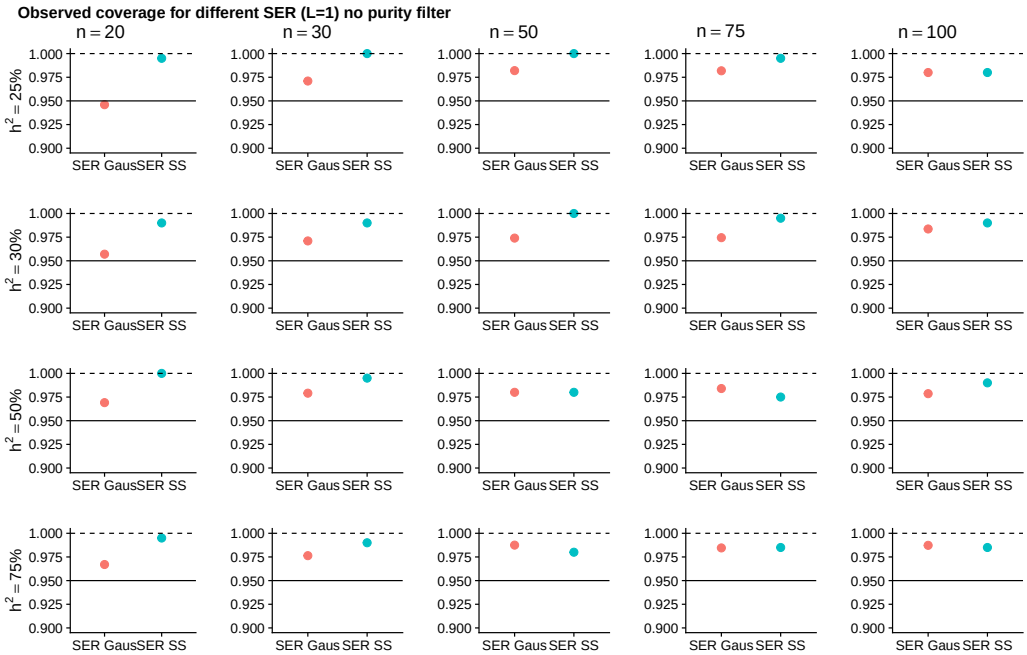

Fig. 6. Comparison of coverage between SuSiE using the default SER (orange) and SuSiE using the SS-SER (blue) under a single causal SNP scenario without purity filtering. The solid line denotes the target 95% coverage level, while the dashed line indicates 100% coverage. Simulations were conducted with the L parameter fixed to 1, a configuration under which SuSiE computes exact posteriors for both SERs.

Simulations

In all simulations, we randomly sample a genomic region from the 1000 Genomes Phase 3 dataset [15] and select N unrelated individuals within this locus. We retain only biallelic SNPs with a minor allele frequency (MAF) of at least 5% among the selected individuals.

Multiple causal SNPs scenario

For each simulation, we vary the number of individuals as $N = 20, 30, 50, 75$, and 100 and set the heritability, defined as the variance explained by the genotype, to $h^2 = 25\%, 30\%, 50\%$, and 75%.

For a given sample size N and heritability h^2, we randomly select N individuals as described above. We then randomly determine the number of causal SNPs, L, where L is drawn uniformly from $1, \ldots, 10$. The phenotype for individual i is generated as: $y_i = \mu_i + \epsilon_i$, where $\mu_i = \sum_{l=1}^{L} x_{il}$. Where x_{il} denotes the genotype at the lth causal SNP for individual i, and ϵ_i is a Gaussian noise term. To ensure that the *in-sample variance* explained by the genotype matches the specified h^2, we rescale the error term ϵ_i such that:

$$\frac{\sum_i (\mu_i - \bar{\mu})^2}{\sum_i (y_i - \bar{y})^2} = h^2, \tag{7}$$

where $\bar{\mu}$ and $\bar{y}$ are the empirical means of μ and y, respectively. This adjustment ensures the in-sample heritability is consistent with the target value, which may otherwise deviate when in small sample study. In these simulations, we set the upper bound for the number of effects to 10 when running SuSiE, regardless of the selected single-effect regression (SER) variant, and allow the method to infer the optimal L.

Single causal SNP scenario

Simulations for the single causal SNP scenario follow the same setup as in the multiple causal SNP scenario, except that we fix the number of causal SNPs at $L = 1$. For each combination of N and h^2, we evaluate the performance of SuSiE and its variants using the following configurations:

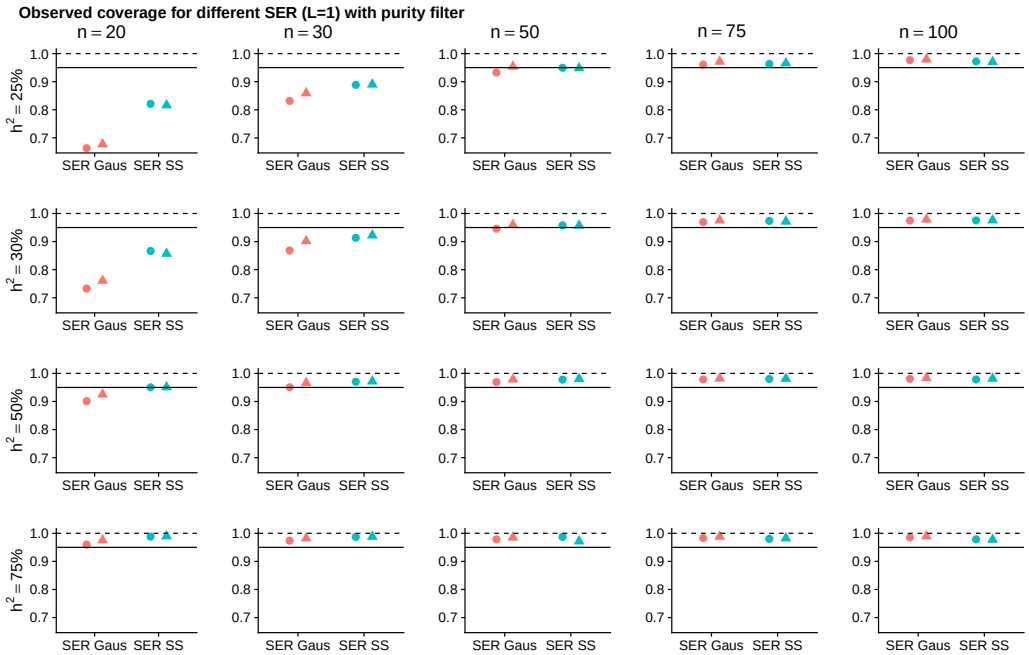

Fig. 7. Comparison of coverage between SuSiE using the default SER (orange) and SuSiE using the SS-SER (blue) under a single causal SNP scenario with purity filtering. The solid line denotes the target 95% coverage level, while the dashed line indicates 100% coverage. Round markers represent simulations where SuSiE dynamically estimates the number of causal effects (with an upper bound of 10), whereas triangle markers correspond to simulations with the LL parameter fixed to 1.

(a) SuSiE with default SER, setting the upper bound on the number of effects to 1, with and without purity filtering.

(b) SuSiE with default SER, setting the upper bound on the number of effects to 10, with and without purity filtering.

(c) SuSiE with Servin and Stephens SER (SS-SER), setting the upper bound on the number of effects to 1, with and without purity filtering.

(d) SuSiE with SS-SER, setting the upper bound on the number of effects to 10, with and without purity filtering.

Summaries of these simulations are presented in Figures 2 to 6. The scripts to replicate these simulations are available at `https://github.com/william-denault/susie_small_sample`

Methods

We present our fitting procedure for the SuSiE using SS-SER in the online material. Below we present the key quantities for estimating to estimate this model.

Modeling Uncertainty in σ^2

For simplicity, we present all derivations assuming κ and λ are strictly positive, ensuring a well-defined prior. However, the quantities can be computed with $\kappa = \lambda = 0$. Consider the following univariate

regression model where σ^2 is treated as unknown:

$$\boldsymbol{y} = \boldsymbol{x}\beta + \epsilon \tag{8}$$
$$\epsilon \sim N(0, \sigma^2) \tag{9}$$
$$\beta \sim N(0, \sigma_0^2\sigma^2) \tag{10}$$
$$\sigma^2 \sim \mathrm{IG}(\kappa/2, \lambda/2) \tag{11}$$

where $\boldsymbol{y}$ and $\boldsymbol{x}$ are in $\in \mathbb{R}^n$, β is an unknown parameter, ϵ is Gaussian noise, and σ_0^2 is a scaling factor on the effect size. $\mathrm{IG}(\kappa, \lambda)$ denotes the inverse gamma distribution with shape κ and scale λ.

The posterior distribution of β conditional on the other inputs and σ^2 is

$$\beta \mid \boldsymbol{x}, \boldsymbol{y}, \sigma^2, \sigma_0^2 \sim N(\mu_1, \sigma^2 s_1^2), \tag{12}$$

where

$$\begin{aligned} s_1^2 &= \frac{r_0^2}{\boldsymbol{x}^T\boldsymbol{x}}, \\ \mu_1 &= r_0^2\hat{b}, \\ r_0^2 &:= \frac{\sigma_0^2}{\sigma_0^2 + 1/(\boldsymbol{x}^T\boldsymbol{x})}. \end{aligned} \tag{13}$$

The posterior distribution of σ^2 is

$$\sigma^2 \mid \boldsymbol{x}, \boldsymbol{y}, \sigma_0^2 \sim \mathrm{IG}\left(\frac{\kappa_1}{2}, \frac{\lambda_1}{2}\right), \tag{14}$$

where

$$\begin{aligned} \kappa_1 &= \kappa + n, \\ \lambda_1 &= \lambda + \mathrm{RSS}(\sigma_0^2), \\ \mathrm{RSS}(\sigma_0^2) &= \boldsymbol{y}^T\boldsymbol{y}\left(1 - r_0^2\sigma_{xy}^2\right), \\ \sigma_{xy} &= \frac{\boldsymbol{x}^T\boldsymbol{y}}{\sqrt{(\boldsymbol{x}^T\boldsymbol{x})(\boldsymbol{y}^T\boldsymbol{y})}}. \end{aligned} \tag{15}$$

Thus, the posterior distribution of β is

$$\beta \mid \boldsymbol{x}, \boldsymbol{y}, \sigma_0^2 \sim t_{\kappa_1}\left(\mu_1, \frac{\lambda_1}{\kappa_1}s_1^2\right). \tag{16}$$

The Bayes factor (BF) for this model is given by

$$\mathrm{BF} = \sqrt{1 - r_0^2} \times \left(\frac{\lambda_1}{\lambda + \boldsymbol{y}^T\boldsymbol{y}}\right)^{-\kappa_1/2}. \tag{17}$$

Further details of these derivations are provided in the supplementary material.Further derivations are provided in the supplementary material.

Single-Effect Regression

Let's now consider the following single-effect regression (SER) model, a multiple linear regression model, where exactly one of the p variables affects the outcome:

$$\boldsymbol{y} = \boldsymbol{X}\boldsymbol{\beta} + \epsilon \tag{18}$$
$$\boldsymbol{\beta} = \boldsymbol{\gamma}\beta \tag{19}$$
$$\boldsymbol{\gamma} = Mult(1, \boldsymbol{\pi}) \tag{20}$$
$$\epsilon \sim N(0, \sigma^2) \tag{21}$$
$$\beta \sim N(0, \sigma_0^2\sigma^2). \tag{22}$$

where $\boldsymbol{X} \in \mathbb{R}^{n \times p}$, $\boldsymbol{y} \in \mathbb{R}^n$, and $\boldsymbol{\gamma}$ is a p vector in which exactly one entry is equal to 1 and the rest are equal to 0. We assume that columns of $\boldsymbol{X}$ and $\boldsymbol{y}$ are centered to have mean zero.

Under this model, the posterior distribution of $\boldsymbol{\beta}_j$ is a mixture of a point mass at zero and a t-distribution:

$$\boldsymbol{\beta}_j | \boldsymbol{X}, \boldsymbol{y}, \sigma_0^2 = (1 - \pi_{1j})\delta_0 + \pi_{1j} t_{\kappa_{1j}} \left(\mu_{1j}, \frac{\lambda_{1j}}{\kappa_{1j}} s_{1j}^2 \right) \tag{23}$$

where $s_{1j}^2, \mu_{1j}, \kappa_{1j}$, and λ_{1j} follow the same definitions as in the univariate case, substituting $\boldsymbol{x}$ with $\boldsymbol{X}_j$ in equation (8). Here, $y = \boldsymbol{X}_j \beta_j + \epsilon$. The mixture weights are given by:

$$\pi_{1j} = \frac{\pi_j \mathrm{BF}_j}{\sum_{j'=1}^{p} \pi_{j'} \mathrm{BF}_{j'}}, \tag{24}$$

with where BF_j is the Bayes factor for the univariate association between column j of $\boldsymbol{X}$ and $\boldsymbol{y}$ (see equation 17).