

A Fast Linear Regression via SVD and Marginalization

Philip Greengard*, Andrew Gelman†, Aki Vehtari‡

November 9, 2020

Abstract

We describe a numerical scheme for evaluating the posterior moments of Bayesian linear regression models with partial pooling of the coefficients. The principal analytical tool of the evaluation is a change of basis from coefficient space to the space of singular vectors of the matrix of predictors. After this change of basis and an analytical integration, we reduce the problem of finding moments of a density over $k + m$ dimensions, to finding moments of an m -dimensional density, where k is the number of coefficients and $k + m$ is the dimension of the posterior. Moments can then be computed using, for example, MCMC, the trapezoid rule, or adaptive Gaussian quadrature. An evaluation of the SVD of the matrix of predictors is the dominant computational cost and is performed once during the precomputation stage. We demonstrate numerical results of the algorithm. The scheme described in this paper generalizes naturally to multilevel and multi-group hierarchical regression models where normal-normal parameters appear.

1 Introduction

Linear regression is a ubiquitous tool for statistical modeling in a range of applications including social sciences, epidemiology, biochemistry, and environmental sciences ([Gelman et al., 2013], [Gelman and Hill, 2007], [Greenland, 2000], [Merlo et al., 2005], [Bardini et al., 2017]).

A common bottleneck for applied statistical modeling workflow is the computational cost of model evaluation. Since posterior distributions in statistical models are often high dimensional and computationally intractable, various techniques have been used to approximate posterior moments. Standard approaches often involve a variety of techniques including Markov chain Monte Carlo (MCMC) or using a suitable approximation of the posterior.

In this paper, we describe an approach for reducing the computational costs for a particular class of regression models — those that contain parameters $\theta \in \mathbb{R}^k$ such that θ has a normal prior and normal likelihood. These models represent only a subset of regression models that appear in applications. We focus our attention in this paper on normal-normal models because they have well known analytical properties and are more computationally tractable than the vast majority of multilevel models. A broader class

*Columbia University

†Department of Statistics and Department of Political Science, Columbia University

‡Department of Computer Science, Aalto University

of models, including logistic regression, contain distributions that are less amenable to the techniques of this paper and will require other analytical and computational tools. Mathematically, marginalization of normal-normal parameters is well-known and has been applied to the posterior by, for example, [Lindley and Smith, 1972]. Our contribution is to provide a stable, accurate, and fast algorithm for marginalization.

The primary numerical tool used in the algorithm is the singular value decomposition (SVD) of the data matrix. As a mathematical and statistical tool, SVD has been known since at least 1936 (see [Eckart and Young, 1936]). Use of the SVD as a practical and efficient numerical algorithm only started gaining popularity much later, with the first widely used scheme introduced in [Golub and Kahan, 1965]. Due in large part to advances in computing power, use of the SVD as a tool in applied mathematics, statistics, and data science has been gaining significant popularity in recent years, however efficient evaluation of SVDs and related matrix decompositions is still an active area of research (see [Hastie et al. 2015], [Halko et al., 2011], [Shamir et al., 2016]).

Similar schemes to ours are used in the software packages lme4 [Bates et al., 2015] and INLA [Rue et al., 2017]. There are several differences between the problems they address and their computational techniques, and those that we shall discuss here. While lme4 finds maximum likelihood and restricted maximum likelihood estimates, our goal is to find posterior moments. The software package INLA uses Laplace approximation on the posterior for a general choice of likelihood functions, whereas our algorithm is focused on fast and accurate solutions for only a particular class of densities: those with normal-normal parameters.

The approach presented in this paper analytically marginalizes the normal-normal parameters of a model using a change of variables. After marginalization, posterior moments can be computed using standard techniques on the lower-dimensional density. In particular, for a model that contains $k + m$ total variables, k of which are normal-normal, our scheme converts the problem of evaluating expectations of a density in $k + m$ dimensions to finding expectations of an m -dimensional density. After marginalization, we evaluate the m -dimensional posterior density in $O(k)$ operations. Without the change of variables, standard evaluation of marginal densities that relies on determinant evaluation requires at least $O(k^3)$ operations.

We illustrate our scheme on the problem of evaluating the marginal expectations of the unnormalized density

$$q(\sigma_1, \sigma_2, \beta) = \sigma_1^{-(k+1)} \sigma_2^{-n} \exp\left(-\gamma \log^2(\sigma_1) - \frac{\sigma_2^2}{2} - \frac{\|X\beta - y\|^2}{2\sigma_2^2} - \frac{\|\beta\|^2}{2\sigma_1^2}\right), \quad (1)$$

where $\gamma > 0$ is a constant, $\sigma_1, \sigma_2 > 0$, and $\beta \in \mathbb{R}^k$. We assume that X is a fixed $n \times k$ matrix, $y \in \mathbb{R}^n$ is fixed, and the normalizing constant of [1] is unknown. For fixed $n, k \in \mathbb{N}$, the algorithm is nearly identical when X is an $n \times k$ matrix to when X is a $k \times n$ matrix. In the case where $k \gg n$, see [Kwon et al., 2011] for a similar approach. Using the standard notation of Bayesian models, density q is the unnormalized posterior of the model

$$\begin{aligned} \sigma_1 &\sim \text{lognormal}(0, \sqrt{\gamma}) \\ \sigma_2 &\sim \text{normal}^+(0, 1) \\ \beta &\sim \text{normal}(0, \sigma_1) \\ y &\sim \text{normal}(X\beta, \sigma_2). \end{aligned} \quad (2)$$

In Appendix A, we include Stan code that can be used to sample from density [\(1\)](#).

Statistical model [\(2\)](#) is a standard model of Bayesian statistics and appears when seeking to model an outcome, y , as a linear combination of related predictors, the columns of X . In [Gelman and Hill, 2007](#), these models are described in detail and are used in the estimation of the distribution of radon levels in houses in Minnesota.

Density [\(1\)](#) is also closely related to posterior densities that appear in genome-wide association studies (GWAS; see [Zhu and Stephens, 2017](#), [Meuwissen, et al., 2001](#), [Azevedo et al., 2015](#)) which can be used to identify genomic regions containing genes linked with a specific trait, such as height. Using the notation of [\(1\)](#), each row of matrix X corresponds to a person, each column of X represents a genomic location, entries of X indicate genotypes, and y corresponds to the trait. Due to technical advances in genome sequencing over the last ten years, it is now feasible to collect large amounts of sequencing data. GWAS models can contain data on up to millions of people and often between hundreds and thousands of genome locations (see [Linner et al., 2019](#)). As a result, efficient computational tools are required for model evaluation.

The number of operations required by the scheme of this paper scales like $O(nk^2)$ with a small constant. The key analytical tool is a change of variables of β such that the terms,

$$-\frac{1}{2\sigma_2^2}\|X\beta - y\|^2 - \frac{1}{2\sigma_1^2}\|\beta\|^2, \tag{3}$$

in [\(1\)](#) are converted to a diagonal quadratic form in \mathbb{R}^k . After that change of variables, expectations over q are analytically converted from integrals over \mathbb{R}^{k+2} to integrals over \mathbb{R}^2 . The remaining 2-dimensional integrals can be computed to high accuracy using classical numerical techniques including, for example, adaptive Gaussian quadrature or even the 2-dimensional trapezoid rule.

The schemes used to evaluate the expectations of [\(1\)](#) generalize naturally to evaluation of expectations of multilevel and multigroup posterior distributions including, for example, the two-group posterior of the form,

$$q(\sigma_1, \sigma_2, \sigma_3, \mu, \beta) = \exp\left(-\frac{1}{2\sigma_1^2}\|X\beta - y\|^2 - \frac{1}{2\sigma_2^2}\sum_{i=1}^{k_1}(\mu - \beta_i)^2 - \frac{1}{2\sigma_3^2}\sum_{i=k_1+1}^{k_1+k_2}\beta_i^2\right), \tag{4}$$

where X is a $n \times k$ matrix, $y \in \mathbb{R}^n$, k_1 and k_2 are non-negative integers satisfying $k_1 + k_2 = k$, and the vector $t \in \mathbb{R}^m$.

For models where m is large, MCMC can be used to evaluate the m -dimensional expectations, with, for example, Stan [Carpenter et al., 2017](#). The m -dimensional distribution has two qualities that make it preferable to its high-dimensional counterpart. First, it requires only $O(k)$ operations to evaluate the integrand, and second, the geometry of the m dimensional marginal distribution will allow for more efficient sampling.

The structure of this paper is as follows. In the following section we describe the analytic integration that transforms [\(1\)](#) from a $k + 2$ -dimensional problem to a 2-dimensional problem. Section [3](#) includes formulas that will allow for the evaluation of posterior moments using the 2-dimensional density. In Sections [4](#) and [5](#) we provide formulas for evaluating covariances of [\(1\)](#). In Section [6](#) we discuss the numerical results

of the implementation of the algorithm. Conclusions and generalizations of the algorithm of this paper are presented in Section 7 Appendix A provides Stan code that can be used to sample from (1), and Appendix B includes proofs of the formulas of this paper.

2 Analytic Integration of β

In this section, we describe how we analytically marginalize the normal-normal parameter β of density (1). We include proofs of all formulas in Appendix B.

We start by marginalizing β using a change of variables that converts the quadratic forms in (1) into diagonal quadratic forms. The resulting integral in the new variable, z , is Gaussian, and the coefficients of z_i and z_i^2 are available analytically. The change of variables is given by the right orthogonal matrix of the singular value decomposition (SVD) of X . That is, we set

$$z = V^t \beta \tag{5}$$

where the SVD of X is

$$X = UDV^t. \tag{6}$$

We define λ_i to be the i^{th} element of the diagonal of D . The elements of diagonal need not be sorted. After this change of variables, we obtain the following identity for the last two terms of (1). A proof can be found in Lemma B.5 in Appendix B.

Formula 2.1.

$$-\frac{1}{2\sigma_2^2} \|X\beta - y\|^2 - \frac{1}{2\sigma_1^2} \|\beta\|^2 = a_0 + \sum_{i=1}^k a_{2,i} \left(z_i - \frac{a_{1,i}}{2a_{2,i}} \right)^2 + \frac{a_{1,i}^2}{4a_{2,i}} \tag{7}$$

where

$$a_{2,i} = \frac{\lambda_i^2}{2\sigma_2^2} + \frac{1}{2\sigma_1^2}, \tag{8}$$

$$a_{1,i} = \frac{w_i}{\sigma_2^2}, \tag{9}$$

and

$$a_0 = -\frac{y^t y}{2\sigma_2^2} \tag{10}$$

where

$$w = V^t X^t y. \tag{11}$$

After performing the change of variables $z = V^t \beta$ and using (7), we now have an expression for density (1) in a form that allows us to use the well-known properties of a Gaussian with diagonal covariance. The following identity uses these properties and provides a formula for analytically reducing expectations of (1) from integrals over $k + 2$ dimensions to integrals over 2 dimensions. After the formula is applied, we have a new density, \tilde{q} , over only 2 dimensions. See Theorem B.6 in Appendix B for a proof.

Formula 2.2. For all $\sigma_1, \sigma_2 > 0$ we have

$$\int_{\mathbb{R}^k} q(\sigma_1, \sigma_2, \beta) d\beta = \tilde{q}(\sigma_1, \sigma_2) \quad (12)$$

where $\tilde{q}(\sigma_1, \sigma_2)$ is defined by the formula

$$\tilde{q}(\sigma_1, \sigma_2) = \sigma_1^{-(k+1)} \sigma_2^{-n} \exp\left(-\gamma \log^2(\sigma_1) - \frac{\sigma_2^2}{2} + a_0 + \sum_{i=1}^k \frac{a_{1,i}^2}{4a_{2,i}}\right) \prod_{i=1}^k \frac{1}{\sqrt{2a_{2,i}}} \quad (13)$$

where $a_{2,i}$ is defined in [8], $a_{1,i}$ is defined in [9], a_0 is defined in [10], and γ is a constant.

Remark 2.1. Certain Bayesian models might contain correlated priors on β that will result in posteriors such as [28] of Section 4. For such models, we perform the change of variables that uses the fact that two diagonal forms over β can be simultaneously diagonalized.

We include in Figure 1 a plot of the density of q as a function of σ_1 and β_1 for fixed σ_2 and randomly chosen X and y . Figure 2 shows a plot of q as a function of σ_2 and β for fixed σ_1 . Figure 3 provides an illustration of \tilde{q} , obtained after the change of variables and marginalization described in this section.

3 Evaluation of Posterior Means

Now that we have reduced the $k + 2$ -dimensional density q to the 2-dimensional density \tilde{q} , it remains to recover the posterior moments of q using \tilde{q} . We first observe that moments of σ_1 and σ_2 with respect to q are equivalent to moments of σ_1 and σ_2 over \tilde{q} . That is,

$$\mathbb{E}_q(\sigma_1) = \mathbb{E}_{\tilde{q}}(\sigma_1) \quad (14)$$

and

$$\mathbb{E}_q(\sigma_2) = \mathbb{E}_{\tilde{q}}(\sigma_2). \quad (15)$$

As for moments of β , we use [13] and standard properties of Gaussians to obtain the following formula.

Formula 3.1. For all $\sigma_1, \sigma_2 > 0$,

$$\int_{\mathbb{R}^k} z_i q(\sigma_1, \sigma_2, \beta) d\beta = \frac{a_{1,i}}{2a_{2,i}} \tilde{q}(t) \quad (16)$$

where q is defined in [1], \tilde{q} is defined in [13], $a_{2,i}$ is defined in [8], and $a_{1,i}$ is defined in [9].

As an immediate consequence of [16], we are able to evaluate the posterior expectation of z as an expectation of a 2-dimensional density:

$$\mathbb{E}_q(z_i) = \mathbb{E}_{\tilde{q}}\left(\frac{a_{1,i}}{2a_{2,i}}\right). \quad (17)$$




Figure 1: *Density of q (see (1)) with respect to σ_1 and β_1 , where $\gamma = 8$, $n = 100$, $k = 10$, and data were randomly generated.*




Figure 2: *Density of q (see (1)) with respect to σ_2 and β_1 , for the same parameters as Figure 1.*

We then transform those expectations back to expectations over the desired basis, β using the matrix V computed in (6). Specifically, using linearity of expectation and (17), we




Figure 3: Density of \tilde{q} (see (13)) using the same q as Figure 1, where $n = 100$, $k = 10$, and data were randomly generated.

know

$$\begin{aligned}
 \mathbb{E}_q((\beta_1, \dots, \beta_k)^t) &= \mathbb{E}_q(VV^t(\beta_1, \dots, \beta_k)^t) \\
 &= V\mathbb{E}_q(V^t(\beta_1, \dots, \beta_k)^t) \\
 &= V\mathbb{E}_q((z_1, \dots, z_k)^t) \\
 &= V\mathbb{E}_{\tilde{q}}\left(\left(\frac{a_{1,1}}{2a_{2,1}}, \dots, \frac{a_{1,k}}{2a_{2,k}}\right)^t\right).
 \end{aligned} \tag{18}$$

4 Covariance of β

In addition to facilitating the rapid evaluation of posterior means, the change of variables described in Section 2 is also useful for the evaluation of higher moments.

Equation (7) shows that after the change of variables from β to z , the resulting density is a Gaussian in z with a diagonal covariance matrix. Additionally, for each z_i , using equation (7) and standard properties of Gaussians, we have the following identity.

Formula 4.1. For all $\sigma_1, \sigma_2 > 0$, we have

$$\int_{\mathbb{R}^k} (z_i - \mu_{z_i})^2 q(\sigma_1, \sigma_2, \beta) d\beta = (2a_{2,i})^{-1} \tilde{q}(\sigma_1, \sigma_2) \tag{19}$$

where μ_{z_i} is the expectation of z_i , \tilde{q} is defined in (13), and $a_{2,i}$ is defined in (8).

The second moments of the posterior of β are obtained as a linear transformation of the posterior variances of z . In particular, denoting the expectation of β by μ_β and the expectation of z by μ_z , we have

$$\begin{aligned}\mathbb{E}((\beta - \mu_\beta)(\beta - \mu_\beta)^t) &= VV^t\mathbb{E}((\beta - \mu_\beta)(\beta - \mu_\beta)^t)VV^t \\ &= V\mathbb{E}(V^t(\beta - \mu_\beta)(\beta - \mu_\beta)^tV)V^t \\ &= V\mathbb{E}((z - \mu_z)(z - \mu_z)^t)V^t.\end{aligned}\tag{20}$$

We observe that due to the independence of all z_i ,

$$\mathbb{E}((z - \mu_z)(z - \mu_z)^t)\tag{21}$$

is diagonal and we can therefore evaluate the $k \times k$ posterior covariance matrix of β by evaluating $\text{var}(z_i)$ for $i = 1, \dots, k$ and then applying two orthogonal matrices. Specifically, combining Formula 4.1 and (20), we obtain

$$\text{cov}(\beta) = V\mathbb{E}_{\tilde{q}}\left(\left((2a_{2,1})^{-1}, \dots, (2a_{2,k})^{-1}\right)^t\right)V^t.\tag{22}$$

5 Variance of σ_1 and σ_2

Higher moments of σ_1 and σ_2 with respect to q can be evaluated directly as higher moments of σ_1 and σ_2 with respect to \tilde{q} . That is, for all $j \in \{2, 3, \dots\}$, we have

$$\mathbb{E}_q((\sigma_1 - \mu_{\sigma_1})^j) = \mathbb{E}_{\tilde{q}}((\sigma_1 - \mu_{\sigma_1})^j)\tag{23}$$

and

$$\mathbb{E}_q((\sigma_2 - \mu_{\sigma_2})^j) = \mathbb{E}_{\tilde{q}}((\sigma_2 - \mu_{\sigma_2})^j).\tag{24}$$

In particular, for $j = 2$, we obtain

$$\text{var}_q(\sigma_1) = \text{var}_{\tilde{q}}(\sigma_1)\tag{25}$$

and

$$\text{var}_q(\sigma_2) = \text{var}_{\tilde{q}}(\sigma_2).\tag{26}$$

Algorithm 1: *Evaluation of posterior expectations of normal-normal models*

- 1 Compute SVD of matrix X
 - 2 Compute w (see (11))
 - 3 Compute $V^t\mathbf{1}$ (see (9))
 - 4 Construct evaluator for density \tilde{q} of (13)
 - 5 Evaluate first and second moments with respect to \tilde{q} : $\mathbb{E}_{\tilde{q}}(\sigma_1), \mathbb{E}_{\tilde{q}}(\sigma_2), \mathbb{E}_{\tilde{q}}(\frac{a_{1,i}}{2a_{2,i}})$
 - 6 Compute $\mathbb{E}(\beta)$ via formula (18)
-

6 Numerical Experiments

Algorithm [1](#) was implemented in Fortran. We used the GFortran compiler on a 2.6 GHz 6-Core Intel Core i7 MacBook Pro. All examples were run in double precision arithmetic. The matrix X and vector y were randomly generated as follows. Each entry of X was generated with an independent Gaussian with mean 0 and variance 1. The vector y was created by first randomly generating a vector $\beta \in \mathbb{R}^k$, each entry of which is an independent Gaussian with mean 0 and variance 1. The vector y was set to the value of $X\beta + \epsilon$ where $\epsilon \in \mathbb{R}^n$ contains standard normal iid entries. We generated y this way in order to ensure that the $\mathbb{E}(\beta_i)$ were not all small in magnitude. We set γ of [\(1\)](#) to 8.

In Table [2](#) and Figure [5](#), we compare the performance of Algorithm [1](#) to two alternative schemes for computing posterior expectations — one in which we analytically marginalize via equation [\(12\)](#) and then integrate the 2-dimensional density via MCMC using Stan. In the other, we use Stan’s MCMC sampling on the full $k + 2$ dimensional posterior. When using MCMC with Stan, we took 10,000 posterior draws. In Table [2](#) and Figure [5](#) we denote Algorithm [1](#) by “SVD-Trap”. The algorithm that uses Stan on the marginal 2-dimensional density is labeled “SVD-MCMC”, and “MCMC” corresponds to the algorithm that uses only MCMC sampling in Stan.

In the appendix, we include Stan code to sample from the marginal density \tilde{q} of [\(13\)](#).

Remark 6.1. In the numerical integration stage of algorithm [1](#), we use the trapezoid rule with 200 nodes in each direction. Because the integrand is smooth and vanishes near the boundary, convergence of the integral is super-algebraic when using the trapezoid rule (see [Stoer and Bulirsch, 1992](#)). A rectangular grid with 200 points in each direction is satisfactory for obtaining approximately double precision accuracy. In problems with large numbers of non-normal-normal parameters, MCMC algorithms such as Hamiltonian Monte Carlo or other methods can be used.

In Tables [1](#) and [2](#), n and k represent the size of the $n \times k$ random matrix X .

The column labeled “max error” provides the maximum absolute error of the expectations of σ_1 , σ_2 , and β_i for $i \in \{1, 2, \dots, n\}$. The true solution was evaluated using trapezoid rule with 500 nodes in each direction in extended precision.

In Table [1](#), “Precompute time (s)” denotes the time in seconds of all computations until numerical integration. These times are dominated by the cost of SVD [\(36\)](#). The total time of the numerical integration in addition to the matrix-vector product [\(18\)](#) is given in “integrate time (s).” The final column of Table [1](#) “total time (s)”, provides the total time of precomputation and integration.

7 Generalizations and Conclusions

In this paper, we present a numerical scheme for the evaluation of the expectations of a particular class of distributions that appear in Bayesian statistics; posterior distributions of linear regression problems with normal-normal parameters.

The scheme presented generalizes naturally to several classes of distributions that appear frequently in Bayesian statistics. We list several examples of posteriors whose expectations can be evaluated using this method.

n	k	max error	precompute time (s)	integrate time (s)	total (s)
50	5	0.22×10^{-13}	0.01	0.01	0.02
100	10	0.26×10^{-13}	0.02	0.01	0.03
500	20	0.30×10^{-13}	0.04	0.01	0.05
1000	50	0.34×10^{-13}	0.09	0.03	0.12
5000	100	0.37×10^{-13}	0.29	0.05	0.34
10000	500	0.26×10^{-13}	14	0.3	14.2
10000	1000	0.39×10^{-13}	54	0.6	54.5

Table 1: *Scaling of computation times for evaluation of expectations of q (see (1)) using Algorithm 1*




Figure 4: *Scaling of computation times for evaluation of posterior expectations of q (see (1)) using Algorithm 1 as a function of k with $n = 10,000$.*

n	k	SVD-Trap		SVD-MCMC		MCMC	
		time (s)	error	time (s)	error	time (s)	error
100	100	0.16	0.9×10^{-14}	11	0.4×10^{-4}	16	0.1×10^{-1}
200	100	0.16	0.9×10^{-14}	11	0.3×10^{-2}	24	0.8×10^{-2}
500	100	0.23	0.9×10^{-13}	12	0.2×10^{-2}	40	0.8×10^{-2}
1000	100	0.25	0.2×10^{-13}	12	0.6×10^{-3}	88	0.7×10^{-2}
5000	100	0.30	0.4×10^{-13}	14	0.2×10^{-3}	617	0.3×10^{-2}
10000	100	0.65	0.2×10^{-13}	13	0.4×10^{-3}	2552	0.2×10^{-2}

Table 2: *Scaling of computation times for evaluation of expectations of q (see (1)) using three different algorithms: i) SVD-Trap: Algorithm 1 of this paper, ii) SVD-MCMC: marginalization with MCMC integration of \tilde{q} using Stan, and iii) MCMC: full MCMC integration of q using Stan.*

Timings for $k = 100$




Figure 5: *Scaling of computation times for evaluation of posterior expectations of q (see (1)) as a function of sample size n with $k = 100$. The three algorithms compared are i) SVD-Trap: Algorithm 1 of this paper, ii) SVD-MCMC: marginalization with MCMC integration of \tilde{q} using Stan, and iii) MCMC: full MCMC integration of q using Stan.*

1. The choice of priors for σ_1 , and σ_2 in this document were log normal and half-normal. This choice did not substantially impact the algorithm and can be generalized. Adaptive Gaussian quadrature can be used for the numerical integration step of the algorithm for a more general choice of prior on σ_1 and σ_2 .

2. Multilevel regression problems with more than two levels.

3. Regression problems with multiple groups such as the two-group model with posterior

$$\exp \left(-\frac{1}{2\sigma_1^2} \|X\beta - y\|^2 - \frac{1}{2\sigma_2^2} \sum_{i=1}^{k_1} (\mu_1 - \beta_i)^2 - \frac{1}{2\sigma_3^2} \sum_{i=k_1+1}^{k_1+k_2} \beta_i^2 \right) \quad (27)$$

where X is a $n \times k$ matrix, $y \in \mathbb{R}^n$, and k_1 and k_2 are non-negative integers satisfying $k_1 + k_2 = k$.

3. Regression problems with correlated priors on β :

$$\exp \left(-\frac{1}{2\sigma_2^2} \|X_1\beta - y\|^2 - \frac{1}{2\sigma_1^2} \|X_2\beta\| \right) \quad (28)$$

For regression problems with large numbers of non-normal-normal parameters, marginal expectations can be computed using, for example, MCMC in Stan. For such problems, the algorithm of this paper would convert an MCMC evaluation from $k + m$ dimensions to m dimensions, where k is the number of normal-normal parameters.

8 Acknowledgements

The authors are grateful to Ben Bales and Mitzi Morris for useful discussions.

A Code

The following Stan code allows for sampling from the distribution corresponding to the probability density function proportional to (1).

```
data {
  int n;
  int k;
  vector[n] y;
  matrix[n,k] X;
}
parameters {
  real<lower=0> sigma1;
  real<lower=0> sigma2;
  vector<offset=0, multiplier=sigma1>[k] beta;
}
model {
  y ~ normal(X*beta, sigma2);
  beta ~ normal(0, sigma1);
  sigma1 ~ lognormal(0, 0.25);
  sigma2 ~ normal(0, 1);
}
```

The following Stan program samples from the marginal density \tilde{q} (see (13)). The data input `yty` corresponds to $y^t y$ of (10), `lam` is the vector of singular values of X , and `w` is the vector w in equation (11). We include R code for computing `yty`, `lam`, and `w` after the following Stan code.

```
functions {
  real q_tilde_lpdf(real sig1, real sig2, vector w, vector lam, real yty,
                   int k, int n) {
    vector[min(n,k)] a2 = lam^2/(sig2^2) + 1/(sig1^2);
    real sol = sum(w^2 ./a2)/2/sig2^4 - sum(log(a2))/2 - yty/(2*sig2^2);
    sol += -min(n,k)*log(sig1) - n*log(sig2);
    return sol;
  }
}
data {
  int n;
  int k;
  vector[min(n,k)] w;
  vector[min(n,k)] lam;
  real yty;
  matrix[min(n,k),k] V;
}
```

```

}
parameters {
  real<lower=0> sigma1;
  real<lower=0> sigma2;
}
model {
  sigma1 ~ q_tilde(sigma2, w, lam, yty, k, n);
  sigma1 ~ lognormal(0, 0.25);
  sigma2 ~ normal(0, 1);
}
generated quantities {
  vector[k] beta;
  {
    vector[min(n,k)] zvar = 1 ./((2*(lam^2 ./((2*sigma2^2) + 1/(2*sigma1^2))));
    vector[min(n,k)] zmu = w./sigma2^2 .* zvar;
    vector[min(n,k)] z = to_vector(normal_rng(zmu, sqrt(zvar)));
    beta = V * z;
  }
}

```

The following is a sample of code from R that can be used for the precomputation stage of Algorithm [1](#)

```

udv <- svd(X)
V <- udv$v
lam <- as.vector(udv$d)
w <- t(V) %*% t(X) %*% y
w <- as.vector(w)
yty <- t(y) %*% y
yty <- yty[1]

```

B Proofs

In this appendix, we include proofs of the formulas provided in this paper. For increased readability, this appendix is self-contained.

B.1 Mathematical Preliminaries and Notation

In this section, we introduce notation and elementary mathematical identities that will be used throughout the remainder of this section.

We define $C \in \mathbb{R}$ by the equation

$$C = \int_{\sigma_1 \in \mathbb{R}^+} \int_{\sigma_2 \in \mathbb{R}^+} \int_{\beta \in \mathbb{R}^k} q(\sigma_1, \sigma_2, \beta) d\beta d\sigma_2 d\sigma_1, \quad (29)$$

and define $\mathbb{E}(\sigma_1)$, $\mathbb{E}(\sigma_2)$, and $\mathbb{E}(\beta_i)$ by the formulas

$$\mathbb{E}(\sigma_1) = \frac{1}{C} \int_{\sigma_1 \in \mathbb{R}^+} \int_{\sigma_2 \in \mathbb{R}^+} \int_{\beta \in \mathbb{R}^k} \sigma_1 q(\sigma_1, \sigma_2, \beta) d\beta d\sigma_2 d\sigma_1, \quad (30)$$

$$\mathbb{E}(\sigma_2) = \frac{1}{C} \int_{\sigma_1 \in \mathbb{R}^+} \int_{\sigma_2 \in \mathbb{R}^+} \int_{\beta \in \mathbb{R}^k} \sigma_2 q(\sigma_1, \sigma_2, \beta) d\beta d\sigma_2 d\sigma_1, \quad (31)$$

and

$$\mathbb{E}(\beta_i) = \frac{1}{C} \int_{\sigma_1 \in \mathbb{R}^+} \int_{\sigma_2 \in \mathbb{R}^+} \int_{\beta \in \mathbb{R}^k} \beta_i q(\sigma_1, \sigma_2, \beta) d\beta d\sigma_2 d\sigma_1 \quad (32)$$

for $i \in \{1, 2, \dots, k\}$.

We provide algorithms for the evaluation of (29), (30), (31), and (32).

We will be denoting by $\mathbb{1}$ the vector of ones

$$\mathbb{1} = (1, 1, \dots, 1)^t. \quad (33)$$

We denote the i^{th} component of a vector v by v_i .

The following two well-known identities give the normalizing constant and expectation of a Gaussian distribution.

Lemma B.1. *For all $\sigma_1, \sigma_2 > 0$ we have*

$$\sqrt{2\pi}\sigma = \int_{\mathbb{R}} e^{-\frac{(\beta-\mu)^2}{2\sigma^2}} d\beta \quad (34)$$

Lemma B.2. *For all μ in \mathbb{R} and $\sigma > 0$, we have*

$$\mu\sqrt{2\pi}\sigma = \int_{\mathbb{R}} \beta e^{-\frac{(\beta-\mu)^2}{2\sigma^2}} d\beta \quad (35)$$

B.2 Analytic Integration of β

We denote the singular value decomposition of X by

$$X = UDV^t \quad (36)$$

where U is an orthogonal $n \times k$ matrix, V is an orthogonal $k \times k$ matrix, and D is a $k \times k$ diagonal matrix. We define $z \in \mathbb{R}^k$ by the formula

$$z = V^t \beta. \quad (37)$$

The following lemma, which will be used in the proof of Lemma B.5, gives an expression for the second to last term of the exponent in (1) after a change of variables.

Lemma B.3. *For all $\beta \in \mathbb{R}^k$, and $y \in \mathbb{R}^n$,*

$$-\frac{1}{2\sigma_2^2} \|X\beta - y\|^2 = -\frac{y^t y}{2\sigma_2^2} + \sum_{i=1}^k -\frac{\lambda_i^2}{2\sigma_2^2} z_i^2 + \frac{w_i}{\sigma_2^2} z_i \quad (38)$$

where

$$w = V^t X^t y, \quad (39)$$

z is defined in (37), and λ_i is the i^{th} entry on the diagonal of D (see (36)).

Proof. Clearly,

$$\|X\beta - y\|^2 = \beta^t X^t X \beta - 2y^t X \beta + y^t y. \quad (40)$$

Substituting (36) and (37) into (40), we obtain

$$\begin{aligned} \|X\beta - y\|^2 &= \beta^t (UDV^t)^t (UDV^t) \beta - 2y^t X V V^t \beta + y^t y \\ &= (\beta^t V) D^2 (V^t \beta) - 2y^t (V^t X^t)^t z + y^t y. \end{aligned} \quad (41)$$

where z is defined in (37). Substituting (39) and (37) into (41), we have

$$\|X\beta - y\|^2 = z^t D^2 z - 2w^t z + y^t y \quad (42)$$

Equation (38) follows immediately from (42). \blacksquare

The following lemma provides an equation for the last term of the exponent in (1). The identity will be used in Lemma B.5

Lemma B.4. For all $\sigma_1 > 0$,

$$-\frac{\|\beta\|^2}{2\sigma_1^2} = \sum_{i=1}^k -\frac{z_i^2}{2\sigma_1^2} \quad (43)$$

where $\beta \in \mathbb{R}^k$, z is defined in (37), and V is defined in (36).

Proof. Clearly,

$$\frac{\|\beta\|^2}{2\sigma_1^2} = \frac{1}{2\sigma_1^2} (Vz)^t (Vz) = \frac{z^t z}{2\sigma_1^2} \quad (44)$$

where V is defined in (36). Equation (43) follows immediately from (44). \blacksquare

The following formula combines Lemma B.3 and Lemma B.4 to convert the final two terms of (1) into a Gaussian in k dimensions.

Lemma B.5.

$$-\frac{\|X\beta - y\|^2}{2\sigma_2^2} - \frac{\|\beta\|^2}{2\sigma_1^2} = a_0 + \sum_{i=1}^k a_{2,i} \left(z_i - \frac{a_{1,i}}{2a_{2,i}}\right)^2 + \frac{a_{1,i}^2}{4a_{2,i}} \quad (45)$$

where

$$a_{2,i} = \frac{\lambda_i^2}{2\sigma_2^2} + \frac{1}{2\sigma_1^2}, \quad (46)$$

$$a_{1,i} = \frac{w_i}{\sigma_2^2} \quad (47)$$

and

$$a_0 = -\frac{y^t y}{2\sigma_2^2} \quad (48)$$

where z is defined in (37), w is defined in (39) and V is defined in (36).

Proof. By combining Lemma [B.3](#) and Lemma [B.4](#), we have

$$-\frac{1}{2\sigma_2^2}\|X\beta - y\|^2 - \frac{1}{2\sigma_1^2}\|\beta\|^2 = a_0 + \sum_{i=1}^k (a_{1,i}z_i - a_{2,i}z_i^2). \quad (49)$$

We obtain equation [\(45\)](#) by completing the square in equation [\(49\)](#). \blacksquare

The following theorem is the principal analytical apparatus of this note. It provides a formula for the k -dimensional integrals that appear in [\(29\)](#), [\(30\)](#), and [\(31\)](#).

Theorem B.6. For all $\sigma_1, \sigma_2 > 0$

$$\int_{\mathbb{R}^k} q(\sigma_1, \sigma_2, \beta) d\beta = \tilde{q}(\sigma_1, \sigma_2) \quad (50)$$

where $\tilde{q}(\sigma_1, \sigma_2)$ is defined by the formula

$$\tilde{q}(\sigma_1, \sigma_2) = \sigma_1^{-(k+1)} \sigma_2^{-n} \exp\left(-\log^2(\sigma_1) - \frac{\sigma_2^2}{2} + a_0 + \sum_{i=1}^k \frac{a_{1,i}^2}{4a_{2,i}}\right) \sqrt{2\pi}^k \prod_{i=1}^k \frac{1}{\sqrt{2a_{2,i}}} \quad (51)$$

where $a_{2,i}$ is defined in [\(46\)](#), $a_{1,i}$ is defined in [\(47\)](#) and a_0 is defined in [\(48\)](#).

Proof. Using [\(1\)](#), clearly

$$\int_{\mathbb{R}^k} q(\sigma_1, \sigma_2, \beta) d\beta = \sigma_1^{-(k+1)} \int_{\mathbb{R}^k} \exp\left(-\log^2(\sigma_1) - \frac{\sigma_2^2}{2} - \frac{1}{2\sigma_2^2}\|X\beta - y\|^2 - \frac{1}{2\sigma_1^2}\|\beta\|^2\right) d\beta \quad (52)$$

Performing the change of variables [\(37\)](#) and substituting [\(45\)](#) into [\(52\)](#), we have

$$\int_{\mathbb{R}^k} q(\sigma_1, \sigma_2, \beta) d\beta = \exp\left(-\log^2(\sigma_1) - \frac{\sigma_2^2}{2} + a_0 + \sum_{i=1}^k \frac{a_{1,i}^2}{4a_{2,i}}\right) \int_{\mathbb{R}^k} \exp\left(\sum_{i=1}^k a_{2,i} \left(z_i - \frac{a_{1,i}}{2a_{2,i}}\right)^2\right) dz \quad (53)$$

Since the integrand on the right side of [\(53\)](#) is a Gaussian in z_i , equation [\(50\)](#) follows from applying Lemma [B.1](#) to [\(53\)](#). \blacksquare

The following theorem provides a formula for the expectation of z (see [\(37\)](#)). We use this formula, in combination with an orthogonal transformation, to obtain the expectation of β .

Theorem B.7. For all $\sigma_1 > 0$ and $\sigma_2 \in \mathbb{R}$,

$$\int_{\mathbb{R}^k} (V^t x)_i q(\sigma_1, \sigma_2, \beta) d\beta = \frac{a_{1,i}}{2a_{2,i}} \tilde{q}(t) \quad (54)$$

where q is defined in [\(1\)](#), \tilde{q} is defined in [\(51\)](#), $a_{2,i}$ is defined in [\(46\)](#), $a_{1,i}$ is defined in [\(47\)](#), a_0 is defined in [\(48\)](#).

Proof. Combining (53) and (37), we have

$$\int_{\mathbb{R}^k} (V^t \beta)_i q(\sigma_1, \sigma_2, \beta) d\beta = \exp\left(-\log^2(\sigma_1) - \frac{\sigma_2^2}{2} + a_0 + \sum_{i=1}^k \frac{a_{1,i}^2}{4a_{2,i}}\right) \int_{\mathbb{R}^k} z_i \exp\left(\sum_{i=1}^k a_{2,i} \left(z_i - \frac{a_{1,i}}{2a_{2,i}}\right)^2\right) dz. \quad (55)$$

Applying Lemma B.2 to (55), we obtain (54). ■

References

- [Azevedo et al., 2015] Azevedo, Camila Ferreira et al. “Ridge, Lasso and Bayesian additive-dominance genomic models.” *BMC genetics* 16, 105 (2015)
- [Bardini et al., 2017] Bardini, R. G. Politano, A. Benso, S. Di Carlo. “Multi-level and hybrid modelling approaches for systems biology”. *Computational and Structural Biotechnology Journal*. 15 (2017).
- [Bates et al., 2015] Bates, Douglas, Martin Machler, Ben Bolker, Steve Walker. “Fitting Linear Mixed-Effects Models Using lme4”. *Journal of Statistical Software*. (2015).
- [Carpenter et al., 2017] Carpenter, Bob, Andrew Gelman, Matthew D. Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, Allen Riddell. “Stan: A Probabilistic Programming Language”. *Journal of Statistical Software*. (2017).
- [Eckart and Young, 1936] Eckart, Carl and Gale Young. “The Approximation of One Matrix with Another of Lower Rank”. *Psychometrika*. 1, 3. (1936).
- [Gelman et al., 2013] Gelman, Andrew, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin. *Bayesian Data Analysis*. Third Edition. New York: U.S. CRC, 2013.
- [Gelman and Hill, 2007] Gelman, Andrew and Jennifer Hill. *Data Analysis using Regression and Multilevel/Hierarchical Models* Cambridge: UK. Cambridge University Press, 2007.
- [Golub and Kahan, 1965] Golub, Gene and William Kahan. “Calculating the Singular Values and Psuedo-Inverse of a Matrix.” *J. SIAM Numer. Anal.* 2 (3). (1965)
- [Greenland, 2000] Greenland, Sander. “Principles of Multilevel Modelling”. *International Journal of Epidemiology*. 29 (2000).
- [Halko et al., 2011] Halko, N., P.G. Martinsson, J. A. Tropp. “Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions.” *SIAM Rev.* 53 (2). (2011)

- [Hastie et al. 2015] Hastie, Trevor, Rahul Mazumder, Jason D. Lee, Reza Zadeh. “Matrix Completion and Low-Rank SVD via Fast Alternating Least Squares.” *The Journal of Machine Learning Research*. 16 (1). (2015)
- [Kwon et al., 2011] Kwon, Soonil, Xiaofei Yan, Jinrui Cui, Jie Yao, Kai Yang, Donald Tsiand, Xiaohui Li, Jerome Rotter, Xiuqing Guo. “Application of Bayesian Regression with Singular Value Decomposition method in Association Studies for Sequence Data.” *BMC Proceedings*. 5 (9) (2011).
- [Lindley and Smith, 1972] Lindley, D. V., and Smith, A. F. M. (1972). Bayes estimates for the linear model. *Journal of the Royal Statistical Society B*. 34, 1–41.
- [Linner et al., 2019] Linner, Karlsson, et al. “Genome-wide association analyses of risk tolerance and risky behaviors in over 1 million individuals identify hundreds of loci and shared genetic influences.” *Nature Genetics*. 51, 2. (2019)
- [Merlo et al., 2005] Merlo, J., B. Chaix, M. Yang, J. Lynch, and L. Rastam. “A Brief Conceptual Tutorial of Multilevel Analysis in Social Epidemiology: Linking the Statistical Concept of Clustering to the Idea of Contextual Phenomenon”. *J Epidemiol Community Health*. 59 (2005).
- [Meuwissen, et al., 2001] Meuwissen, T.H., B. J. Hayes and M. E. Goddard. “Prediction of Total Genetic Value Using Genome-Wide Dense Marker Maps.” *Genetics* 157, 4 (2001).
- [Neal, 2003] Neal, Radford. “Slice Sampling”. *Annals of Statistics*. 31 (2003).
- [Rue et al., 2017] Rue, Havard, Andrea Riebler, Sigrunn H. Sorbye, Janine B. Illian, Daniel P. Simpson, Finn K. Lindgren. “Bayesian Computing with INLA: A Review”. *Annual Review of Statistics and Its Application*. 4 (2017).
- [Shamir et al., 2016] Shamir, Ohad. “Fast Stochastic Algorithms for SVD and PCA: Convergence Properties and Convexity.” *Proceedings of the 33rd ICML*. New York, NY. (2016)
- [Stoer and Bulirsch, 1992] Stoer, Josef and Roland Bulirsch. *Introduction to Numerical Analysis*, 2nd ed., Springer-Verlag, 1992.
- [Zhu and Stephens, 2017] Zhu, Xiang and Matthew Stephens. “Bayesian Large-Scale Multiple Regression with Summary Statistics from Genome-Wide Association Studies”. *The Annals of Applied Statistics*. 11, 3 (2017).