

Fast Model-Fitting of Bayesian Variable Selection Regression Using the Iterative Complex Factorization Algorithm

Quan Zhou^{*‡} and Yongtao Guan^{†§¶}

Abstract. Bayesian variable selection regression (BVSr) is able to jointly analyze genome-wide genetic datasets, but the slow computation via Markov chain Monte Carlo (MCMC) hampered its wide-spread usage. Here we present a novel iterative method to solve a special class of linear systems, which can increase the speed of the BVSr model-fitting tenfold. The iterative method hinges on the complex factorization of the sum of two matrices and the solution path resides in the complex domain (instead of the real domain). Compared to the Gauss-Seidel method, the complex factorization converges almost instantaneously and its error is several magnitude smaller than that of the Gauss-Seidel method. More importantly, the error is always within the pre-specified precision while the Gauss-Seidel method is not. For large problems with thousands of covariates, the complex factorization is 10–100 times faster than either the Gauss-Seidel method or the direct method via the Cholesky decomposition. In BVSr, one needs to repetitively solve large penalized regression systems whose design matrices only change slightly between adjacent MCMC steps. This slight change in design matrix enables the adaptation of the iterative complex factorization method. The computational innovation will facilitate the wide-spread use of BVSr in reanalyzing genome-wide association datasets.

Keywords: Cholesky decomposition, exchange algorithm, fastBVSr, Gauss-Seidel method, heritability.

1 Introduction

Bayesian variable selection regression (BVSr) can jointly analyze genome-wide genetic data to produce the posterior probability of association for each covariate and estimate hyperparameters such as heritability and the number of covariates that have nonzero effects (Guan and Stephens, 2011). But the slow computation due to model averaging

^{*}Department of Statistics, Rice University, 6100 Main St, Houston TX, 77005, qz9@rice.edu

[†]USDA/ARS Children's Nutrition Research Center, Department of Pediatrics, and Department of Molecular and Human Genetics of Baylor College of Medicine, 1100 Bates, Room 2070, Houston TX, 77030

[‡]Quan Zhou did most of his work when he was a PhD student in the program of Quantitative and Computational Biosciences, Baylor College of Medicine.

[§]Yongtao Guan's current affiliation: Department of Biostatistics and Bioinformatics, Duke University School of Medicine, Durham, NC, 27707, yongtao.guan@duke.edu

[¶]This work was supported by United States Department of Agriculture/Agriculture Research Service under contract number 6250-51000-057 and National Institutes of Health under award number R01HG008157.

using Markov chain Monte Carlo (MCMC) hampered its otherwise warranted widespread usage. Here we present a novel iterative method to solve a special class of linear systems, which can increase the speed of the BVS model-fitting tenfold.

1.1 Model and priors

We first briefly introduce the BVS method. Our model, prior specification and notation follow closely those of Guan and Stephens (2011). Consider the linear regression model

$$\mathbf{y} = \mu \mathbf{1} + \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim \text{MVN}(\mathbf{0}, \tau^{-1} \mathbf{I}), \quad (1)$$

where \mathbf{X} is an $n \times M$ column-centered matrix with $M \gg n$, \mathbf{I} denotes an identity matrix of proper dimension, \mathbf{y} and $\boldsymbol{\varepsilon}$ are n -vectors, $\boldsymbol{\beta}$ is an M -vector, and MVN stands for multivariate normal distribution. Let γ_j be an indicator of the j -th covariate having a nonzero effect and write $\boldsymbol{\gamma} = \{\gamma_1, \dots, \gamma_j, \dots, \gamma_M\}$. A spike-and-slab prior for β_j (the j -th component of $\boldsymbol{\beta}$) is specified below,

$$\begin{aligned} \gamma_j &| \pi \sim \text{Bernoulli}(\pi), \\ \beta_j &| \gamma_j = 0 \sim \delta_0, \\ \beta_j &| \gamma_j = 1, \sigma_\beta, \tau \sim N(0, \sigma_\beta^2/\tau), \end{aligned} \quad (2)$$

where π is the proportion of covariates that have non-zero effects and σ_β^2 is the variance of prior effect size (scaled by τ^{-1}). We will specify priors for both later. We use noninformative priors on the parameters μ and τ ,

$$\begin{aligned} \mu &| \tau \sim N(0, \sigma_\mu^2/\tau), & \sigma_\mu &\rightarrow \infty, \\ \tau &\sim \text{Gamma}(\kappa_1/2, \kappa_2/2), & \kappa_1, \kappa_2 &\rightarrow 0, \end{aligned} \quad (3)$$

where Gamma is in the shape-rate parameterization. As pointed out in Guan and Stephens (2011), prior (3) is equivalent to $P(\mu, \tau) \propto \tau^{-1/2}$, which is known as Jeffreys' prior (Ibrahim and Laud, 1991; O'Hagan and Forster, 2004). In practice, this means to use a diffuse prior for both μ and τ . Some may favor a simpler form $P(\mu, \tau) \propto 1/\tau$, which makes no practical difference (Berger et al., 2001; Liang et al., 2008).

Given $\boldsymbol{\gamma}$ and σ_β^2 , after integrating out $\boldsymbol{\beta}$, τ , μ and letting $\sigma_\mu \rightarrow \infty$, $\kappa_1, \kappa_2 \rightarrow 0$, the Bayes factor with reference to the null model can be computed in closed form,

$$\text{BF}(\boldsymbol{\gamma}, \sigma_\beta^2) = |\mathbf{I} + \sigma_\beta^2 \mathbf{X}_\gamma^t \mathbf{X}_\gamma|^{-1/2} \left(1 - \frac{\mathbf{y}^t \mathbf{X}_\gamma \hat{\boldsymbol{\beta}}}{\mathbf{y}^t \mathbf{y} - n\bar{y}^2} \right)^{-n/2}, \quad (4)$$

where \mathbf{X}_γ denotes the submatrix of \mathbf{X} with columns for which $\gamma_j = 1$, $|\cdot|$ denotes matrix determinant, and $\hat{\boldsymbol{\beta}}$ is the posterior mean for $\boldsymbol{\beta}$ given by

$$\hat{\boldsymbol{\beta}} = (\mathbf{X}_\gamma^t \mathbf{X}_\gamma + \sigma_\beta^{-2} \mathbf{I})^{-1} \mathbf{X}_\gamma^t \mathbf{y}. \quad (5)$$

The null-based Bayes factor $\text{BF}(\boldsymbol{\gamma}, \sigma_\beta^2)$ is proportional to the marginal likelihood $P(\mathbf{y} | \boldsymbol{\gamma}, \sigma_\beta^2)$, and evaluating BF is easier than evaluating the marginal likelihood due to cancellation of constants. The limiting prior (3) not only makes the Bayes factor expression

simpler (compared to that with finite value of σ_μ and positive values of κ_1 and κ_2), but also makes it invariant with respect to the shifting and scaling of \mathbf{y} .

We now discuss the prior specification for the two hyperparameters, π and σ_β^2 . To specify the prior for σ_β^2 , Guan and Stephens (2011) introduced a hyperparameter h (which stands for heritability) such that

$$h = \sigma_\beta^2 \sum_{j=1}^M \gamma_j s_j / (1 + \sigma_\beta^2 \sum_{j=1}^M \gamma_j s_j), \quad (6)$$

where s_j denotes the variance of the j -th covariate. Conditional on γ , specifying a prior on h will induce a prior on σ_β^2 , so henceforth we may write $\sigma_\beta^2(\gamma, h)$ to emphasize σ_β^2 is a function of h and γ . Since h is motivated by the narrow-sense heritability, its prior is easy to specify and we use

$$h \sim \text{Unif}(0, 1), \quad (7)$$

by default to reflect our lack of knowledge of heritability, although one can impose a strong prior by specifying a uniform distribution on a narrow support. A bonus of specifying the prior on σ_β^2 through h is that when $h \sim \text{Unif}(0, 1)$, the induced prior on σ_β^2 is heavy-tailed (Guan and Stephens, 2011).

Up till now, we follow faithfully the model and the prior specification of Guan and Stephens (2011). The prior on γ can be induced by the prior on π . Guan and Stephens (2011) specified the prior on π as uniform on its log scale, $\log \pi \sim \text{Unif}(\log \pi_{\min}, \log \pi_{\max})$, which is equivalent to $P(\pi) \propto 1/\pi$ for $\pi \in (\pi_{\min}, \pi_{\max})$, and sampled π, h, γ . But here we do something slightly different by integrating out π analytically. This is sensible because γ is very informative on π . Specifically, we integrate $P(\gamma, \pi)$ over $P(\pi)$ such that $P(\gamma) = \int P(\gamma | \pi) P(\pi) d\pi$ to obtain the marginal prior on γ ,

$$P(\gamma) = \frac{1}{\log(\pi_{\max}/\pi_{\min})} \int_{\pi_{\min}}^{\pi_{\max}} P(\gamma | \pi) / \pi d\pi, \quad (8)$$

where the finite integral is related to the truncated Beta distribution and can be evaluated conveniently. If π_{\min} goes to 0 and π_{\max} goes to 1 we have an improper prior $P(\pi) \propto 1/\pi$ and the marginal prior on γ becomes

$$P(\gamma) \propto \Gamma(|\gamma|) \Gamma(M + 1 - |\gamma|), \quad (9)$$

where we recall M is the total number of covariates, $|\gamma| = \sum \gamma_j$ is the number of selected covariates in the model, and Γ denotes the Gamma function. $P(\gamma)$ is always a proper probability distribution because it is defined on a finite set.

1.2 Posterior inference and computation

The joint posterior distribution of (γ, h) is given by

$$P(\gamma, h | \mathbf{y}) \propto P(\mathbf{y} | \gamma, h) P(\gamma) P(h). \quad (10)$$

The posterior inferences typically include computing the posterior inclusion probability $P(\gamma_j = 1 \mid \mathbf{y})$, which measures the strength of marginal association of the j -th covariate, the posterior distribution of the model size $|\gamma|$, and the posterior distribution of the heritability h , which measures the proportion of phenotypic variance explained by the selected models. We use MCMC to sample this joint posterior of (γ, h) . Our sampling scheme follows closely that of Guan and Stephens (2011). In each MCMC iteration, to evaluate (10) for a proposed parameter pair (γ', h') , we need to compute the marginal likelihood $P(\mathbf{y} \mid \gamma', h')$, which is proportional to (4). Two time-consuming calculations are the matrix determinant $|\mathbf{I} + \sigma_\beta^2 \mathbf{X}_\gamma^t \mathbf{X}_\gamma|$ and $\hat{\beta}$ defined in (5), both of which have cubic complexity (in $|\gamma|$). The computation of the determinant can be avoided by using a MCMC sampling trick which we will discuss later. The main focus of the paper is a novel algorithm to evaluate (5), which reduces its complexity from cubic to quadratic.

The rest of the paper is structured as follows. Section 2 introduces the iterative complex factorization (ICF) algorithm. Section 3 describes how to incorporate the ICF algorithm into BVSR. Both sections contain numerical examples, including a real dataset from genome-wide association studies. A short discussion concludes the paper.

2 The iterative complex factorization

In this section, we propose a novel algorithm for solving the following linear system

$$(\mathbf{X}^t \mathbf{X} + \Sigma^2) \hat{\beta} = \mathbf{z}, \quad (11)$$

where Σ is a diagonal matrix with positive (but not necessarily identical) entries on the diagonal, and \mathbf{X} is an $n \times p$ matrix. Clearly (5) is a special case of (11). In the context of BVSR, \mathbf{X} should be understood as \mathbf{X}_γ and $p = |\gamma|$. We assume $p < n$, where the sample size n ranges from several hundreds to tens of thousands. The computational advancement we will introduce, however, can be applied to scenarios where $p > n$ (see discussion in Section 4).

Note that $\hat{\beta}$ is the familiar ridge regression estimator (Draper and Van Nostrand, 1979). It may appear that an algorithm designed for solving ridge regression can be borrowed to solve (11). But a unique feature of BVSR is that in each iteration of MCMC, the design matrix \mathbf{X} usually changes only by one or a few columns. Thus, $\mathbf{X}^t \mathbf{X}$ and its Cholesky decomposition can be obtained conveniently (details will follow). This unique feature allows us to design a much more efficient algorithm.

2.1 Existing methods

In Guan and Stephens (2011), the linear system (11) was solved using the Cholesky decomposition of $\mathbf{X}^t \mathbf{X} + \Sigma^2$, which requires $p^3/3$ flops (Trefethen and Bau III, 1997, Lec. 23). Although computing $\mathbf{X}^t \mathbf{X}$ from \mathbf{X} requires $O(np^2)$ flops, when \mathbf{X} only changes by a few columns, the majority of the entries in $\mathbf{X}^t \mathbf{X}$ do not change and updating $\mathbf{X}^t \mathbf{X}$ only requires $O(np)$ flops.

Iterative methods sometimes can be used to reduce the computational time. Define

$$\mathbf{A} = \mathbf{X}^t \mathbf{X} + \mathbf{\Sigma}^2 = \mathbf{L} + \mathbf{D} + \mathbf{U}, \quad (12)$$

where \mathbf{L} (\mathbf{U}) is the strictly lower (upper) triangular component and \mathbf{D} contains only the diagonals. Then three popular iterative procedures can be summarized as follows:

$$\text{Jacobi method:} \quad \hat{\boldsymbol{\beta}}^{(k+1)} = \mathbf{D}^{-1} \left[-(\mathbf{L} + \mathbf{U})\hat{\boldsymbol{\beta}}^{(k)} + \mathbf{z} \right];$$

$$\text{Gauss-Seidel method:} \quad \hat{\boldsymbol{\beta}}^{(k+1)} = (\mathbf{D} + \mathbf{L})^{-1} \left(-\mathbf{U}\hat{\boldsymbol{\beta}}^{(k)} + \mathbf{z} \right);$$

$$\text{successive over-relaxation:} \quad \hat{\boldsymbol{\beta}}^{(k+1)} = (\mathbf{D} + \omega \mathbf{L})^{-1} \left[-(\omega \mathbf{U} - (1 - \omega) \mathbf{D})\hat{\boldsymbol{\beta}}^{(k)} + \omega \mathbf{z} \right].$$

The successive over-relaxation (SOR) method is a generalization of the Gauss-Seidel method, where ω is called the relaxation parameter. When \mathbf{A} is positive definite, the Gauss-Seidel method always converges and the SOR method converges for $\omega \in (0, 2)$ (Golub and Van Loan, 2012, Chap. 10.1.2). For all three iterative methods, each iteration requires $2p^2$ flops. Thus, whether an iterative method is more efficient than the Cholesky decomposition depends on how many iterations it takes to converge. Another notable class of iterative methods is called Krylov subspace methods. Two famous examples are the steepest descent and the conjugate gradient (Trefethen and Bau III, 1997, Lec. 38).

In principle, all methods developed to solve ridge regression can be used here to solve (11), as we alluded to earlier. For example, the methods of Eldén (1977), Lawson and Hanson (1995) and Turlach (2006) were developed for solving (11) with fixed \mathbf{X} but changing $\mathbf{\Sigma}$, while Hawkins and Yin (2002) devised a method for fixed $\mathbf{\Sigma}$ but changing \mathbf{X} . Modern least square solvers (Rokhlin and Tygert, 2008; Avron et al., 2010; Meng et al., 2014) typically considered the case where \mathbf{X} is extremely large, sparse, or ill-conditioned, and under such conditions, these least square solvers outperform solving directly via Cholesky decomposition; see Meng et al. (2014, Sec. 5) for how to apply least square solvers to ridge regression. But all methods quoted above are less effective (or not applicable) in the context of BVSR, because they are not designed for solving (11) millions of times, each time with a slightly different \mathbf{X} and a different $\mathbf{\Sigma}$, and they did not take advantage of the feature of BVSR that the Cholesky decomposition of $\mathbf{X}^t \mathbf{X}$ can be obtained efficiently.

2.2 The iterative complex factorization (ICF) algorithm

Let $\mathbf{R}^t \mathbf{R}$ be the Cholesky decomposition of $\mathbf{X}^t \mathbf{X}$, where \mathbf{R} is upper triangular. In the context of BVSR, given the Cholesky decomposition of $\mathbf{X}^t \mathbf{X}$, the Cholesky decomposition of a new matrix $(\mathbf{X}')^t \mathbf{X}'$ can be obtained efficiently since \mathbf{X}' (the proposed new design matrix) differs from \mathbf{X} only by one or a few columns (see Section 3.1). So we consider solving the linear system

$$(\mathbf{R}^t \mathbf{R} + \mathbf{\Sigma}^2) \hat{\boldsymbol{\beta}} = \mathbf{z}. \quad (13)$$

Contrary to our intuition, the Cholesky decomposition of $\mathbf{R}^t \mathbf{R} + \mathbf{\Sigma}^2$ cannot be obtained efficiently. This was also noticed in Zhou et al. (2013, p. 6). We instead perform the

following decomposition

$$\begin{aligned} \mathbf{R}^t \mathbf{R} + \Sigma^2 &= \mathbf{H} - i\mathbf{S}, \\ \mathbf{H} &= (\mathbf{R}^t - i\Sigma)(\mathbf{R} + i\Sigma), \\ \mathbf{S} &= \mathbf{R}^t \Sigma - \Sigma \mathbf{R}, \end{aligned} \tag{14}$$

where i is the imaginary unit. Then we have the update

$$\mathbf{H}\hat{\boldsymbol{\beta}}^{(k+1)} = i\mathbf{S}\hat{\boldsymbol{\beta}}^{(k)} + \mathbf{z},$$

where the right-hand side is a complex vector, and $\hat{\boldsymbol{\beta}}^{(k+1)}$ can be obtained by a forward and a backward substitution involving two complex triangular matrices, $\mathbf{R}^t - i\Sigma$ and $\mathbf{R} + i\Sigma$. This update, however, diverges from time to time. Examining the details of the observed divergent cases reveals that the culprit is the imaginary part of $\hat{\boldsymbol{\beta}}^{(k)}$. Because the solution $\hat{\boldsymbol{\beta}}$ is real, discarding the imaginary part of $\hat{\boldsymbol{\beta}}^{(k)}$ at the end of each iteration will not affect the fixed point to which the iterative method converges. Denoting the real part of a complex entity (scalar or vector) by Re , the generalized update of our algorithm ICF (Iterative Complex Factorization) becomes

$$\hat{\boldsymbol{\beta}}^{(k+1)} = \text{Re}[(1 - \omega)\hat{\boldsymbol{\beta}}^{(k)} + \omega\mathbf{H}^{-1}(i\mathbf{S}\hat{\boldsymbol{\beta}}^{(k)} + \mathbf{z})], \tag{15}$$

where we have also introduced a relaxation parameter ω . Intuitively ω makes the update *lazy* to avoid over-shooting. The revised update converges almost instantaneously, in a few iterations, compared to a few dozen to a few hundred iterations with the Gauss-Seidel method. Each iteration of (15) requires $6p^2$ flops, thrice that required by a Gauss-Seidel iteration, because ICF operates complex (instead of real) matrices and vectors. The right-hand side of (15) can be reorganized as $\text{Re}[(1 - \omega)\mathbf{I} + i\omega\mathbf{H}^{-1}\mathbf{S}]\hat{\boldsymbol{\beta}}^{(k)} + \omega\mathbf{H}^{-1}\mathbf{z}$. Note that $\mathbf{H}^{-1}\mathbf{z}$ can be computed via forward and backward substitutions and does not require matrix inversion.

2.3 ICF converges to the right target

Proposition 1. Denote $\Psi(\omega) = \text{Re}[(1 - \omega)\mathbf{I} + i\omega\mathbf{H}^{-1}\mathbf{S}]$. Then ICF in (15) converges for any starting point if $\rho(\Psi(\omega)) < 1$ where ρ denotes the spectral radius.

Proof. The true solution $\hat{\boldsymbol{\beta}}$ satisfies $\mathbf{H}\hat{\boldsymbol{\beta}} = i\mathbf{S}\hat{\boldsymbol{\beta}} + \mathbf{z}$, which, after some algebra, gives $\hat{\boldsymbol{\beta}}^{(k+1)} - \hat{\boldsymbol{\beta}} = \Psi(\omega)(\hat{\boldsymbol{\beta}}^{(k)} - \hat{\boldsymbol{\beta}})$. The statement then follows faithfully from Theorem 10.1.1 of Golub and Van Loan (2012). \square

Theorem 2. There exists $\omega \in (0, 1]$ such that the ICF update detailed in (15) converges to the true solution.

Proof. Using the notation defined in (12) and (14), we have $\mathbf{H} = \mathbf{A} + i\mathbf{S}$. Since $[\text{Re}(\mathbf{H}^{-1}) + i\text{Im}(\mathbf{H}^{-1})](\mathbf{A} + i\mathbf{S}) = \mathbf{I}$, we have

$$\begin{cases} \text{Re}(\mathbf{H}^{-1})\mathbf{A} - \text{Im}(\mathbf{H}^{-1})\mathbf{S} = \mathbf{I}, \\ \text{Im}(\mathbf{H}^{-1})\mathbf{A} + \text{Re}(\mathbf{H}^{-1})\mathbf{S} = \mathbf{0}. \end{cases}$$

Using the fact that \mathbf{A} is invertible, we can solve the above system and obtain

$$\operatorname{Im}(\mathbf{H}^{-1}) = -\mathbf{A}^{-1}\mathbf{S}(\mathbf{A} + \mathbf{S}\mathbf{A}^{-1}\mathbf{S})^{-1}.$$

Both \mathbf{A} and \mathbf{S} are real matrices, and by the Woodbury identity we have

$$\Psi(\omega) = \mathbf{I} - \omega(\mathbf{I} + (\mathbf{A}^{-1}\mathbf{S})^2)^{-1}. \quad (16)$$

Immediately, for a fixed ω , the spectrum of $\Psi(\omega)$ is fully determined by the spectrum of $\mathbf{A}^{-1}\mathbf{S}$. Because \mathbf{S} is skew-symmetric, we have $-\mathbf{S} = \mathbf{S}^t$. Since \mathbf{A} is symmetric, so does $\mathbf{A}^{-1/2}$. Then we have $-\mathbf{A}^{-1/2}\mathbf{S}\mathbf{A}^{-1/2} = \mathbf{A}^{-1/2}\mathbf{S}^t\mathbf{A}^{-1/2}$, which means $\mathbf{A}^{-1/2}\mathbf{S}\mathbf{A}^{-1/2}$ is also skew-symmetric. Hence the eigenvalues of $\mathbf{A}^{-1}\mathbf{S}$, which are identical to those of $\mathbf{A}^{-1/2}\mathbf{S}\mathbf{A}^{-1/2}$, are conjugate pairs of pure imaginary numbers or zero. Let $\pm\eta i$ be such a pair with $\eta \geq 0$ and \mathbf{u} be the eigenvector corresponding to the eigenvalue ηi . We have $\mathbf{A}^{-1}\mathbf{S}\mathbf{u} = i\eta\mathbf{u}$, which implies $i\mathbf{u}^*\mathbf{S}\mathbf{u} = -\eta\mathbf{u}^*\mathbf{A}\mathbf{u}$, where \mathbf{u}^* denotes the conjugate transpose of \mathbf{u} . Since \mathbf{H} is a Hermitian positive definite matrix,

$$\mathbf{u}^*\mathbf{H}\mathbf{u} = \mathbf{u}^*(\mathbf{A} + i\mathbf{S})\mathbf{u} = (1 - \eta)\mathbf{u}^*\mathbf{A}\mathbf{u} > 0.$$

Since \mathbf{A} is also positive definite, we have $\mathbf{u}^*\mathbf{A}\mathbf{u} > 0$ and

$$0 \leq \eta < 1. \quad (17)$$

By (16), the eigenvalues of $\Psi(\omega)$ are identical pairs equal to $1 - \omega/(1 - \eta^2)$ (this includes the case $\eta = 0$). Proposition 1 just requires $|1 - \omega/(1 - \eta^2)| < 1$, or equivalently,

$$0 < \omega < 2(1 - \eta^2), \quad (18)$$

holds for all η . Thus the existence of ω follows from (17). \square

By Proposition 1, the spectral radius of $\Psi(\omega)$ determines how fast the error converges to zero. We provide a theory-guided procedure to adaptively tune the relaxation parameter ω , which relies on the following proposition that connects $\rho(\Psi(\omega))$ with ω and η .

Proposition 3. Denote the eigenvalues of $\mathbf{A}^{-1}\mathbf{S}$ as $\pm\eta i$ ($\eta \geq 0$), where \mathbf{A} and \mathbf{S} are defined in (12) and (14), and obtain η_{\min} and η_{\max} . Then the spectral radius of $\Psi(\omega)$ is

$$\rho(\Psi(\omega)) = \max \left\{ 1 - \frac{\omega}{1 - \eta_{\min}^2}, \frac{\omega}{1 - \eta_{\max}^2} - 1 \right\}, \quad (19)$$

and the optimal value for ω to achieve the minimum of $\rho(\Psi(\omega))$ is

$$\omega^* = 2 \left(\frac{1}{1 - \eta_{\min}^2} + \frac{1}{1 - \eta_{\max}^2} \right)^{-1}. \quad (20)$$

Proof. By (16) and (17), the smallest and the largest eigenvalue of $\Psi(\omega)$ are $1 - \omega/(1 - \eta_{\max}^2)$ and $1 - \omega/(1 - \eta_{\min}^2)$ with $\eta \in [0, 1)$. After adjusting for their signs, we obtain (19). The right-hand side of (19) is a function of ω , with the first item decreasing linearly in ω and the second item increasing linearly in ω . Hence the minimum of (19) is attained when the two quantities in the braces are equal, which proves (20). \square

Our adaptive strategy for choosing ω assumes that η_{min} is zero, which holds trivially for odd p by the property of skew-symmetric matrices. When p is even, our numerical studies found $\eta_{min} = 0$ is still a good approximation in practice (Supplementary S3 (Zhou and Guan, 2018)). We start the ICF update with $\omega^{(0)} = 1$. Suppose at the k -th iteration we can produce an estimate $\hat{\rho}^{(k)}$ for the spectral radius of $\Psi(\omega^{(k)})$. Then using (19), η_{max}^2 can be estimated by $\eta_{max}^2 \approx 1 - \omega^{(k)} / (1 + \hat{\rho}^{(k)})$. Plugging this into (20) we obtain an update for ω

$$\omega^{(k+1)} = \frac{2\omega^{(k)}}{1 + \omega^{(k)} + \hat{\rho}^{(k)}}.$$

Note that the update does not involve η_{max}^2 , but only $\omega^{(k)}$ and $\hat{\rho}^{(k)}$, and $\omega^{(k+1)}$ is a decreasing function of $\rho^{(k)}$. Finally, to estimate the spectral radius of $\rho^{(k)}$ we use

$$\hat{\rho}^{(k)} = \frac{\|\hat{\beta}^{(k)} - \hat{\beta}^{(k-1)}\|_2}{\|\hat{\beta}^{(k-1)} - \hat{\beta}^{(k-2)}\|_2},$$

where $\|\cdot\|_2$ denotes the ℓ^2 -norm. This update strategy borrows the idea of power iteration and is motivated by the observation that $\hat{\beta}^{(k)} - \hat{\beta}^{(k-1)} = \Psi(\omega)(\hat{\beta}^{(k-1)} - \hat{\beta}^{(k-2)})$ if ω were fixed (see also Proposition 1). The procedure works well in our numerical studies. To take care of the boundary conditions, we use $\omega^{(k)} = 1$ for $k = 0, 1, 2$.

2.4 ICF outperforms other methods

Our numerical comparison studies were based on real datasets of genome-wide association studies downloaded from dbGaP. The details of the datasets can be found in Section 3.3. Because the convergence of iterative methods is sensitive to the collinearity in the design matrix \mathbf{X} , our comparison studies used two datasets: the first one contains 20K SNPs sampled across the whole genome, and the other contains 20K SNPs that are physically adjacent. The first dataset has little or no collinearity (henceforth referred to as IND), and the second has collinearity due to linkage disequilibrium (henceforth referred to as DEP). The sample size is $n = 3000$ for both datasets. Our numerical studies compared different methods (detailed below) for their speed and accuracy of solving (11). Given p , for one experiment we sampled without replacement p columns from the IND (or DEP) dataset to obtain \mathbf{X} , simulated under the null $\mathbf{z} \sim \text{MVN}(\mathbf{0}, \mathbf{I})$, and solved (11) using different methods with $\Sigma = \text{diag}(4, \dots, 4)$, which corresponded to $\sigma_\beta = 0.5$. For each p we conducted 1000 independent experiments.

Our initial studies compared ICF with six other methods: the Cholesky decomposition (Chol), the Jacobi method, the Gauss-Seidel method (GS), the successive over-Relaxation (SOR) method, the steepest descent method, and the conjugate gradient (CG) method. We excluded the Jacobi method and the steepest descent method due to their poor performance. To ensure a fair comparison in the context of BVS, the starting point for Chol, GS, and SOR was $\mathbf{A} = \mathbf{X}^t \mathbf{X} + \sigma_\beta^{-2} \mathbf{I}$ being obtained, and the starting point for ICF was the upper triangular matrix \mathbf{R} such that $\mathbf{X}^t \mathbf{X} = \mathbf{R}^t \mathbf{R}$ being obtained. For GS, we tried the preconditioning method of Kohno et al. (1997), which is the most efficient among the methods surveyed in Niki et al. (2004), but we observed no improvement, most likely because \mathbf{A} is well-conditioned due to the regularization. For

SOR, we need choose a value for the relaxation parameter (denoted by ω_{SOR}), which is known to be very difficult. A solution was provided by Young (1954) (see also Yang and Matthias, 2007), but we observed that it did not apply when $p > 500$. After trial and error, we settled on using $\omega_{\text{SOR}} = 1.2$, which appeared to be optimal in our numerical studies. For all iterative methods, we started from $\hat{\beta}^{(0)} = \mathbf{0}$ and stopped if

$$\max_j |\hat{\beta}_j^{(k)} - \hat{\beta}_j^{(k-1)}| < 10^{-6}, \quad (21)$$

or the number of iterations exceeded 200. The computer code for different methods was written in C++ and was run in the same environment. The Cholesky decomposition was implemented using GSL (GNU Scientific Library) (Gough, 2009), and GS and SOR were implemented in a manner that accounted for the sparsity of the triangular matrices \mathbf{L} and \mathbf{U} to obtain maximum efficiency (Golub and Van Loan, 2012, Chap. 10.1.2). Lastly, we included in the comparison the LAPACK routine DGELS, the most widely used least square solver, as a baseline reference.

Dataset	p	Time (in seconds)						Convergence failures			
		DGELS	Chol	ICF	GS	SOR	CG	ICF	GS	SOR	CG
IND	50	7.3	0.034	0.019	0.016	0.025	0.054	0	0	0	0
	100	26	0.20	0.07	0.07	0.09	0.24	0	0	0	0
	200	115	1.38	0.30	0.34	0.39	1.15	0	1	1	0
	500	679	21.0	2.7	3.6	2.9	9.8	0	2	1	0
	1000	2948	161	14	36	25	60	0	11	7	0
DEP	50	7.3	0.035	0.020	0.029	0.031	0.056	0	8	6	0
	100	27	0.20	0.08	0.23	0.19	0.26	0	29	20	0
	200	116	1.39	0.45	2.13	1.69	1.34	0	125	93	0
	500	683	21.0	5.5	36.1	32.8	15.1	0	621	514	0
	1000	2984	160	36	183	180	133	0	979	951	0

Table 1: Wall time usage (in seconds) and numbers of convergence failures. The top half is for the IND dataset and the bottom half is for the DEP dataset. The statistics for each of the six methods were obtained from 1,000 independent repeats. DGELS: LAPACKE_dgels routine; Chol: Cholesky decomposition; ICF: iterative complex factorization; GS: Gauss-Seidel method; SOR: successive over-relaxation; CG: conjugate gradient. “Convergence failures” columns give the numbers of experiments that fail to converge within 200 iterations for the four iterative methods.

The results are summarized in Table 1. For the IND dataset, three iterative methods (ICF, GS and SOR) appear on par with each other for smaller p . For $p = 1000$, ICF outperforms the other two, and is 10 times faster than the Cholesky decomposition. On average it took ICF 5 iterations to converge, and ICF never failed to converge in all experiments. On the other hand, both GS and SOR failed to converge, at least once, for large p . For the DEP dataset, we note that ICF is the fastest among all the methods compared. For larger p ($p = 500, 1000$), ICF is 4–5 times faster than the Cholesky decomposition, and 5–6 times faster than the other three methods. Both GS and SOR had difficulty in converging within 200 iterations for large p . DGELS is always the slowest

since it assumes $\mathbf{X}^t \mathbf{X}$ is unknown and solves (11) by QR decomposition. Advanced least square solvers (Avron et al., 2010; Meng et al., 2014) have similar performance to DGELS and can only beat it by a small margin when \mathbf{X} is very large. We also tweaked simulation conditions to check whether the results were stable. For example, we tried to simulate \mathbf{z} under the alternative instead of the null, and to initialize $\hat{\beta}^{(0)}$ in different manners, such as using unpenalized linear estimates. The results remained essentially unchanged under different tweaks.




Figure 1: Comparison of the accuracy using the IND dataset. The three methods compared here are ICF (iterative complex factorization), CG (conjugate gradient), and SOR (successive over-relaxation). Each panel shows the distributions of the maximum (entry-wise) absolute error on \log_{10} scale.

Next, we compared the accuracy of different iterative methods, where the truth was obtained from the Cholesky decomposition. In this experiment we used the IND dataset and compared for $p = 500$ and $p = 1000$. Each method was repeated for 10,000 experiments. The maximum entry-wise deviation (denoted by d) was obtained for each method, each experiment, under each simulation condition. Only those converged experiments were included in the comparison. Figure 1 shows the distributions of $\log_{10} d$ for three methods. Clearly, ICF outperforms CG and SOR by a large margin. (GS is omitted since its accuracy is poorer than that of SOR in almost every experiment.) Because in real applications we are oblivious to the truth (or the truth is expensive to obtain), a method is more desirable if its deviation from the truth is within the pre-specified precision. Figure 1 shows that ICF always achieves the pre-specified precision (10^{-6}) while the other two methods do not. Moreover, the deviation of ICF is 2–3 orders of magnitude smaller than that of SOR, and 3–4 orders of magnitude smaller than that of the CG method. We repeated the experiments for the DEP dataset and made the same observations.

Figure 2 compares the number of iterations used to converge for different iterative methods for large p . ICF takes a few iterations to converge, SOR a dozen iterations, CG about 50 iterations. More numerical experiments can be found in Supplementary S2 where we investigated how p and σ_β affects convergence rates. We also compared the performance of the iterative methods when the design matrix \mathbf{X} is beyond the counts




Figure 2: Comparison of the convergence speed using the IND dataset. Each panel shows the distributions of the number of iterations needed to stop by rule (21).

of reference alleles, such as normal or log-normal distributed (Supplementary S1), and when X has a different degree of collinearity (Supplementary S2.2). ICF exhibits an overwhelming advantage in every scenario.

3 ICF dramatically increases the speed of BVSR

3.1 Incorporating ICF into BVSR

As we mentioned in the introduction, for BVSR the most time-consuming step in MCMC is the computation of (4) for every proposed (γ', h') (superscript ' denotes the proposed value). To successfully incorporate ICF into BVSR, we need to overcome two hurdles: avoiding the computation of the determinant term in (4) and efficiently obtaining the Cholesky decomposition $\mathbf{R}^t \mathbf{R} = \mathbf{X}_\gamma^t \mathbf{X}_\gamma$.

Computing matrix determinant has a cubic complexity. If this is not avoided, using ICF to compute $\hat{\beta}$ becomes pointless. When evaluating (4), we need to compute the determinant $|\mathbf{X}_\gamma^t \mathbf{X}_\gamma + \sigma_\beta^{-2}(\gamma, h) \mathbf{I}|$, which takes $O(|\gamma|^3)$ operations. Identifying that the determinant is a normalization constant, avoiding computation of the determinant becomes a well studied problem in the MCMC literature that deals with the so-called “doubly-intractable distributions,” i.e. distributions with two nested unknown normalization constants (Møller et al., 2006; Murray et al., 2012). Recall that we want to sample

$$\begin{aligned}
 P(\gamma, h \mid \mathbf{y}) &\propto P(\mathbf{y} \mid \gamma, h) P(\gamma, h) \\
 &\propto |\Omega(\gamma, h)|^{-1/2} \sigma_\beta^{-|\gamma|}(\gamma, h) \left(1 - \frac{\mathbf{y}^t \mathbf{X}_\gamma \Omega^{-1}(\gamma, h) \mathbf{X}_\gamma^t \mathbf{y}}{\mathbf{y}^t \mathbf{y} - n \bar{y}^2} \right)^{-n/2} P(\gamma, h), \quad (22)
 \end{aligned}$$

where $\mathbf{\Omega}(\gamma, h) = \mathbf{X}_\gamma^t \mathbf{X}_\gamma + \sigma_\beta^{-2}(h, \gamma) \mathbf{I}$, and it is the computation of $|\mathbf{\Omega}(\gamma, h)|$ that we want to avoid. For a naive Metropolis-Hastings algorithm, if the current state is (γ, h) and the proposed move is (γ', h') , we need to evaluate *the Hastings ratio*

$$\alpha = \frac{K(\gamma, h \mid \gamma', h')}{K(\gamma', h' \mid \gamma, h)} \frac{Z(\gamma', h')}{Z(\gamma, h)} \frac{L(\mathbf{y}, \gamma', h')}{L(\mathbf{y}, \gamma, h)} \frac{P(\gamma', h')}{P(\gamma, h)}, \quad (23)$$

where

$$\begin{aligned} Z(\gamma, h) &= |\mathbf{\Omega}(\gamma, h)|^{-1/2} \sigma_\beta^{-|\gamma|}(\gamma, h), \\ L(\mathbf{y}, \gamma, h) &= (\mathbf{y}^t \mathbf{y} - \mathbf{y}^t \mathbf{X}_\gamma \mathbf{\Omega}^{-1}(\gamma, h) \mathbf{X}_\gamma^t \mathbf{y} - n\bar{y}^2)^{-n/2}, \end{aligned} \quad (24)$$

and $K(\cdot' \mid \cdot)$ is the proposal distribution for proposing \cdot' from \cdot . The proposed move (γ', h') is accepted with probability $\min(1, \alpha)$, which is called *the Metropolis rule*. Apparently, both $|\mathbf{\Omega}(\gamma, h)|$ and $|\mathbf{\Omega}(\gamma', h')|$ need to be evaluated in a naive MCMC implementation.

To avoid computing $Z(\gamma, h)$, notice that

$$\frac{Z(\gamma', h')}{Z(\gamma, h)} = \frac{P(\tilde{\mathbf{y}} \mid \gamma', h')}{P(\tilde{\mathbf{y}} \mid \gamma, h)} \frac{L(\tilde{\mathbf{y}}, \gamma, h)}{L(\tilde{\mathbf{y}}, \gamma', h')} \quad (25)$$

holds for all $\tilde{\mathbf{y}}$, and if $\tilde{\mathbf{y}}$ is sampled from $P(\cdot \mid \gamma', h')$, the ratio $L(\tilde{\mathbf{y}}, \gamma, h)/L(\tilde{\mathbf{y}}, \gamma', h')$ becomes a one-sample importance-sampling estimate of $Z(\gamma', h')/Z(\gamma, h)$. Hence we can plug in $L(\tilde{\mathbf{y}}, \gamma, h)/L(\tilde{\mathbf{y}}, \gamma', h')$ to replace $Z(\gamma', h')/Z(\gamma, h)$ and compute the Hastings ratio by

$$\alpha(\gamma, h, \gamma', h', \tilde{\mathbf{y}}) = \frac{K(\gamma, h \mid \gamma', h')}{K(\gamma', h' \mid \gamma, h)} \frac{L(\tilde{\mathbf{y}}, \gamma, h)}{L(\tilde{\mathbf{y}}, \gamma', h')} \frac{L(\mathbf{y}, \gamma', h')}{L(\mathbf{y}, \gamma, h)} \frac{P(\gamma', h')}{P(\gamma, h)}. \quad (26)$$

This is the exchange algorithm of Murray et al. (2012), which can be viewed as a special case of the pseudo-marginal method (Andrieu and Roberts, 2009). In the Appendix we prove that the stationary distribution of this MCMC is the desired posterior distribution $P(\gamma, h \mid \mathbf{y})$. Since the Bayes factor (4) is invariant to the scaling and shifting of \mathbf{y} , we can simply assume $\mu = 0$ and $\tau = 1$ when sampling $\tilde{\mathbf{y}}$.

To tackle the second difficulty, notice that in each MCMC iteration, usually only one or a few entries of γ are flipped between 0 and 1. Such proposals are often called “local” proposals, and Markov chains using local proposals usually have high acceptance ratios, which indicate the chains are “mixing” well (Guan and Krone, 2007). When one covariate is added into the model, the new Cholesky decomposition, $(\mathbf{R}')^t \mathbf{R}'$, can be obtained from the previous decomposition $\mathbf{R}^t \mathbf{R}$ by a forward substitution. When one covariate is deleted from the model, we simply delete the corresponding column from \mathbf{R} and introduce new zeros by Givens rotation (Golub and Van Loan, 2012, Chap. 5.1), which requires $\approx 3k^2$ flops if the $(|\gamma| - k)$ -th predictor is removed. We provide a toy example in Supplementary to demonstrate the update on the Cholesky decomposition.

3.2 Numerical examples

We developed a software package, fastBVSR, to fit the BVSR model by incorporating ICF and the exchange algorithm into the MCMC procedure. The algorithm is




Figure 3: Heritability estimates. The left panel is the result of fastBVSr, obtained using 20,000 burn-in steps and 100,000 sampling steps, and the right panel is the result of GCTA. In each iteration of BVSr, the heritability is estimated by computing the proportion of explained variance of \mathbf{y} using the sampled parameter values. The grey bars represent 95% credible intervals for fastBVSr and ± 2 standard error for GCTA.

summarized in Supplementary S4. The software is written in C++ and available at <http://www.haploTYPE.org/software.html>. In fastBVSr, we also implemented Rao-Blackwellization, as described in Guan and Stephens (2011), to reduce the variance of the estimates for γ and β . By default, Rao-Blackwellization is done every 1000 iterations.

To check the performance of fastBVSr, we performed simulation studies based on a real dataset described in Yang et al. (2010) (henceforth the Height dataset). The Height dataset contains 3,925 subjects and 294,831 common SNPs (minor allele frequency $\geq 5\%$) after routine quality control (c.f. Xu and Guan, 2014). We sampled 10,000 SNPs across the genome to perform simulation studies. Our aim was to check whether fastBVSr can reliably estimate the heritability in the simulated phenotypes. To simulate phenotypes of different heritability, we randomly selected 200 causal SNPs (out of 10,000) to obtain γ . For each selected SNP we drew its effect size from the standard normal distribution to obtain β_γ (the subvector of β that contains nonzero entries). Then we simulated the standard normal error term to obtain ϵ , and scaled the simulated effect sizes simultaneously using λ such that $\mathbf{y} = \lambda \mathbf{X}_\gamma \beta_\gamma + \epsilon$ and the heritability, $1 - \text{Var}(\epsilon)/\text{Var}(\mathbf{y})$, was h (taking values in 0.01, 0.02, \dots , 0.99). This was the same procedure as that of Guan and Stephens (2011). For each simulated phenotype, we ran fastBVSr to obtain the posterior estimate of heritability. We compared fastBVSr with GCTA (Yang et al., 2011), which is a software package to estimate heritability and do prediction using the linear mixed model. For heritability estimation, GCTA has been shown to be unbiased and accurate in a wide range of settings (Yang et al., 2010; Lee et al., 2011). The result is shown in Figure 3. Both fastBVSr and GCTA can estimate the heritability accurately; the mean absolute error is 0.014 for fastBVSr and 0.029 for

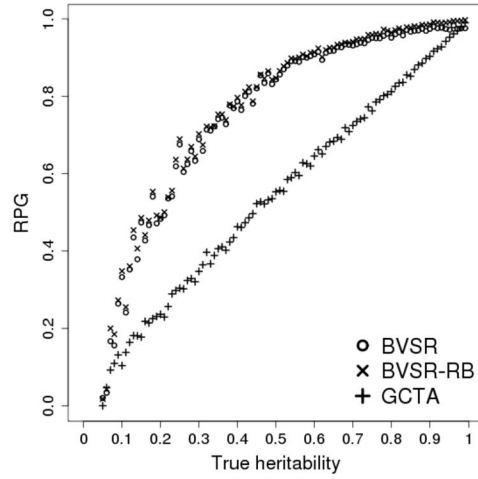


Figure 4: Relative prediction gain (RPG) of fastBVSR and GCTA. BVSR stands for the RPG of the crude posterior mean estimates from fastBVSR; BVSR-RB represents the RPG of the Rao-Blackwellized estimates from fastBVSR; for GCTA, RPG is computed using the BLUPs (best linear unbiased predictors) from linear mixed model.

GCTA. Noticeably, fastBVSR has a smaller variance but a slight bias when the true heritability is large.

Next we compare the predictive performance of fastBVSR and GCTA. We first define the mean squared error as a function of $\hat{\beta}$,

$$\text{MSE}(\hat{\beta}) = \frac{1}{n} \|\mathbf{X}\beta - \mathbf{X}\hat{\beta}\|_2^2.$$

Then following Guan and Stephens (2011) we define relative prediction gain (RPG) to measure the predictive performance of $\hat{\beta}$

$$\text{RPG} = \frac{\text{MSE}(\mathbf{0}) - \text{MSE}(\hat{\beta})}{\text{MSE}(\mathbf{0}) - \text{MSE}(\beta)}.$$

The advantage of RPG is that the scaling of \mathbf{y} does not contribute to RPG so that simulations with different heritability can be compared fairly. Clearly when $\hat{\beta} = \beta$, $\text{RPG} = 1$; when $\hat{\beta} = \mathbf{0}$, $\text{RPG} = 0$. Figure 4 shows that fastBVSR has much better predictive performance than GCTA, which reflects the advantage of BVSR over the linear mixed model. This advantage owes to the model averaging used in BVSR (Raftery et al., 1997; Broman and Speed, 2002). Besides, note that BVSR with Rao-Blackwellization performs slightly better than BVSR alone.

Lastly we check the calibration of the posterior inclusion probability (PIP), $P(\gamma_j = 1 \mid \mathbf{y})$. We pool the results of the 99 simulation sets with different heritability (from 0.01 to 0.99 at increment of 0.01). In each set there are 10,000 estimated PIPs, one




Figure 5: Calibration of the posterior inclusion probability (PIP). In the left panel PIP is estimated directly from MCMC samples, and in the right panel PIP is estimated using Rao-Blackwellization to reduce variance. All the PIP estimates are pooled into 20 bins (see the main text for details). X-axis denotes the mean PIP of each bin and the Y-axis denotes the fraction of true positives in each bin. The grey bars represent ± 2 standard error of the estimated proportion of true positives in each bin. The grey line is $y = x$.

for each SNP, and in total there are 990,000 PIP estimates. We group these PIPs into 20 bins, $[0.05 \times (i - 1), 0.05 \times i)$ for $i = 1, \dots, 20$, and for each bin we compute the fraction of true positives. If the PIP estimates are well calibrated, we expect that in each bin, the average of the PIPs roughly equals the fraction of true positives. Figure 5 shows that the PIP estimated by BVSR is conservative, which agrees with the previous study (Guan and Stephens, 2011), and the PIP estimated by Rao-Blackwellization is well calibrated.

3.3 Analysis of real GWAS datasets

We applied fastBVSR to analyze GWAS datasets concerning intraocular pressure (IOP). We applied and downloaded two GWAS datasets from the database of Genotypes and Phenotypes (dbGaP), one for glaucoma (dbGaP accession number: phs000238.v1.p1) and the other for intraocular hypertension (dbGaP accession number: phs000240.v1.p1). The combined dataset contains 301,143 autosomal SNPs and 3,226 subjects, and was used previously by Zhou and Guan (2017) to examine the relationship between p-values and Bayes factors.

We conducted five independent runs with different random seeds. The sparsity parameters for BVSR were set as $\pi_{min} = 0.0001$ and $\pi_{max} = 0.01$, which reflected a prior of 30–3000 marginally associated SNPs. The posterior mean model size in each run ranges from 430 to 622 (Table 2). The average number of iterations for ICF to converge

is 7.3 (95% interval = $[4, 13]$), which suggests that ICF is effective in analyzing real datasets. The speed improvement of fastBVSR over piMASS (a BVSR implementation that uses the Cholesky decomposition to fit Bayesian linear regression, a companion software package of Guan and Stephens (2011)) is dramatic: with model size around 500 and more than 3000 individuals, fastBVSR took 14 hours to run 2.1 million MCMC iterations, compared to more than 1000 hours used by piMASS on problems of matching size. Out of concern for the cumulative error in obtaining the matrix \mathbf{R} by updating the Cholesky decomposition (detailed in Section 3.1), we compared every 1000 steps the updated Cholesky decomposition with the directly computed Cholesky decomposition, and results suggested that the entry-wise difference was absolutely negligible.

We examine the inference of the hyperparameters, namely, the narrow-sense heritability h , the model size $|\gamma|$, and the sparsity parameter π . Table 2 shows the inferences of the three parameters are more or less consistent across five independent runs, and thus we combine the results from five independent runs in the following discussion. The posterior mean for h is 0.28 with 95% credible interval (0.1, 0.44). This estimate of the heritability is smaller than those reported in the literature: 0.62 from a twin study (Carbonaro et al., 2008), 0.29 from a sibling study (Chang et al., 2005), and 0.35 from an extended pedigree study (van Koolwijk et al., 2007). For comparison, using the same dataset, the heritability estimated by GCTA is 0.47 with standard error 0.11. The underestimation of h using fastBVSR is perhaps due to over-shrinking of the effect sizes (the posterior mean of prior effect size σ_β is 0.048). The posterior mean for π is 0.0016, which suggests that IOP is a very much polygenic phenotype, agreeing with the complex etiology of glaucoma (Weinreb et al., 2014), to which high IOP is a precursor phenotype.

	Combined	Run1	Run2	Run3	Run4	Run5
h	0.28	0.27	0.26	0.32	0.28	0.26
$ \gamma $	474	446	395	622	476	430
π	0.0016	0.0015	0.0013	0.0021	0.0016	0.0014

Table 2: Posterior mean estimates for the hyperparameters: heritability h , model size $|\gamma|$, and sparsity parameter π . In each MCMC iteration, heritability is re-estimated by computing the proportion of explained variance of \mathbf{y} using the sampled parameters. The first column gives the combined results from the five independent runs; each run has 2.1 million MCMC steps, including 0.1 million burn-in steps.

Lastly, we examine the top marginal signals detected by fastBVSR. Table 3 lists top 11 SNPs based on PIP inferred from fastBVSR, at an arbitrary PIP threshold of 0.25. Table 3 also includes PPA (posterior probability of association) based on the single SNP test for two choices of prior effect sizes: $\sigma_\beta = 0.05$ which is the posterior mean of σ_β inferred from fastBVSR and $\sigma_\beta = 0.5$ which is a popular choice for single SNP analysis. We use the posterior mean $\pi = 0.0016$ inferred from fastBVSR as the prior odds of association. Multiplying a Bayes factor and the prior odds we obtain the posterior odds, from which we obtain the PPA = $(\text{posterior odds})/(1 + \text{posterior odds})$. First the two columns of PPAs are highly similar in both ranking and magnitude. This observation agrees with our experience and theoretical studies (Zhou and Guan,

2017) that modest prior effect sizes produce similar evidence for association. Second, PPA and PIP have significantly different rankings (Wilcoxon rank test p-value = 0.025). There are two related explanations for this observation: if two SNPs are correlated, the PIPs tend to split between the two; conditioning on the other associated SNPs, the marginal associations tend to change significantly. Three known GWAS signals, *rs12150284*, *rs2025751*, *rs7518099/rs4656461*, which were replicated in our previous single-SNP analysis (Zhou and Guan, 2017), all show appreciable PIPs. Moreover, we note a potential novel association first reported in Zhou and Guan (2017): two SNPs located in *PEX14*, *rs12120962*, *rs12127400*, have PIPs of 0.57 and 0.27 respectively.

SNP	Chr	Pos (Mb)	PPA $_{\sigma_\beta=0.5}$	PPA $_{\sigma_\beta=0.05}$	PIP _{BVSR}
rs12120962	1	10.53	0.85	0.84	0.57
rs12127400	1	10.54	0.74	0.74	0.27
rs4656461	1	163.95	1.0	0.96	0.36
rs7518099	1	164.00	1.0	0.98	0.42
rs7645716	3	46.31	0.62	0.56	0.35
rs2025751	6	51.73	0.81	0.82	0.68
rs10757601	9	26.18	0.47	0.55	0.31
rs9783190	10	106.76	0.39	0.31	0.32
rs1381143	12	86.76	0.27	0.37	0.31
rs4984577	15	93.76	0.43	0.49	0.42
rs12150284	17	9.97	0.99	0.97	0.92

Table 3: SNPs with (Rao-Blackwellized) PIP > 0.25. Chr is short for chromosome, and Pos is short for position, which is measured in mega-basepair (Mb) according to HG18. PPA stands for the posterior inclusion probability. Computing PPA requires the Bayes factor and the prior odds. For prior odds we use 0.0016, which is the posterior mean of π inferred from fastBVSR. Bayes factors are computed using two priors effect sizes ($\sigma_\beta = 0.05, 0.5$). SNPs that are mentioned in the main text are highlighted in bold.

4 Discussion

We developed a novel algorithm, iterative complex factorization, to solve a class of penalized linear systems, proved its convergence, and demonstrated its effectiveness and efficiency through simulation studies. The novel algorithm can dramatically increase the speed of BVSR, which we demonstrated by analyzing a real dataset. In our simulation studies we only considered $n > p$ (in BVSR $p = |\gamma|$ is the size of the selected model). When $n \leq p$, the ICF can be implemented with the dual variable method of Saunders et al. (1998) (see also Lu et al., 2013), which hinges on the identity $(\mathbf{X}_\gamma^t \mathbf{X}_\gamma + \sigma_\beta^{-2} \mathbf{I})^{-1} \mathbf{X}_\gamma^t \mathbf{y} = \mathbf{X}_\gamma^t (\mathbf{X}_\gamma \mathbf{X}_\gamma^t + \sigma_\beta^{-2} \mathbf{I})^{-1} \mathbf{y}$.

One limitation of ICF is that it requires the Cholesky decomposition to obtain the effective complex factorization. Nevertheless, ICF still has a range of potential applications. The general ridge regression (Draper and Van Nostrand, 1979) is an obvious one, and here we point out two other examples. The first is the Bayesian sparse linear

mixed model (BSLMM) of Zhou et al. (2013), which can be viewed as a generalization of BVSR. BSLMM has two components, one is the linear component that corresponds to \mathbf{X}_γ and the other is the random effect component. The linear component requires one to solve a system that is similar to (11) (Text S2 Zhou et al., 2013); therefore ICF can be applied to BSLMM to make it more efficient. Another example of ICF application is using the variational method to fit the BVSR (Carbonetto and Stephens, 2012; Huang et al., 2016). In particular, Huang et al. (2016) estimated (γ, β) using an iterative method. In each iteration, the posterior mean for β is updated by solving a linear system that has the same form as (11), where the dimension is equal to the number of covariates with nonzero PIP estimates, and the diagonal matrix Σ depends on the PIP estimates. Applying ICF to the variational method might be more beneficial than applying it to BVSR because the dimension of the linear system is exceedingly larger in the variational method.

The iterative complex factorization method converges almost instantaneously, and it is far more accurate than other iterative algorithms such as the Gauss-Seidel method. Why does it converge so fast and at the same time so accurate? Our intuition is that the imaginary part of the update in (15), which involves the skew-symmetric matrix $\mathbf{S} = \mathbf{R}^t \Sigma - \Sigma \mathbf{R}$, whose Youla decomposition (Youla, 1961) suggests that it rotates $\hat{\beta}^{(k)}$, scales and flips its coordinates and rotates back, and which then appends itself to be the imaginary part of \mathbf{z} , provides a “worm hole” through the imaginary dimension in an otherwise purely real, unimaginary, solution path. Understanding how this works mathematically will have a far-reaching effect on Bayesian statistics, statistical genetics, machine learning, and computational biology.

Finally, we hope our new software fastBVSR, which is 10–100 times faster than the previous implementation of BVSR, will facilitate the wide-spread use of BVSR in analyzing or reanalyzing datasets from genome-wide association (GWA) and expression quantitative trait loci (eQTL) studies.

Supplementary Material

Fast model-fitting of Bayesian variable selection regression using the iterative complex factorization algorithm (Supplementary) (DOI: [10.1214/18-BA1120SUPP](https://doi.org/10.1214/18-BA1120SUPP); .pdf). The online appendix includes additional numerical studies (S1–S3), a summary of the MCMC algorithm (S4), examples for updating the Cholesky decomposition (S5) and a proof for the stationary distribution of the exchange algorithm (S6). S1 compares the performance of the iterative methods using simulated datasets; S2 studies the behavior of the convergence rates of the iterative methods; S3 provides numerical evidence for the assumption $\eta_{min}^2 \approx 0$ used in Section 2.3.

References

- Andrieu, C. and Roberts, G. O. (2009). “The pseudo-marginal approach for efficient Monte Carlo computations.” *The Annals of Statistics*, 697–725. [MR2502648](https://doi.org/10.1214/07-AOS574). doi: <https://doi.org/10.1214/07-AOS574>. 584

- Avron, H., Maymounkov, P., and Toledo, S. (2010). “Blendenpik: Supercharging LAPACK’s least-squares solver.” *SIAM Journal on Scientific Computing*, 32(3): 1217–1236. MR2639236. doi: <https://doi.org/10.1137/090767911>. 577, 582
- Berger, J. O., Pericchi, L. R., Ghosh, J., Samanta, T., De Santis, F., Berger, J., and Pericchi, L. (2001). “Objective Bayesian methods for model selection: introduction and comparison.” *Lecture Notes-Monograph Series*, 135–207. MR2000753. doi: <https://doi.org/10.1214/lnms/1215540968>. 574
- Broman, K. W. and Speed, T. P. (2002). “A model selection approach for the identification of quantitative trait loci in experimental crosses.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(4): 641–656. MR1979381. doi: <https://doi.org/10.1111/1467-9868.00354>. 586
- Carbonaro, F., Andrew, T., Mackey, D., Spector, T., and Hammond, C. (2008). “Heritability of intraocular pressure: a classical twin study.” *British Journal of Ophthalmology*, 92(8): 1125–1128. 588
- Carbonetto, P. and Stephens, M. (2012). “Scalable variational inference for Bayesian variable selection in regression, and its accuracy in genetic association studies.” *Bayesian Analysis*, 7(1): 73–108. MR2896713. doi: <https://doi.org/10.1214/12-BA703>. 590
- Chang, T. C., Congdon, N. G., Wojciechowski, R., Muñoz, B., Gilbert, D., Chen, P., Friedman, D. S., and West, S. K. (2005). “Determinants and heritability of intraocular pressure and cup-to-disc ratio in a defined older population.” *Ophthalmology*, 112(7): 1186–1191. 588
- Draper, N. R. and Van Nostrand, R. C. (1979). “Ridge regression and James-Stein estimation: review and comments.” *Technometrics*, 21(4): 451–466. MR0555086. doi: <https://doi.org/10.2307/1268284>. 576, 589
- Eldén, L. (1977). “Algorithms for the regularization of ill-conditioned least squares problems.” *BIT Numerical Mathematics*, 17(2): 134–145. MR0474912. 577
- Golub, G. H. and Van Loan, C. F. (2012). *Matrix computations, 3rd edition*. JHU Press. MR3024913. 577, 578, 581, 584
- Gough, B. (2009). *GNU scientific library reference manual*. Network Theory Ltd. 581
- Guan, Y. and Krone, S. M. (2007). “Small-world MCMC and Convergence to Multi-modal Distributions: From Slow Mixing to Fast Mixing.” *Annals of Applied Probability*, 17: 284–304. MR2292588. doi: <https://doi.org/10.1214/105051606000000772>. 584
- Guan, Y. and Stephens, M. (2011). “Bayesian variable selection regression for genome-wide association studies and other large-scale problems.” *The Annals of Applied Statistics*, 1780–1815. MR2884922. doi: <https://doi.org/10.1214/11-AOAS455>. 573, 574, 575, 576, 585, 586, 587, 588

- Hawkins, D. M. and Yin, X. (2002). “A faster algorithm for ridge regression of reduced rank data.” *Computational Statistics & Data Analysis*, 40(2): 253–262. [MR1924009](#). doi: [https://doi.org/10.1016/S0167-9473\(02\)00034-8](https://doi.org/10.1016/S0167-9473(02)00034-8). 577
- Huang, X., Wang, J., and Liang, F. (2016). “A Variational Algorithm for Bayesian Variable Selection.” *arXiv preprint arXiv:1602.07640*. 590
- Ibrahim, J. G. and Laud, P. W. (1991). “On Bayesian analysis of generalized linear models using Jeffreys’s prior.” *Journal of the American Statistical Association*, 86(416): 981–986. [MR1146346](#). 574
- Kohno, T., Kotakemori, H., Niki, H., and Usui, M. (1997). “Improving the modified Gauss-Seidel method for Z-matrices.” *Linear Algebra and its Applications*, 267: 113–123. [MR1479115](#). 580
- Lawson, C. L. and Hanson, R. J. (1995). *Solving least squares problems*. SIAM. [MR1349828](#). doi: <https://doi.org/10.1137/1.9781611971217>. 577
- Lee, S. H., Wray, N. R., Goddard, M. E., and Visscher, P. M. (2011). “Estimating missing heritability for disease from genome-wide association studies.” *The American Journal of Human Genetics*, 88(3): 294–305. 585
- Liang, F., Paulo, R., Molina, G., Clyde, M. A., and Berger, J. O. (2008). “Mixtures of g priors for Bayesian variable selection.” *Journal of the American Statistical Association*, 103(481). [MR2420243](#). doi: <https://doi.org/10.1198/016214507000001337>. 574
- Lu, Y., Dhillon, P., Foster, D. P., and Ungar, L. (2013). “Faster ridge regression via the subsampled randomized hadamard transform.” In *Advances in neural information processing systems*, 369–377. 589
- Meng, X., Saunders, M. A., and Mahoney, M. W. (2014). “LSRN: A parallel iterative solver for strongly over-or underdetermined systems.” *SIAM Journal on Scientific Computing*, 36(2): C95–C118. [MR3172249](#). doi: <https://doi.org/10.1137/120866580>. 577, 582
- Møller, J., Pettitt, A. N., Reeves, R., and Berthelsen, K. K. (2006). “An efficient Markov chain Monte Carlo method for distributions with intractable normalising constants.” *Biometrika*, 93(2): 451–458. [MR2278096](#). doi: <https://doi.org/10.1093/biomet/93.2.451>. 583
- Murray, I., Ghahramani, Z., and MacKay, D. (2012). “MCMC for doubly-intractable distributions.” *arXiv preprint arXiv:1206.6848*. 583, 584
- Niki, H., Harada, K., Morimoto, M., and Sakakihara, M. (2004). “The survey of preconditioners used for accelerating the rate of convergence in the Gauss–Seidel method.” *Journal of Computational and Applied Mathematics*, 164: 587–600. [MR2056901](#). doi: <https://doi.org/10.1016/j.cam.2003.11.012>. 580
- O’Hagan, A. and Forster, J. J. (2004). *Kendall’s advanced theory of statistics, volume 2B: Bayesian inference*, volume 2. Arnold. [MR3237119](#). 574

- Raftery, A. E., Madigan, D., and Hoeting, J. A. (1997). “Bayesian model averaging for linear regression models.” *Journal of the American Statistical Association*, 92(437): 179–191. [MR1436107](#). doi: <https://doi.org/10.2307/2291462>. 586
- Rokhlin, V. and Tygert, M. (2008). “A fast randomized algorithm for overdetermined linear least-squares regression.” *Proceedings of the National Academy of Sciences*, 105(36): 13212–13217. [MR2443725](#). doi: <https://doi.org/10.1073/pnas.0804869105>. 577
- Saunders, C., Gammerman, A., and Vovk, V. (1998). “Ridge regression learning algorithm in dual variables.” [MR3417807](#). 589
- Trefethen, L. N. and Bau III, D. (1997). *Numerical linear algebra*, volume 50. Siam. [MR1444820](#). doi: <https://doi.org/10.1137/1.9780898719574>. 576, 577
- Turlach, B. A. (2006). “An even faster algorithm for ridge regression of reduced rank data.” *Computational Statistics & Data Analysis*, 50(3): 642–658. [MR2196288](#). doi: <https://doi.org/10.1016/j.csda.2004.09.007>. 577
- van Koolwijk, L. M., Despriet, D. D., van Duijn, C. M., Cortes, L. M. P., Vingerling, J. R., Aulchenko, Y. S., Oostra, B. A., Klaver, C. C., and Lemij, H. G. (2007). “Genetic contributions to glaucoma: heritability of intraocular pressure, retinal nerve fiber layer thickness, and optic disc morphology.” *Investigative Ophthalmology & Visual Science*, 48(8): 3669–3676. 588
- Weinreb, R., Aung, T., and Medeiros, F. (2014). “The pathophysiology and treatment of glaucoma: A review.” *JAMA*, 311(18): 1901–1911. URL + <http://dx.doi.org/10.1001/jama.2014.3192> 588
- Xu, H. and Guan, Y. (2014). “Detecting Local Haplotype Sharing and Haplotype Association.” *Genetics*, 197(3): 823–838. 585
- Yang, J., Benyamin, B., McEvoy, B. P., Gordon, S., Henders, A. K., Nyholt, D. R., Madden, P. A., Heath, A. C., Martin, N. G., Montgomery, G. W., et al. (2010). “Common SNPs explain a large proportion of the heritability for human height.” *Nature Genetics*, 42(7): 565–569. 585
- Yang, J., Lee, S. H., Goddard, M. E., and Visscher, P. M. (2011). “GCTA: a tool for genome-wide complex trait analysis.” *The American Journal of Human Genetics*, 88(1): 76–82. 585
- Yang, S. and Matthias, K. G. (2007). “The optimal relaxation parameter for the SOR method applied to a classical model problem.” Technical report, Technical Report number TR2007-6, Department of Mathematics and Statistics, University of Maryland, Baltimore County. 581
- Youla, D. C. (1961). “A normal form for a matrix under the unitary congruence group.” *Canadian Journal of Mathematics*, 13: 694–704. [MR0132754](#). doi: <https://doi.org/10.4153/CJM-1961-059-8>. 590

- Young, D. (1954). “Iterative methods for solving partial difference equations of elliptic type.” *Transactions of the American Mathematical Society*, 76(1): 92–111. MR0059635. doi: <https://doi.org/10.2307/1990745>. 581
- Zhou, Q. and Guan, Y. (2017). “On the Null Distribution of Bayes Factors in Linear Regression.” *Journal of American Statistical Association*, to appear. 587, 588, 589
- Zhou, Q. and Guan, Y. (2018). “Fast model-fitting of Bayesian variable selection regression using the iterative complex factorization algorithm (Supplementary).” *Bayesian Analysis*. doi: <https://doi.org/10.1214/18-BA1120SUPP>. 580
- Zhou, X., Carbonetto, P., and Stephens, M. (2013). “Polygenic modeling with Bayesian sparse linear mixed models.” *PLoS Genet*, 9(2): e1003264. 577, 590

Acknowledgments

We thank the two anonymous reviewers for their helpful comments that improved the quality of our presentation.