

UC Davis

UC Davis Previously Published Works

Title

A fast and efficient Gibbs sampler for BayesB in whole-genome analyses

Permalink

<https://escholarship.org/uc/item/5cm7m5bq>

Journal

Genetics Selection Evolution, 47(1)

ISSN

0999-193X

Authors

Cheng, Hao
Qu, Long
Garrick, Dorian J
[et al.](#)

Publication Date

2015-12-01

DOI

10.1186/s12711-015-0157-x

Peer reviewed

RESEARCH ARTICLE

Open Access



A fast and efficient Gibbs sampler for BayesB in whole-genome analyses

Hao Cheng^{1,2}, Long Qu³, Dorian J. Garrick^{1,4} and Rohan L. Fernando^{1*}

Abstract

Background: In whole-genome analyses, the number p of marker covariates is often much larger than the number n of observations. Bayesian multiple regression models are widely used in genomic selection to address this problem of $p \gg n$. The primary difference between these models is the prior assumed for the effects of the covariates. Usually in the BayesB method, a Metropolis–Hastings (MH) algorithm is used to jointly sample the marker effect and the locus-specific variance, which may make BayesB computationally intensive. In this paper, we show how the Gibbs sampler without the MH algorithm can be used for the BayesB method.

Results: We consider three different versions of the Gibbs sampler to sample the marker effect and locus-specific variance for each locus. Among the Gibbs samplers that were considered, the most efficient sampler is about 2.1 times as efficient as the MH algorithm proposed by Meuwissen et al. and 1.7 times as efficient as that proposed by Habier et al.

Conclusions: The three Gibbs samplers presented here were twice as efficient as Metropolis–Hastings samplers and gave virtually the same results.

Background

In whole-genome analyses, the number p of marker covariates is often much larger than the number n of observations. Bayesian multiple regression models are widely used in genomic selection to address this problem of $p \gg n$. The primary difference between these models is the prior assumed for the effects of the covariates. These priors and their effects on inference have been recently reviewed by Gianola [1]. In most Bayesian analyses of whole-genome data, inferences are based on Markov chains constructed to have a stationary distribution equal to the posterior distribution of the unknown parameters of interest [2]. This is often done by employing a Gibbs sampler where samples are drawn from the full-conditional distributions of the parameters [3].

It can be shown that in BayesA introduced by Meuwissen et al. [4], the prior for each marker effect follows a scaled t distribution [5]. However when the prior for the marker effect is specified as a t distribution, its

full-conditional is not of a known form. Fortunately, this prior can also be specified as a normally distributed marker effect conditional on a locus-specific variance, which is given a scaled inverted chi-square distribution. When marginalized over the variance, this gives a t distribution for the marker effect [5]. Thus, the posterior for the marker effect would be identical under both these priors. The second form of the prior, however, is more convenient because it results in the full-conditional for the marker effect having a normal distribution.

BayesA is a special case of BayesB, also introduced by Meuwissen et al. [4], where the prior for each marker effect follows a mixture distribution with a point mass at zero with probability π and a univariate- t distribution with probability $1 - \pi$ [5]. When $\pi = 0$, BayesB becomes BayesA. When the marker effect is non-null, as in BayesA, the second form of the prior leads to the full-conditional of the marker effect being normal. Nevertheless, Meuwissen et al. [4] used a Metropolis–Hastings (MH) algorithm to jointly sample the marker effect and the locus-specific variance because they argued that “the Gibbs sampler will not move through the entire sampling space” for BayesB. In their MH algorithm, they use the prior distribution of

*Correspondence: rohan@iastate.edu

¹ Department of Animal Science, Iowa State University, Ames 50011, IA, USA

Full list of author information is available at the end of the article

the locus-specific variance as the proposal distribution. When π is high, the proposed values for the marker effect will be zero with high probability. Thus, for each locus, 100 cycles of MH algorithm were used in their paper, which makes BayesB computationally intensive. Habier et al. [6] used an alternative proposal, where the marker effect was zero with probability 0.5, that leads to a more efficient MH algorithm. For each locus, five cycles of MH were used to sample marker effects in this efficient MH method.

In this paper, we will show how Gibbs samplers without the MH algorithm can be used for the BayesB method. Recall that by introducing a locus-specific variance into BayesA, the full-conditional for the marker effects becomes normal. Similarly, in this paper we show that by introducing a variable δ_j in BayesB, indicating whether the marker effect for a locus is zero or non-zero, the marker effect and locus-specific variance can be sampled using Gibbs. We consider three different versions of the Gibbs sampler to sample each marker effect, locus-specific variance and its indicator variable δ_j . The objectives of this paper are to introduce these samplers and study their performance.

Methods

Statistical methods

BayesB introduced by Meuwissen et al. [4] assumes that each locus-specific variance follows a mixture distribution. However, following Gianola [5], we prefer to specify the mixture at the level of the marker effect instead of the locus specific variance. In this formulation, the prior for the marker effect is a mixture with a point mass at zero and a univariate normal distribution conditional on σ_j^2 :

$$(\alpha_j | \sigma_j^2) \begin{cases} = 0 & \text{with probability } \pi \\ \sim N(0, \sigma_j^2) & \text{with probability } (1 - \pi), \end{cases}$$

where σ_j^2 follows a scaled inverted chi-square distribution and π is treated as known. Employing the concept of data augmentation, it is convenient to write the marker effect α_j as $\alpha_j = \beta_j \delta_j$, where we introduce a Bernoulli variable δ_j with probability of success $1 - \pi$ and normally distributed variable β_j with mean zero and variance σ_j^2 , which has a scaled inverted chi-square distribution. As shown below, Gibbs sampling can be used to draw samples for these unknowns.

Gibbs samplers for BayesB

Here, we present three Gibbs samplers for BayesB. The first is a single-site Gibbs sampler, where all parameters are sampled from their full conditional distributions. The second is a joint Gibbs sampler, where δ_j, β_j are sampled from the joint full-conditional distribution $f(\delta_j, \beta_j | \mathbf{y}, \mu, \boldsymbol{\beta}_{-j}, \boldsymbol{\delta}_{-j}, \boldsymbol{\xi}, \sigma_e^2)$, where $\boldsymbol{\xi} = [\sigma_1^2, \sigma_2^2, \dots, \sigma_k^2]$, because δ_j and β_j are highly dependent. Carlin and Chib

[7] have shown that the prior used for parameters that are not in the model does not affect the Bayes factor. Thus, this prior, which they call a pseudo prior, can be chosen to improve mixing of the sampler. Following Carlin and Chib, the third sampler is a Gibbs sampler where a pseudo prior is used for β_j when δ_j is zero. Godsill [8] has shown that the marginal posterior for a parameter in the model does not depend on the choice of pseudo priors. It has been suggested to choose the full conditional distribution for β_j when it is in the model as the pseudo prior [7, 8]. This choice is justified by showing that using this pseudo prior is equivalent to sampling δ_j from $f(\delta_j | \mathbf{y}, \mu, \boldsymbol{\beta}_{-j}, \boldsymbol{\delta}_{-j}, \boldsymbol{\xi}, \sigma_e^2)$ [8]. However, in BayesB, use of the exact full conditional distribution as the pseudo prior will require MH to sample σ_j^2 and σ_e^2 . Thus in this paper, a distribution close to the full conditional is used.

BayesB model with data augmentation model

$$y_i = \mu + \sum_{j=1}^k X_{ij} \beta_j \delta_j + e_i, \tag{1}$$

where y_i is the phenotype for individual i , μ is the overall mean, k is the number of SNPs, X_{ij} is the genotype covariate at locus j for animal i (coded as 0, 1, 2), β_j is the allele substitution effect for locus j , δ_j is an indicator variable and e_i is the random residual effect for individual i .

Priors

The prior for μ is a constant. The prior for e_i is $e_i | \sigma_e^2 \stackrel{iid}{\sim} N(0, \sigma_e^2)$ and $(\sigma_e^2 | \nu_e, S_e^2) \sim \nu_e S_e^2 \chi_{\nu_e}^{-2}$. The prior for β_j is $\beta_j | \sigma_j^2 \sim N(0, \sigma_j^2)$. The prior for $(\sigma_j^2 | \nu_\beta, S_\beta^2) \sim \nu_\beta S_\beta^2 \chi_{\nu_\beta}^{-2}$. The prior for δ_j is

$$(\delta_j | \pi) \begin{cases} = 1 & \text{probability } (1 - \pi) \\ = 0 & \text{probability } \pi. \end{cases}$$

Single-site Gibbs sampler

The full conditional distributions of μ, σ_j^2 and σ_e^2 are well-known [4, 9, 10]. Thus they are presented here without derivations. The full conditional of μ is a normal distribution with mean $\frac{1}{n} \mathbf{e}_\mu$ and variance $\frac{\sigma_e^2}{n}$, where $\mathbf{e}_\mu = \mathbf{y} - \sum_{j=1}^k \mathbf{X}_j \beta_j \delta_j$ and n is the number of individuals. The full conditional distributions of σ_j^2 and σ_e^2 are both scaled inverted chi-square distributions; for σ_j^2 , the scale parameter is $\frac{(\nu_\beta S_\beta^2 + \beta_j^2)}{\nu_\beta + 1}$ and the degrees of freedom parameter are $\nu_\beta + 1$; for σ_e^2 , the scale parameter is $\frac{\nu_e S_e^2 + \mathbf{e}' \mathbf{e}}{\nu_e + n}$ and the degrees of freedom parameter are $\nu_e + n$, where $\mathbf{e} = \mathbf{y} - \mathbf{1}\mu - \sum_{j=1}^k \mathbf{X}_j \beta_j \delta_j$.

Next, we derive the full conditional distributions of β_j and δ_j . These full conditional distributions are proportional to the joint distribution of all parameters and \mathbf{y} , which can be written as

$$\begin{aligned}
 & f(\mathbf{y}, \mu, \boldsymbol{\beta}, \boldsymbol{\delta}, \boldsymbol{\xi}, \sigma_e^2) \\
 & \propto f(\mathbf{y} \mid \mu, \boldsymbol{\beta}, \boldsymbol{\delta}, \boldsymbol{\xi}, \sigma_e^2) f(\boldsymbol{\beta} \mid \boldsymbol{\xi}) f(\boldsymbol{\delta}) f(\boldsymbol{\xi}) f(\sigma_e^2) \\
 & \propto (\sigma_e^2)^{-\frac{n}{2}} \exp\left(-\frac{\mathbf{e}'\mathbf{e}}{2\sigma_e^2}\right) \\
 & \quad \times \prod_{j=1}^k \frac{1}{\sqrt{2\pi}(\sigma_j^2)^{\frac{1}{2}}} \exp\left(-\frac{\beta_j^2}{2\sigma_j^2}\right) \\
 & \quad \times \prod_{j=1}^k \pi^{(1-\delta_j)} (1-\pi)^{\delta_j} \\
 & \quad \times \prod_{j=1}^k \frac{\left(S_\beta^2 \frac{v_\beta}{2}\right)^{\frac{v_\beta}{2}}}{\Gamma\left(\frac{v_\beta}{2}\right)} (\sigma_j^2)^{\left(-\frac{v_\beta+2}{2}\right)} \exp\left(-\frac{v_\beta S_\beta^2}{2\sigma_j^2}\right) \\
 & \quad \times \frac{\left(S_e^2 \frac{v_e}{2}\right)^{\frac{v_e}{2}}}{\Gamma\left(\frac{v_e}{2}\right)} (\sigma_e^2)^{\left(-\frac{v_e+2}{2}\right)} \exp\left(-\frac{v_e S_e^2}{2\sigma_e^2}\right).
 \end{aligned} \tag{2}$$

The full conditional distribution of β_j is now obtained by dropping factors that do not involve β_j , which gives

$$\begin{aligned}
 & f(\beta_j \mid \mathbf{y}, \mu, \boldsymbol{\beta}_{-j}, \boldsymbol{\delta}, \boldsymbol{\xi}, \sigma_e^2) \\
 & \propto f(\mathbf{y} \mid \mu, \boldsymbol{\beta}, \boldsymbol{\delta}, \boldsymbol{\xi}, \sigma_e^2) f(\beta_j \mid \sigma_j^2) \\
 & \propto \exp\left(-\frac{\mathbf{e}'\mathbf{e}}{2\sigma_e^2}\right) \exp\left(-\frac{\beta_j^2}{2\sigma_j^2}\right) \\
 & \propto \exp\left[-\frac{(\mathbf{w} - \mathbf{X}_j \beta_j \delta_j)' (\mathbf{w} - \mathbf{X}_j \beta_j \delta_j)}{2\sigma_e^2}\right] \\
 & \quad \times \exp\left(-\frac{\beta_j^2}{2\sigma_j^2}\right),
 \end{aligned}$$

where $\mathbf{w} = \mathbf{y} - \mathbf{1}\mu - \sum_{j' \neq j} \mathbf{X}_{j'} \beta_{j'} \delta_{j'}$. When $\delta_j = 1$,

$$\begin{aligned}
 & f(\beta_j \mid \mathbf{y}, \mu, \boldsymbol{\beta}_{-j}, \boldsymbol{\delta}, \boldsymbol{\xi}, \sigma_e^2) \\
 & \propto \exp\left[-\frac{(\mathbf{w} - \mathbf{X}_j \beta_j)' (\mathbf{w} - \mathbf{X}_j \beta_j)}{2\sigma_e^2}\right] \exp\left(-\frac{\beta_j^2}{2\sigma_j^2}\right) \\
 & \propto \exp\left[-\frac{1}{2} \frac{\left(\beta_j - \frac{\mathbf{X}_j' \mathbf{w}}{c_j}\right)^2}{\frac{\sigma_e^2}{c_j}}\right]
 \end{aligned} \tag{3}$$

where $c_j = \mathbf{X}_j' \mathbf{X}_j + \frac{\sigma_e^2}{\sigma_j^2}$. Now, (3) can be recognized as the kernel of a normal distribution with mean $\frac{\mathbf{X}_j' \mathbf{w}}{c_j}$ and variance $\frac{\sigma_e^2}{c_j}$. When $\delta_j = 0$,

$$\begin{aligned}
 & f(\beta_j \mid \mathbf{y}, \mu, \boldsymbol{\beta}_{-j}, \boldsymbol{\delta}, \boldsymbol{\xi}, \sigma_e^2) \\
 & \propto \exp\left(-\frac{\mathbf{w}'\mathbf{w}}{2\sigma_e^2}\right) \exp\left(-\frac{\beta_j^2}{2\sigma_j^2}\right),
 \end{aligned}$$

and dropping the factor $\left[\exp\left(-\frac{\mathbf{w}'\mathbf{w}}{2\sigma_e^2}\right)\right]$, which is free of β_j , gives

$$f(\beta_j \mid \mathbf{y}, \mu, \boldsymbol{\beta}_{-j}, \boldsymbol{\delta}, \boldsymbol{\xi}, \sigma_e^2) \propto \exp\left(-\frac{\beta_j^2}{2\sigma_j^2}\right),$$

which is the kernel of a normal distribution with null mean and variance σ_j^2 . Thus,

$$p(\beta_j \mid ELSE) = \begin{cases} \sim N\left(\frac{\mathbf{X}_j' \mathbf{w}}{c_j}, \frac{\sigma_e^2}{c_j}\right) & \text{when } \delta_j = 1, \\ \sim N(0, \sigma_j^2) & \text{when } \delta_j = 0, \end{cases}$$

where ELSE stands for all the other parameters and \mathbf{y} . This means when $\delta_j = 1$, the sampling of β_j is identical to that in BayesA; when $\delta_j = 0$, β_j is sampled from its prior. This is different from the original implementation of BayesB introduced by Meuwissen et al. [4].

Similarly, the full conditional distribution of δ_j can be obtained from (2) by dropping all factors free of δ_j , which gives

$$\begin{aligned}
 & \Pr(\delta_j = 1 \mid ELSE) \\
 & \propto f(\mathbf{y} \mid \mu, \boldsymbol{\beta}, \boldsymbol{\delta}, \boldsymbol{\xi}, \sigma_e^2) \Pr(\delta_j = 1) \\
 & \propto \exp\left[-\frac{(\mathbf{w} - \mathbf{X}_j \beta_j \delta_j)' (\mathbf{w} - \mathbf{X}_j \beta_j \delta_j)}{2\sigma_e^2}\right] (1-\pi) \\
 & \propto \exp\left[-\frac{(\mathbf{w} - \mathbf{X}_j \beta_j)' (\mathbf{w} - \mathbf{X}_j \beta_j)}{2\sigma_e^2}\right] (1-\pi),
 \end{aligned}$$

$$\begin{aligned}
 & \Pr(\delta_j = 0 \mid ELSE) \\
 & \propto f(\mathbf{y} \mid \mu, \boldsymbol{\beta}, \boldsymbol{\delta}, \boldsymbol{\xi}, \sigma_e^2) \Pr(\delta_j = 0) \\
 & \propto \exp\left[-\frac{(\mathbf{w} - \mathbf{X}_j \beta_j \delta_j)' (\mathbf{w} - \mathbf{X}_j \beta_j \delta_j)}{2\sigma_e^2}\right] \pi \\
 & \propto \exp\left(-\frac{\mathbf{w}'\mathbf{w}}{2\sigma_e^2}\right) \pi.
 \end{aligned}$$

Thus,

$$\begin{aligned} & \Pr(\delta_j = 1 \mid ELSE) \\ &= \frac{(1 - \pi)d_j}{(1 - \pi)d_j + \pi \exp\left(-\frac{\mathbf{w}'\mathbf{w}}{2\sigma_e^2}\right)}, \end{aligned}$$

$$\text{where } d_j = \exp\left[-\frac{(\mathbf{w}-X_j\beta_j)'(\mathbf{w}-X_j\beta_j)}{2\sigma_e^2}\right].$$

Joint Gibbs sampler

The same priors as in single-site Gibbs sampler are used here. The only difference is that δ_j, β_j are sampled from their joint full conditional distribution, which can be written as the product of the full conditional distribution of β_j given δ_j and marginal full conditional distribution of δ_j :

$$\begin{aligned} & f(\delta_j, \beta_j \mid \mathbf{y}, \mu, \boldsymbol{\beta}_{-j}, \boldsymbol{\delta}_{-j}, \boldsymbol{\xi}, \sigma_e^2) \\ &= f(\beta_j \mid \delta_j, \mathbf{y}, \mu, \boldsymbol{\beta}_{-j}, \boldsymbol{\delta}_{-j}, \boldsymbol{\xi}, \sigma_e^2) \\ & \quad \times \Pr(\delta_j \mid \mathbf{y}, \mu, \boldsymbol{\beta}_{-j}, \boldsymbol{\delta}_{-j}, \boldsymbol{\xi}, \sigma_e^2). \end{aligned}$$

Thus, δ_j is first sampled from $f(\delta_j \mid \mathbf{y}, \mu, \boldsymbol{\beta}_{-j}, \boldsymbol{\delta}_{-j}, \boldsymbol{\xi}, \sigma_e^2)$. Then β_j is sampled from $f(\beta_j \mid \delta_j, \mathbf{y}, \mu, \boldsymbol{\beta}_{-j}, \boldsymbol{\delta}_{-j}, \boldsymbol{\xi}, \sigma_e^2)$, which is identical to the sampling of β_j in BayesB with single-site Gibbs sampler. The marginal full conditional for δ_j can be written as:

$$\begin{aligned} & f(\delta_j \mid \mathbf{y}, \mu, \boldsymbol{\beta}_{-j}, \boldsymbol{\delta}_{-j}, \boldsymbol{\xi}, \sigma_e^2) \\ & \propto f(\mathbf{y} \mid \delta_j, \mu, \boldsymbol{\beta}_{-j}, \boldsymbol{\delta}_{-j}, \boldsymbol{\xi}, \sigma_e^2) \Pr(\delta_j) \\ & \propto f(\mathbf{w} \mid \delta_j, \sigma_j^2, \sigma_e^2) \Pr(\delta_j), \end{aligned} \tag{4}$$

where $\mathbf{w} = \mathbf{y} - \mathbf{1}\mu - \sum_{j' \neq j} X_{j'}\beta_{j'}\delta_{j'} = \mathbf{e} + X_j\beta_j\delta_j$. Now

$f(\mathbf{w} \mid \delta_j, \sigma_j^2, \sigma_e^2)$ is a multivariate normal distribution with mean $E(\mathbf{e} + X_j\beta_j\delta_j \mid \delta_j, \sigma_j^2, \sigma_e^2)$ and variance $Var(X_j\beta_j\delta_j + \mathbf{e} \mid \delta_j, \sigma_j^2, \sigma_e^2)$. When $\delta_j = 1$, it becomes a multivariate normal with null mean and variance $X_jX_j'\sigma_j^2 + I\sigma_e^2$; when $\delta_j = 0$, it becomes a multivariate normal with null mean and variance $I\sigma_e^2$. Thus, samples can be drawn using

$$\begin{aligned} & f(\delta_j = 1 \mid \mathbf{y}, \mu, \boldsymbol{\beta}_{-j}, \boldsymbol{\delta}_{-j}, \boldsymbol{\xi}, \sigma_e^2) \\ &= \frac{h_1 \Pr(\delta_j = 1)}{h_1 \Pr(\delta_j = 1) + h_0 \Pr(\delta_j = 0)}, \end{aligned} \tag{5}$$

$$= \frac{1}{1 + \frac{h_0 \Pr(\delta_j=0)}{h_1 \Pr(\delta_j=1)}} \tag{6}$$

where $h_i = f(\mathbf{w} \mid \delta_j = i, \sigma_j^2, \sigma_e^2)$. However, evaluating the multivariate normal distribution $f(\mathbf{w} \mid \delta_j, \sigma_j^2, \sigma_e^2)$ is computationally intense. An efficient way is to use the univariate distribution of $X_j'\mathbf{w}$, which contains all the information from \mathbf{w} about β_j , instead of the distribution of \mathbf{w} , which is a multivariate. Thus, (5) can be written as

$$\begin{aligned} & f(\delta_j = 1 \mid \mathbf{y}, \mu, \boldsymbol{\beta}_{-j}, \boldsymbol{\delta}_{-j}, \boldsymbol{\xi}, \sigma_e^2) \\ &= \frac{m_1 \Pr(\delta_j = 1)}{m_1 \Pr(\delta_j = 1) + m_0 \Pr(\delta_j = 0)}, \end{aligned}$$

where $m_i = f(X_j'\mathbf{w} \mid \delta_j = i, \mu, \boldsymbol{\beta}_{-j}, \boldsymbol{\delta}_{-j}, \boldsymbol{\xi}, \sigma_e^2)$, m_1 is an univariate normal distribution with null mean and variance $(X_j'X_j)^2\sigma_j^2 + X_j'X_j\sigma_e^2$, and m_0 is an univariate normal distribution with null mean and variance $X_j'X_j\sigma_e^2$.

Gibbs sampler with pseudo priors

Here, following Carlin and Chib [7], a pseudo prior is used for β_j when δ_j is zero. They proposed to use the full conditional distribution of β_j when $\delta_j = 1$ as the pseudo prior for β_j when $\delta_j = 0$ [7, 8], which results in the prior for β_j as

$$\begin{aligned} & \beta_j \mid \delta_j \\ &= \begin{cases} \sim N(\mathbf{0}, \sigma_j^2) & \delta_j = 1 \\ \sim f(\beta_j \mid \delta_j = 1, \mathbf{y}, \mu, \boldsymbol{\beta}_{-j}, \boldsymbol{\delta}_{-j}, \boldsymbol{\xi}, \sigma_e^2) & \delta_j = 0. \end{cases} \end{aligned}$$

We show below that the posterior mean of the marker effect, $\alpha_j = \beta_j\delta_j$, does not depend on the pseudo prior. This posterior mean can be written as:

$$\begin{aligned} & E(\beta_j\delta_j \mid \mathbf{y}) \\ &= \sum_{\delta_j} \int \beta_j\delta_j f(\beta_j, \delta_j \mid \mathbf{y}) d\beta_j \\ &= \frac{\sum_{\delta_j} \int \beta_j\delta_j f(\mathbf{y} \mid \beta_j, \delta_j) f(\beta_j \mid \delta_j) \Pr(\delta_j) d\beta_j}{f(\mathbf{y})} \tag{7} \\ &= \frac{\int \beta_j f(\mathbf{y} \mid \beta_j, \delta_j = 1) f(\beta_j \mid \delta_j = 1) \Pr(\delta_j = 1) d\beta_j}{f(\mathbf{y})}. \end{aligned}$$

The numerator in (7) is free of the pseudo prior: $f(\beta_j \mid \delta_j = 0)$. Furthermore, it can be seen from the model equation (1) that the value of \mathbf{y} is free of β_j when $\delta_j = 0$. Thus, the marginal distribution of \mathbf{y} , the denominator of (7), does not depend on the pseudo prior, which is the distribution of β_j when $\delta_j = 0$. As both the numerator and denominator of (1) are free of the pseudo prior, it follows that the posterior mean of α_j does not depend on the pseudo prior for β_j .

We show here that, given this pseudo prior, the full conditional for δ_j is identical to the marginal full conditional

distribution of δ_j , $\Pr(\delta_j | \mathbf{y}, \mu, \boldsymbol{\beta}_{-j}, \boldsymbol{\delta}_{-j}, \boldsymbol{\xi}, \sigma_e^2)$, which is used in the joint Gibbs sampler.

The full conditional probability of $\delta_j = 1$ can be written as:

$$\begin{aligned} \Pr(\delta_j = 1 | ELSE) &= \frac{g_1}{g_1 + g_0} \\ &= \frac{1}{1 + \frac{g_0}{g_1}}, \end{aligned} \tag{8}$$

where

$$\begin{aligned} g_i &= \Pr(\delta_j = i | ELSE) \\ &\propto f(\mathbf{y} | \delta_j = i, \boldsymbol{\beta}_j, \boldsymbol{\delta}_{-j}, \boldsymbol{\beta}_{-j}, \mu, \boldsymbol{\xi}, \sigma_e^2) \\ &\quad \times f(\boldsymbol{\beta}_j | \delta_j = i, \sigma_j^2) \times p(\delta_j = i) \\ &\propto f(\mathbf{w} | \delta_j = i, \boldsymbol{\beta}_j, \sigma_e^2) \\ &\quad \times f(\boldsymbol{\beta}_j | \delta_j = i, \sigma_j^2) \times p(\delta_j = i) \end{aligned}$$

with $i = 0$ or 1 .

The ratio in the denominator of (8) is:

$$\begin{aligned} \frac{g_0}{g_1} &= f(\mathbf{w} | \delta_j = 0, \boldsymbol{\beta}_j, \sigma_e^2) \\ &\times \frac{f(\boldsymbol{\beta}_j | \delta_j = 0)}{f(\mathbf{w} | \delta_j = 1, \boldsymbol{\beta}_j, \sigma_e^2) \times f(\boldsymbol{\beta}_j | \delta_j = 1, \sigma_j^2)} \\ &\times \frac{\Pr(\delta_j = 0)}{\Pr(\delta_j = 1)}. \end{aligned} \tag{10}$$

In the above equation, (9) is identical to h_0 in (5) used in the joint Gibbs sampler, because $f(\mathbf{w} | \delta_j = 0, \boldsymbol{\beta}_j, \sigma_e^2)$ is also a multivariate normal with null mean and variance $I\sigma_e^2$. We show below that (10) is identical to h_1^{-1} . Our proposed prior for $\boldsymbol{\beta}_j$ when $\delta_j = 0$ (the pseudo prior) can be written as:

$$\begin{aligned} p(\boldsymbol{\beta}_j | \delta_j = 0) &= f(\boldsymbol{\beta}_j | \delta_j = 1, \mathbf{y}, \mu, \boldsymbol{\beta}_{-j}, \boldsymbol{\delta}_{-j}, \boldsymbol{\xi}, \sigma_e^2) \\ &= f(\boldsymbol{\beta}_j | \delta_j = 1, \mathbf{w}, \sigma_e^2, \sigma_j^2) \\ &= \frac{f(\mathbf{w} | \delta_j = 1, \boldsymbol{\beta}_j, \sigma_e^2) f(\boldsymbol{\beta}_j | \delta_j = 1, \sigma_j^2)}{f(\mathbf{w} | \delta_j = 1, \sigma_j^2, \sigma_e^2)}. \end{aligned} \tag{11}$$

After replacing the pseudo prior in (10) with (11), it becomes $\frac{1}{f(\mathbf{w} | \delta_j = 1, \sigma_j^2, \sigma_e^2)}$, which is identical to h_1^{-1} .

Thus, the ratio $\frac{g_0}{g_1}$ in (8) is identical to $\frac{h_0 \Pr(\delta_j=1)}{h_1 \Pr(\delta_j=0)}$ in (6),

which proves the full conditional probability (8) of $\delta_j = 1$, when the proposed prior is used, is identical to (6), the marginal full conditional probability of $\delta_j = 1$, which is used in the joint Gibbs sampler.

Use of the exact full conditional distribution as the pseudo prior in BayesB, however, will require MH to sample σ_j^2 and σ_e^2 . Thus, a distribution close to the full conditional is used here. Here, we use a normal distribution with mean $\frac{\mathbf{X}'_j \mathbf{w}}{\mathbf{X}'_j \mathbf{X}_j + \tilde{\lambda}}$ and variance $\frac{\tilde{\sigma}_e^2}{\mathbf{X}'_j \mathbf{X}_j + \tilde{\lambda}}$, where $\tilde{\lambda} = \frac{\tilde{\sigma}_e^2}{\tilde{\sigma}_j^2}$, and $\tilde{\sigma}_e^2$ and $\tilde{\sigma}_j^2$ are means of the prior distributions for the residual and the marker effect variances, respectively.

Next, we will show the derivation of the full conditionals, which are proportional to the joint distribution of all parameters and \mathbf{y} . Here, the joint distribution of all parameters and \mathbf{y} can be written as:

$$\begin{aligned} f(\mathbf{y}, \mu, \boldsymbol{\beta}, \boldsymbol{\delta}, \boldsymbol{\xi}, \sigma_e^2) &\propto f(\mathbf{y} | \mu, \boldsymbol{\beta}, \boldsymbol{\delta}, \boldsymbol{\xi}, \sigma_e^2) f(\boldsymbol{\beta} | \boldsymbol{\delta}, \boldsymbol{\xi}) f(\boldsymbol{\delta}) f(\boldsymbol{\xi}) f(\sigma_e^2) \end{aligned} \tag{12}$$

$$\propto (\sigma_e^2)^{-\frac{n}{2}} \exp\left(-\frac{\mathbf{e}'\mathbf{e}}{2\sigma_e^2}\right) \tag{13}$$

$$\times \prod_{j=1}^k \left[\frac{1}{\sqrt{2\pi}(\sigma_j^2)^{\frac{1}{2}}} \exp\left(-\frac{\beta_j^2}{2\sigma_j^2}\right) \right]^{\delta_j} \tag{14}$$

$$\times \prod_{j=1}^k \left[\frac{1}{\sqrt{2\pi} \left(\frac{\tilde{\sigma}_e^2}{\mathbf{X}'_j \mathbf{X}_j + \tilde{\lambda}}\right)^{\frac{1}{2}}} \right]^{1-\delta_j} \tag{15}$$

$$\times \prod_{j=1}^k \left\{ \exp\left[-\frac{\left(\beta_j - \frac{\mathbf{X}'_j \mathbf{w}}{\mathbf{X}'_j \mathbf{X}_j + \tilde{\lambda}}\right)^2}{2 \frac{\tilde{\sigma}_e^2}{\mathbf{X}'_j \mathbf{X}_j + \tilde{\lambda}}}\right] \right\}^{1-\delta_j} \tag{16}$$

$$\times \prod_{j=1}^k \pi^{(1-\delta_j)} (1-\pi)^{\delta_j} \tag{17}$$

$$\times \prod_{j=1}^k \frac{\left(\frac{\sigma_\beta^2 v_\beta}{2}\right)^{\frac{v_\beta}{2}}}{\Gamma\left(\frac{v_\beta}{2}\right)} (\sigma_j^2)^{\left(-\frac{v_\beta+2}{2}\right)} \exp\left(-\frac{v_\beta S_\beta^2}{2\sigma_j^2}\right) \tag{18}$$

Table 1 Efficiency of alternative MCMC samplers for BayesB

	Alternative MCMC samplers				
	MH	Efficient MH	Single-site Gibbs	Joint Gibbs	Gibbs with pseudo prior
Computing time	90,009	70,714	52,452	44,726	47,043
Effective sample size	25,262	24,588	24,684	26,757	25,036
Effective samples/s	0.280	0.347	0.471	0.598	0.532

Efficiency of alternative MCMC samplers for BayesB. Results are given for the computing time in seconds to obtain 50,000 samples, effective sample size and effective samples/s for BayesB using Metropolis–Hastings (MH), single-site Gibbs sampler, joint Gibbs sampler and Gibbs sampler with pseudo priors

$$\times \frac{(S_e^2 \frac{\nu_e}{2})^{\frac{\nu_e}{2}}}{\Gamma(\frac{\nu_e}{2})} (\sigma_e^2)^{-\frac{\nu_e+2}{2}} \exp\left(-\frac{\nu_e S_e^2}{2\sigma_e^2}\right). \quad (19)$$

It is easy to see that the full conditional distribution of σ_e^2 , which does not involve δ_j , is the same as that in the single-site Gibbs sampler. Although μ also appears in w in (16), (16) has no effect on the full conditional of μ because the columns of X , which are always centered, are orthogonal to the column vectors of ones so that $X'_j \mathbf{1} \mu = \mathbf{0}$. Thus, the full conditional of μ is the same as that in the single-site Gibbs sampler. When $\delta_j = 1$, the full conditional distribution of β_j is identical to that in the single-site Gibbs sampler. When $\delta_j = 0$, (16) is the only part that includes β_j . Thus the full conditional distribution of β_j is:

$$p(\beta_j | ELSE) = \begin{cases} \sim N\left(\frac{X'_j w}{X'_j X_j + \frac{\sigma_e^2}{\sigma_j^2}}, \frac{\sigma_e^2}{X'_j X_j + \frac{\sigma_e^2}{\sigma_j^2}}\right) & \delta_j = 1, \\ \sim N\left(\frac{X'_j w}{X'_j X_j + \tilde{\lambda}}, \frac{\tilde{\sigma}_e^2}{X'_j X_j + \tilde{\lambda}}\right) & \delta_j = 0. \end{cases}$$

When $\delta_j = 1$, the full conditional distribution of σ_j^2 is the same as that in the single-site Gibbs sampler. When $\delta_j = 0$, (18) is the only part that contains σ_j^2 , which means it should be sampled from its prior. Thus when $\delta_j = 1$, the full conditional distribution of σ_j^2 is a scaled inverted chi-square distribution with scale parameter $\frac{\nu_\beta S_\beta^2 + \beta_j^2}{\nu_\beta + 1}$ and degrees of freedom parameter $\nu_\beta + 1$; when $\delta_j = 0$, it is a scaled inverted chi-square distribution with scale parameter S_β^2 and degrees of freedom parameter ν_β . The full conditional distribution of δ_j can be obtained from the joint distribution of all parameters and y by dropping all factors free of δ_j , which gives

$$\Pr(\delta_j = 1 | ELSE) \propto f(y | \mu, \beta, \delta, \xi, \sigma_e^2) \Pr(\delta_j = 1) f(\beta_j | \delta_j = 1),$$

$$\Pr(\delta_j = 0 | ELSE) \propto f(y | \mu, \beta, \delta, \xi, \sigma_e^2) \Pr(\delta_j = 0) f(\beta_j | \delta_j = 0).$$

Compared to the full conditional distributions for δ_j in the single-site Gibbs sampler, the difference is the extra factor $f(\beta_j | \delta_j = 1)$ and $f(\beta_j | \delta_j = 0)$, because they cannot be canceled out as in the single-site Gibbs sampler. Thus,

$$\Pr(\delta_j = 1 | ELSE) = \frac{(1 - \pi) d_j f(\beta_j | \delta_j = 1)}{(1 - \pi) d_j f(\beta_j | \delta_j = 1) + \pi \exp\left(-\frac{w' w}{2\sigma_e^2}\right) f(\beta_j | \delta_j = 0)},$$

where $f(\beta_j | \delta_j = 1) = \frac{1}{\sqrt{2\pi}(\sigma_j^2)^{\frac{1}{2}}} \exp\left(-\frac{\beta_j^2}{2\sigma_j^2}\right)$ and

$$f(\beta_j | \delta_j = 0) = \frac{1}{\sqrt{2\pi} \left(\frac{\tilde{\sigma}_e^2}{X'_j X_j + \tilde{\lambda}}\right)^{\frac{1}{2}}} \exp\left[-\frac{\left(\beta_j - \frac{X'_j w}{X'_j X_j + \tilde{\lambda}}\right)^2}{2 \frac{\tilde{\sigma}_e^2}{X'_j X_j + \tilde{\lambda}}}\right].$$

Data

Real genotypic data and simulated phenotypic data were used here to compare BayesB using MH, efficient MH or the three different Gibbs samplers as described above. The genotypic data included 3961 individuals with 55,734 SNPs. The heritability of the simulated trait was 0.25. The training data contained 3206 individuals and the remaining individuals were used for testing. A chain of length of 50,000 was used to estimate parameters of interest. Prediction accuracies were calculated using different samplers. The effective sample sizes [11], which estimate the number of independent samples from a chain, were calculated for σ_e^2 to compare convergence rates for different methods. Computing time for different methods with the same number of iterations were also compared.

Results

The number of effective samples per second of computing time was obtained for BayesB using MH, efficient MH or the three different Gibbs samplers. These three Gibbs samplers were almost twice as efficient as Metropolis–Hastings (Table 1). The prediction accuracies for different samplers, which are calculated as the correlation between

estimated breeding values and simulated phenotypes, are all equal to 0.296. Posterior means of μ for these four samplers are all equal to 2.508. Posterior means of σ_e^2 for different samplers are almost equal, ranging from 0.955 to 0.957.

Discussion

In the joint Gibbs sampler, δ_j and β_j are sampled jointly, which addresses the problem of dependence between δ_j and β_j . Thus, the joint sampler had the largest effective sample size. However, in the single-site Gibbs sampler, δ_j and β_j are sampled from their full conditionals, and thus due to the dependence between β_j and δ_j , the single-site Gibbs sampler had the smallest effective sample size. These differences in effective sample size, however, were small.

In the Gibbs sampler with pseudo priors, β_j and δ_j are also sampled from their full conditionals. Recall that we have shown that the posterior mean of the marker effects does not depend on the pseudo prior. Furthermore, Godsill [8] has shown that the marginal posterior for parameters in the model do not depend on the pseudo prior, which is the prior for β_j when $\delta_j = 0$. As suggested by Carlin and Chib [7], when the full conditional distribution of β_j when $\delta_j = 1$ is chosen to be the pseudo prior, we have shown that the samples of β_j and δ_j are identically distributed to those from the joint Gibbs sampler. Thus, the Gibbs sampler with pseudo priors will have a similar effective sample size as the joint Gibbs sampler.

However, when the full conditional distribution for β_j when $\delta_j = 1$ is used as the pseudo prior in BayesB, the full conditional distributions of σ_e^2 and σ_j^2 are not of known forms because σ_e^2 and σ_j^2 are in the pseudo prior for the marker effect. In contrast to BayesB, in the model used by Godsill [8] to justify the use of full conditional distributions as the pseudo priors, for simplicity, hyper-parameters such as σ_e^2 were omitted [8]. Here, we have replaced σ_e^2 and σ_j^2 in the pseudo prior with constants such that the full conditionals for σ_e^2 and σ_j^2 have scaled inverted chi-square distributions. This modification will give a pseudo prior whose distribution is close to that of the full conditional. In the Gibbs sampler with this pseudo prior, the effective sample size was smaller than in the joint Gibbs sampler but still larger than in the single-site Gibbs sampler.

Conclusions

When a MH algorithm is used to jointly sample the marker effect and the locus-specific variance, the BayesB method is computationally intensive. After introducing

a variable δ_j , indicating whether the marker effect for a locus is zero or non-zero, the marker effect and locus-specific variance can be sampled using Gibbs sampler without MH. Among the Gibbs samplers that were considered here, the joint Gibbs sampler is the most efficient. This sampler is about 2.1 times as efficient as the MH algorithm proposed by Meuwissen et al. [4] and 1.7 times as efficient as that proposed by Habier et al. [6].

Authors' contributions

HC, LQ, DJG, RLF contributed to the development of the statistical methods. HC wrote the program code and conducted the analyses. The manuscript was prepared by RLF and HC. All authors read and approved the final manuscript.

Author details

¹ Department of Animal Science, Iowa State University, Ames 50011, IA, USA. ² Department of Statistics, Iowa State University, Ames, IA 50011, USA. ³ Department of Mathematics and Statistics, Wright State University, Dayton, OH 45435, USA. ⁴ Institute of Veterinary, Animal and Biomedical Science, Massey University, Palmerston North, New Zealand.

Acknowledgements

This work was supported by the US Department of Agriculture, Agriculture and Food Research Initiative National Institute of Food and Agriculture Competitive Grant No. 2012-67015-19420 and by National Institutes of Health Grant R01GM099992.

Competing interests

The authors declare that they have no competing interests.

Received: 27 February 2015 Accepted: 23 September 2015

Published online: 14 October 2015

References

- Gianola D. Priors in whole-genome regression: the bayesian alphabet returns. *Genetics*. 2013;157:1819–29.
- Norris JR. Markov chains. Cambridge series on statistical and probabilistic mathematics. New York: Cambridge University Press; 1997.
- Sorensen DA, Gianola D. Likelihood, Bayesian, and MCMC methods in quantitative genetics. New York: Springer; 2002.
- Meuwissen THE, Hayes BJ, Goddard ME. Prediction of total genetic value using genome-wide dense marker maps. *Genetics*. 2001;157:1819–29.
- Gianola D, de los Campos G, Hill WG, Manfredi E, Fernando R. Additive genetic variability and the bayesian alphabet. *Genetics*. 2009;183:347–63.
- Habier D, Fernando RL, Kizilkaya K, Garrick DJ. Extension of the bayesian alphabet for genomic selection. *BMC Bioinformatics*. 2011;12:186.
- Carlin BP, Chib S. Bayesian model choice via Markov chain Monte Carlo methods. *J R Stat Soc Series B Stat Methodol*. 1995;57:473–84.
- Godsill SJ. On the relationship between Markov chain Monte Carlo methods for model uncertainty. *J Comput Graph Stat*. 2001;10:230–48.
- Habier D, Fernando RL, Kizilkaya K, Garrick DJ. Extension of the Bayesian alphabet for genomic selection. In: Proceedings of the 9th world congress on genetics applied to livestock production, 1–6 August 2010, Leipzig, 2010. <http://www.kongressband.de/wcgalp2010/assets/pdf/0468.pdf>.
- Fernando RL, Garrick D. Bayesian methods applied to GWAS. *Methods Mol Biol*. 2013;1019:237–74.
- Geyer CJ. Practical Markov chain Monte Carlo. *Stat Sci*. 1992;7:473–511.