

Fast Exact Bayesian Inference for Sparse Signals in the Normal Sequence Model

Tim van Erven^{*,†} and Botond Szabó^{*,†}

Abstract. We consider exact algorithms for Bayesian inference with model selection priors (including spike-and-slab priors) in the sparse normal sequence model. Because the best existing exact algorithm becomes numerically unstable for sample sizes over $n = 500$, there has been much attention for alternative approaches like approximate algorithms (Gibbs sampling, variational Bayes, etc.), shrinkage priors (e.g. the Horseshoe prior and the Spike-and-Slab LASSO) or empirical Bayesian methods. However, by introducing algorithmic ideas from online sequential prediction, we show that exact calculations are feasible for much larger sample sizes: for general model selection priors we reach $n = 25\,000$, and for certain spike-and-slab priors we can easily reach $n = 100\,000$. We further prove a de Finetti-like result for finite sample sizes that characterizes exactly which model selection priors can be expressed as spike-and-slab priors. The computational speed and numerical accuracy of the proposed methods are demonstrated in experiments on simulated data, on a differential gene expression data set, and to compare the effect of multiple hyper-parameter settings in the beta-binomial prior. In our experimental evaluation we compute guaranteed bounds on the numerical accuracy of all new algorithms, which shows that the proposed methods are numerically reliable whereas an alternative based on long division is not.

AMS 2000 subject classifications: Primary 62G05; secondary 62F15.

Keywords: spike-and-slab prior, model selection, high-dimensional statistics.

1 Introduction

In the sparse normal sequence model we observe a sequence $Y = (Y_1, \dots, Y_n)$ that satisfies

$$Y_i = \theta_i + \varepsilon_i, \quad i = 1, \dots, n, \quad (1)$$

for independent standard normal random variables ε_i , where $\theta = (\theta_1, \dots, \theta_n)$ is the unknown signal of interest. It is assumed that the number of non-zero signal components s in θ is small compared to the size of the whole sample (i.e. $s = o(n)$). Applications of this model include detecting differentially expressed genes [45, 36, 39, 56, 23], bankruptcy prediction for publicly traded companies using Altman's Z-score in finance [2, 3], separation of the background and source in astronomical images [18, 28], and wavelet analysis [1, 31]. The model is further of interest to sanity check (approximate) inference methods for the more general sparse linear regression model (see [15] and references therein), which reduces to the normal sequence model when the design is the identity matrix.

^{*}Leiden University, tim@timvanerven.nl, b.t.szabo@math.leidenuniv.nl

[†]Research supported by the Netherlands Organization for Scientific Research: grant numbers VI.Vidi.192.095 for Van Erven and 016.Veni.173.040 for Szabó.

The sparse normal sequence model, which is also called the sparse normal means model, has been extensively studied from a frequentist perspective (see, for instance, [27, 7, 1]), but here we consider Bayesian approaches, which endow θ with a prior distribution. This prior serves as a natural way to introduce sparsity into the model and the corresponding posterior can be used for model comparison and uncertainty quantification (see [24, 53, 37, 5] and references therein). One natural and well-understood class of priors are *model selection priors* that take the following hierarchical form:

- i.) First a sparsity level s is chosen from a prior π_n on $\{0, 1, \dots, n\}$.
- ii.) Then, given s , a subset of nonzero coordinates $\mathcal{S} \subset \{0, 1, \dots, n\}$ of size $|\mathcal{S}| = s$ is selected uniformly at random.
- iii.) Finally, given s and \mathcal{S} , the means $\theta_{\mathcal{S}} = (\theta_i)_{i \in \mathcal{S}}$ corresponding to the nonzero coordinates in \mathcal{S} are endowed with a prior $G_{\mathcal{S}}$, while the remaining coefficients $\theta_{\mathcal{S}^c} = (\theta_i)_{i \notin \mathcal{S}}$ are set to zero.

As is common, we will choose the prior $G_{\mathcal{S}}$ on the nonzero coordinates in a factorized form; i.e. $\theta_i \sim G$ for all $i \in \mathcal{S}$, where G is a fixed one-dimensional prior, which we assume to have a density g (with respect to the Lebesgue measure). Under suitable conditions on π_n and G , the posterior has good frequentist properties and contracts around the true parameter at the minimax rate, as shown by Castillo and Van der Vaart [17]. Notably, they require the prior π_n to decrease at an exponential rate.

A special case of the model selection priors are the spike-and-slab priors developed by Mitchell and Beauchamp [42], George and McCulloch [26], where the coefficients of θ are assigned prior probabilities

$$\begin{aligned} \theta_i \mid \alpha &\sim (1 - \alpha)\delta_0 + \alpha G, & i = 1, \dots, n, \\ \alpha &\sim \Lambda_n, \end{aligned} \tag{2}$$

with δ_0 the Dirac-delta measure at 0 (a *spike*) and G the same one-dimensional prior as above (called the *slab* in this context). The a priori likelihood of nonzero coefficients is controlled by the *mixing parameter* $\alpha \in [0, 1]$, and finally Λ_n is a hyperprior on α . A typical choice for Λ_n is the beta distribution: $\alpha \sim \text{Beta}(\kappa, \lambda)$. In this case the prior on the sparsity level in the model selection formulation takes the form $\pi_n(s) = \binom{n}{s} \frac{B(\kappa+s, \lambda+n-s)}{B(\kappa, \lambda)}$, where $B(\kappa, \lambda)$ denotes the beta function with parameters κ and λ . The resulting prior is called the *beta-binomial prior*. A natural choice is $\kappa = \lambda = 1$ [54], which corresponds to a uniform prior on α , but this choice does not satisfy the exponential decrease condition on π_n . Castillo and Van der Vaart therefore propose $\kappa = 1$ and $\lambda = n + 1$, which does satisfy their exponential decrease condition [17, Example 2.2], and in Section 5 we confirm empirically that the latter indeed leads to better posterior estimates for θ .

Model selection priors set certain signal components to zero, which is desirable for model selection, but makes computation of the posterior difficult since the number of possible sets \mathcal{S} is exponentially large (i.e. 2^n). Castillo and Van der Vaart [17] do provide

an exact algorithm, based on multiplication of polynomials (see Appendix A), but this algorithm runs into numerical problems for sample sizes over $n = 250$ or sometimes $n = 500$ (see Section 3) and it also requires $O(n^3)$ computation steps, which makes it too slow to handle large n .

The computational difficulty of model selection priors has given rise to a variety of alternative priors based on shrinkage. These include the horseshoe prior [11], for which multiple scalable implementations are available [58, 30]. The corresponding posterior achieves the minimax contraction rate and, under mild conditions, also provides reliable uncertainty quantification [61, 60, 59]. The posterior median and draws from the horseshoe posterior are not sparse, but one can use it for model selection after post-processing the posterior. An alternative is to replace the spike in the spike-and-slab prior with a Laplace distribution with very small variance, as in the Spike-and-Slab LASSO [48]. One can efficiently compute the maximum a posteriori (MAP) estimator of the corresponding posterior distribution by convex optimization.

Another way to deal with the computational problems for model selection priors is to consider approximations. The available options include Stochastic Search Variable Selection (SSVS) [26], variational Bayes approximation [65], Langevin Markov Chain Monte Carlo [44], Expectation Maximization [50], Hamiltonian Monte Carlo [53] or empirical Bayes methods [31, 40, 6].

In this paper we return to the goal of exactly computing the posterior for model selection priors, without changing the prior or introducing approximations. In Section 2 we propose a new approach based on a representation of model selection priors by a Hidden Markov Model (HMM) that comes from the literature on online sequential prediction and data compression [64], for which we can apply the standard Forward-Backward algorithm [46]. The computational complexity of this algorithm is $O(n^2)$. To appreciate the speed-up compared to $O(n^3)$ run time, see Section 4, where this method runs in under 15 minutes while the previous algorithm of Castillo and Van der Vaart would take approximately 20 days. Furthermore, in Section 2.2 we specialize to spike-and-slab priors and introduce an even faster algorithm based on a discretization of the α hyper-parameter, which has only $O(n^{3/2})$ run time. Using results from online sequential prediction [20], we show that this discretization provides an accurate approximation of the posterior that can be made exact to arbitrary precision, provided that the density of Λ_n varies sufficiently slowly. Our conditions do not directly allow κ or λ to depend on n in the beta-binomial prior, so we provide an extra result to cover the important case that $\kappa = 1$ and $\lambda = n + 1$.

Our two new approaches allow us to easily handle data sets of size $n = 25\,000$ for general model selection priors and $n = 100\,000$ for the subclass of spike-and-slab priors with sufficiently regular Λ_n , which both substantially exceed the earlier limit of $n = 500$. These results are obtained on a standard laptop within a maximum time limit of half an hour. Run times for larger sample sizes can be estimated by extrapolating from Figure 2.

In Section 2.3 we further derive sufficient and necessary conditions to decide whether a model selection prior can be written in the more efficient spike-and-slab form. Since the distribution of the binary indicators for whether $\theta_i = 0$ or not is exchangeable

under the model selection priors, this amounts to a finite sample de Finetti result for a restricted class of exchangeable distributions.

In Section 3, we demonstrate the scalability and numerical accuracy of the proposed methods on simulated data. We also show there that our deterministic algorithm can be used as a benchmark to test the accuracy of approximation methods: we compare the approximate posterior from Gibbs sampling and variational Bayes to the exact posterior computed by our algorithm, which shows the surprisingly limited number of decimal places to which their answers are reliable. Then, in Section 4, we compare our methods to other approaches suggested in the literature in an application to differential gene expression for Ulcerative Colitis and Crohn's Disease. In Section 5 we further use our new algorithms to empirically investigate the importance of the exponential decrease condition on π_n by varying the hyper-parameters κ and λ of the beta-binomial prior. We find that exponential decrease is not just a sufficient condition for minimax posterior contraction, but it also leads to better posterior estimates of θ . The paper is concluded by Section 6, where we discuss possible extensions of our algorithms.

In addition to the main paper, we provide an accompanying R package that implements our new methods [62], and supplementary material with several appendices [63]. In Appendix A we first recall the exact algorithm by Castillo and Van der Vaart [17]. We show how to resolve its numerical stability issues by performing all intermediate computations in a logarithmic representation. The bottleneck then becomes its computational complexity, because it requires $O(n^3)$ steps, which is prohibitive for large n . Two natural ideas to speed up the algorithm have been proposed by [17, 12], one based on fast polynomial multiplication and one based on long division. Surprisingly, although both approaches look very promising in theory, it turns out that neither of them works well in practice: the theoretical speed-ups for fast polynomial multiplication turn out to be so asymptotic that they do not provide significant gains for any reasonable n ; and the long division approach becomes numerically unstable again. In Appendix B we provide an additional variation on an experiment from Section 5. Finally, Appendix C contains all proofs.

2 Exact Algorithms for Model Selection Priors

In this section we propose novel, exact algorithms for computing (marginal statistics of) the posterior distribution corresponding to model selection priors. For general model selection priors we propose a model selection HMM algorithm, and for spike-and-slab priors we introduce a faster method based on discretization of the α hyper-parameter. The section is concluded with a characterization of the subclass of model selection priors that can be expressed in the more efficient spike-and-slab form.

Marginal Statistics We are interested in computing the marginal posterior probabilities that the coordinates of θ are nonzero:

$$q_{n,i} := \Pi_n(\theta_i \neq 0 \mid Y) \quad \text{for } i = 1, \dots, n.$$

These are sufficient to compute any other marginal statistics of interest, because, conditionally on whether θ_i is 0 or not, the pair (Y_i, θ_i) is independent of all other pairs $(Y_j, \theta_j)_{j \neq i}$. For instance, the marginal posterior means can be expressed as

$$\mathbb{E}[\theta_i | Y] = q_{n,i} \mathbb{E}[\theta_i | Y_i, \theta_i \neq 0] = q_{n,i} \frac{\zeta(Y_i)}{\psi(Y_i)},$$

where $\psi(y) = \int \phi(y - t)g(t) dt$ is the slab density and $\zeta(y) = \int t\phi(y - t)g(t) dt$, with ϕ the standard normal density. We may also obtain marginal quantiles by inverting the marginal posterior distribution functions

$$\Pi_n(\theta_i \leq u | Y) = (1 - q_{n,i})\mathbf{1}_{u \geq 0} + q_{n,i} \frac{\psi(Y_i, u)}{\psi(Y_i)},$$

where $\psi(y, u) = \int_{-\infty}^u \phi(y - t)g(t) dt$. In particular, the marginal medians correspond to

$$\hat{\theta}_i^{\text{med}} = \left[H_{n,i}^{-1} \left(\frac{1}{2q_{n,i}} \right) \wedge 0 \right] + \left[H_{n,i}^{-1} \left(1 - \frac{1}{2q_{n,i}} \right) \vee 0 \right],$$

where $H_{n,i}^{-1}$ is the inverse of the function $H_{n,i}(u) = \frac{\psi(Y_i, u)}{\psi(Y_i)}$ and we use the conventions that $H_{n,i}^{-1}(v) = -\infty$ for $v \leq 0$ and $H_{n,i}^{-1}(v) = \infty$ for $v \geq 1$, see [17].

2.1 The Model Selection HMM Algorithm

Our first computationally efficient approach is based on a Hidden Markov Model (HMM) that comes from the literature on online sequential prediction and data compression [64]. This approach makes it possible to reliably compute all marginal posterior probabilities $q_{n,i}$ in only $O(n^2)$ operations for any model selection prior.

To define the HMM, we will encode the subset of nonzero coordinates $\mathcal{S} \subset \{0, 1, \dots, n\}$ as a binary vector $B = (B_1, \dots, B_n)$, where $B_i = 1$ if $i \in \mathcal{S}$ and $B_i = 0$ otherwise. The crucial observation is that the conditional probabilities of the model selection prior

$$\Pi_n(B_{i+1} | B_1, \dots, B_i) = \Pi_n(B_{i+1} | M_i) \tag{3}$$

only depend on the total number of nonzeros $M_i = \sum_{j=1}^i B_j \in \{0, \dots, i\}$ in the first i coordinates and not on the locations of these coordinates. We can use this observation to interpret the model selection prior as the model selection HMM shown in Figure 1, where each hidden state $H_i = (B_i, M_i)$ contains sufficient information to compute both the transition probabilities

$$P(H_{i+1} | H_i) = \begin{cases} \Pi_n(B_{i+1} | M_i) & \text{if } M_{i+1} = M_i + B_{i+1}, \\ 0 & \text{otherwise,} \end{cases}$$

and the conditional distribution of θ_i given H_i :

$$\begin{aligned} \theta_i &= 0 \quad (\text{a.s.}) && \text{if } B_i = 0, \\ \theta_i &\sim G && \text{if } B_i = 1. \end{aligned}$$

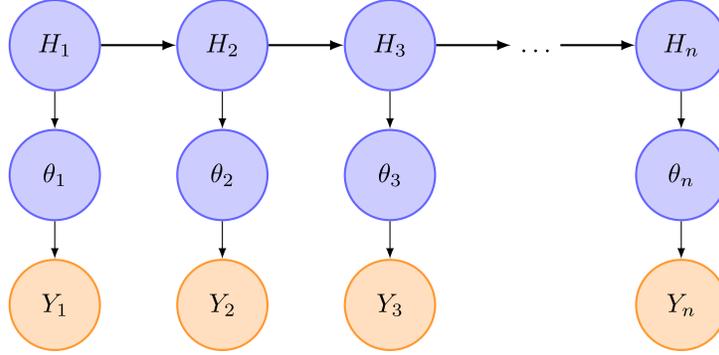


Figure 1: The model selection prior as a Hidden Markov Model.

In fact, in our implementation we will integrate out θ_i to directly obtain the conditional density

$$p(Y_i | H_i) = \begin{cases} \phi(Y_i) & \text{if } B_i = 0, \\ \psi(Y_i) & \text{if } B_i = 1. \end{cases}$$

(Note that $\psi(Y_i)$ is the conditional density of observation Y_i for slabs, while $\phi(Y_i)$ is the density of Y_i for spikes.) Finally, the initial probabilities of H_1 are

$$P(H_1) = \begin{cases} \Pi_n(B_1) & \text{if } M_1 = B_1, \\ 0 & \text{otherwise.} \end{cases}$$

We note that the sequence of hidden states H_1, \dots, H_n is in one-to-one correspondence with \mathcal{S} . Consequently, since the model selection HMM expresses the same joint distribution on H_1, \dots, H_n as the model selection prior, and the conditional distribution of θ and Y given H_1, \dots, H_n is also the same, it follows that the model selection HMM is equivalent to the corresponding model selection prior.

What we gain is that, for HMMs, standard efficient algorithms are available, whose run times depend on the number of state transitions with nonzero probabilities $P(H_{i+1} | H_i)$ [46]. For our purposes, we will use the Forward-Backward algorithm to compute $\Pi_n(H_i | Y)$ for all i in $O(n^2)$ steps, from which we can obtain $q_{n,i} = \Pi_n(B_i = 1 | Y)$ for all i in another $O(n^2)$ steps by marginalizing. For numerical accuracy, we perform all calculations using the logarithmic representation discussed in Appendix A.2.

Let $Y_a^b = (Y_a, \dots, Y_b)$. Then the Forward phase in this algorithm computes the densities $p(Y_1^i, H_i = h_i)$ from $p(Y_1^{i-1}, H_{i-1} = h_{i-1})$ for all $i = 1, \dots, n$ and all values h_i of the hidden states using the recursion

$$p(Y_1^i, h_i) = \begin{cases} p(Y_1 | h_1)P(h_1) & \text{for } i = 1, \\ p(Y_i | h_i) \sum_{h_{i-1}} p(Y_1^{i-1}, h_{i-1})P(h_i | h_{i-1}) & \text{for } 1 < i \leq n. \end{cases}$$

After the Forward phase, the Forward-Backward algorithm performs the Backward phase, which computes $p(Y_{i+1}^n | H_i = h_i)$ from $p(Y_{i+2}^n | H_{i+1} = h_{i+1})$ for all $i = n, \dots, 1$ using the recursion

$$p(Y_{i+1}^n | h_i) = \begin{cases} 1 & \text{for } i = n, \\ \sum_{h_{i+1}} p(Y_{i+2}^n | h_{i+1})p(Y_{i+1} | h_{i+1})P(h_{i+1} | h_i) & \text{for } 1 \leq i < n. \end{cases}$$

Combining the results from the Forward and Backward phases, we can compute

$$\Pi_n(h_i | Y) \propto p(Y_1^i, h_i)p(Y_{i+1}^n | h_i)$$

for all i and h_i as desired.

The HMM described here was introduced by [64] for the Beta(1/2, 1/2)-binomial prior (i.e. the spike-and-slab prior with $\Lambda_n = \text{Beta}(1/2, 1/2)$) in the context of the Switching Method for data compression. See [35] for an overview of many variations on this HMM. Indeed, for any Beta(κ, λ)-binomial prior this HMM is particularly natural, because the transition probabilities of the hidden states have a closed-form expression:

$$\Pi_n(B_{i+1} = 1 | B_1, \dots, B_i) = \Pi_n(B_{i+1} = 1 | M_i) = \frac{\kappa + M_i}{\kappa + \lambda + i}.$$

Here we add the observations that, even when the conditional probabilities $\Pi_n(B_{i+1} | B_1, \dots, B_i)$ are not available in closed form for a given model selection prior, they still satisfy (3) and can be efficiently obtained from

$$\Pi_n(B_{i+1} | B_1, \dots, B_i) = \frac{v_{i+1}(M_i + B_{i+1})}{v_i(M_i)},$$

where $v_i(m) = \Pi_n(B_1 = b_1, \dots, B_i = b_i)$ is the joint probability of any sequence b_1, \dots, b_i with m ones. These joint probabilities can be pre-computed for $i = n, \dots, 1$ in $O(n^2)$ steps using the recursion

$$v_i(m) = \begin{cases} \pi_n(m) / \binom{n}{m} & \text{for } i = n, \\ v_{i+1}(m) + v_{i+1}(m + 1) & \text{for } 1 \leq i < n. \end{cases}$$

Thus we can calculate the marginal posterior probabilities in $O(n^2)$ steps for any model selection prior, not just for beta-binomial priors. The numerical accuracy of this algorithm is demonstrated in Section 3.

2.2 A Faster Algorithm for Spike-and-Slab Priors

In this section we restrict our attention to the spike-and-slab subclass of model selection priors, for which we propose further speed-ups. It is intuitively clear that the mixing hyper-parameter α plays a key role in the behavior of the prior distribution. The optimal choice of α heavily depends on the sparsity parameter s of the model. For instance in case of Cauchy slabs the optimal oracle choice $\alpha = (s/n)\sqrt{\log(n/s)}$ results in minimax

posterior contraction [17] and reliable uncertainty quantification [16] in ℓ_2 -norm. However, in practice the sparsity level s is (typically) not known in advance. Therefore one cannot use the optimal oracle choice for α . In [17] it was also shown that by choosing $\alpha = 1/n$ the posterior contracts around the truth at the nearly optimal rate $s \log(n)$. This seemingly solves the problem of choosing the tuning hyper-parameter. However, a related simulation study in [17] shows that hard-thresholding at the corresponding $\sqrt{2 \log(n)}$ level pairs up with substantially worse practical performance; see Tables 1 and 2 in [17]. Furthermore, in view of [16] the choice of $\alpha = 1/n$ imposes too strong prior assumptions, resulting in overly small posterior spread which leads to unreliable Bayesian uncertainty quantification, i.e. the frequentist coverage of the ℓ_2 -credible set will tend to zero.

Therefore in practice one has to consider a data driven (adaptive) choice of the hyper-parameter α . A computationally appealing approach is the empirical Bayes method, where the maximum marginal likelihood estimator is plugged into the posterior. The corresponding posterior mean achieves a (nearly) minimax convergence rate [31], and for slab distributions with polynomial tails the corresponding posterior contracts around the truth at the optimal rate [13]. However for light-tailed slabs (e.g. Laplace) the empirical Bayes posterior distribution will achieve a highly suboptimal contraction rate around the truth; see again [13].

Another standard (and from a Bayesian perspective more natural) approach is to endow the hyper-parameter α with another layer of prior Λ_n . However, computational problems may arise using standard Gibbs sampling techniques for sampling from the posterior; see Section 3.2 for a demonstration of this problem on a simulated data set. In the literature various speed-ups were proposed. One can for instance focus on relevant sub-sequences of the sequential parameter θ and apply the Gibbs sampler only on them. Another approach is to apply the Hamiltonian Monte Carlo method, see for instance [53]. However, none of these approaches provides an easy way to quantify their approximation error when run for a finite number of iterations. In the next section we propose a deterministic algorithm to approximate the marginal posterior probabilities $q_{n,i}$ for spike-and-slab priors, with a guaranteed bound on its approximation error that can be made arbitrarily close to zero.

Approximation via Discretization of the Mixing Parameter

For general model selection priors the fast HMM algorithm from Section 2.1 requires $O(n^2)$ steps. However, for the special case of spike-and-slab priors we can do even better: we can approximate the corresponding posterior to arbitrary precision using only $O(n^{3/2})$ steps, provided that the density λ_n of the mixing distribution Λ_n on α satisfies certain regularity conditions.

The Algorithm Our approach is to approximate the prior Λ_n by a prior $\tilde{\Lambda}_n$ that is supported on $k = O(n^{1/2})$ discretization points $\alpha_1, \dots, \alpha_k$. Then let Π_n be the original spike-and-slab prior corresponding to a given choice of Λ_n , and let $\tilde{\Pi}_n$ be the prior corresponding to $\tilde{\Lambda}_n$. Conditional on α , the pairs (θ_i, Y_i) are independent. Computing

the likelihood

$$\prod_{i=1}^n \left((1 - \alpha)\phi(Y_i) + \alpha\psi(Y_i) \right)$$

for a single α therefore takes $O(n)$ steps, and consequently we can obtain the posterior probabilities $\tilde{\Pi}_n(\alpha_j | Y)$ of all k discretization points in $O(kn)$ steps. We can then compute

$$\tilde{q}_{n,i} := \tilde{\Pi}_n(\theta_i \neq 0 | Y) = \sum_{j=1}^k \tilde{\Pi}_n(\alpha_j | Y) \frac{\alpha_j \psi(\tilde{Y}_i)}{(1 - \alpha_j)\phi(\tilde{Y}_i) + \alpha_j \psi(\tilde{Y}_i)}$$

in another $O(k)$ steps independently for each i , leading to a total run time of $O(kn) = O(n^{3/2})$ steps. We again perform all calculations using the logarithmic representation from Appendix A.2.

Choice of Discretization Points As in Section 2.1, let $B = (B_1, \dots, B_n)$ be latent binary random variables such that $B_i = 0$ if $\theta_i = 0$ and $B_i = 1$ otherwise. We will choose discretization points $\alpha_1, \dots, \alpha_k$ and the discretized prior $\tilde{\Lambda}_n$ such that the ratio $\Pi_n(B = b) / \tilde{\Pi}_n(B = b)$ is in $[1 - \epsilon, 1 + \epsilon]$ for all realizations b of B , where $\epsilon > 0$ can be made arbitrarily small. Since, conditional on B , the discretized model is the same as the original model, this implies that the posterior probabilities $\Pi_n(\theta | Y)$ and $\tilde{\Pi}_n(\theta | Y)$ must also be within a factor of $(1 + \epsilon) / (1 - \epsilon) \approx 1$.

Conditional on the mixing hyper-parameter α , the sequence B consists of independent, identically distributed Bernoulli random variables, and Π_n and $\tilde{\Pi}_n$ respectively assign hyperpriors Λ_n and $\tilde{\Lambda}_n$ to the success probability α . To discretize α , we will follow an approach introduced by [20] in the context of online sequential prediction with adversarial data. They observe that it is more convenient to reparametrize the Bernoulli model using the *arcsine transformation* [4, 25], which makes the Fisher information constant:

$$\beta(\alpha) = \arcsin \sqrt{\alpha}, \quad \alpha(\beta) = \sin^2 \beta \quad \text{for } \beta \in [0, \pi/2].$$

We will use a uniform discretization of β with k discretization points spaced $\delta_k = \pi / (2k) \propto 1 / \sqrt{n}$ apart, which in the α -parametrization maps to a spacing that is proportional to $1 / \sqrt{n}$ around $\alpha = 1/2$ but behaves like $1/n$ for α near 0 or 1. Specifically, let $\alpha_j = \alpha(\beta_j)$ with

$$\beta_1 = \frac{1}{2}\delta_k, \quad \beta_2 = \frac{3}{2}\delta_k, \quad \beta_3 = \frac{5}{2}\delta_k, \quad \dots, \quad \beta_k = \frac{\pi}{2} - \frac{1}{2}\delta_k.$$

The prior mass of each α under Λ_n is then reassigned to its closest discretization point in the β -parametrization. If Λ_n has no point-masses exactly half-way between discretization points, then this means that

$$\tilde{\Lambda}_n(\alpha_j) = \Lambda_n([\alpha(\beta_j - \delta_k/2), \alpha(\beta_j + \delta_k/2)]). \tag{4}$$

Otherwise, if Λ_n does have such point-masses, their masses may be divided arbitrarily over their neighboring discretization points.

Approximation Guarantees For simplicity we will assume that Λ_n has a Lebesgue-density $\lambda_n(\alpha) = d\Lambda_n(\alpha)/d\alpha$. It will also be convenient to let $\alpha_0 = 0$ and $\alpha_{k+1} = 1$, and to define

$$P_\alpha(n, \hat{\alpha}) = \alpha^{\hat{\alpha}n}(1 - \alpha)^{(1-\hat{\alpha})n}, \quad \text{for } \hat{\alpha} \in [0, 1], n \in \mathbb{R}^+,$$

which may be interpreted as the Bernoulli(α) likelihood of a binary sequence with maximum likelihood parameter $\hat{\alpha}$. In particular, if $\hat{\alpha} = s/n$ with s the number of ones in $b \in \{0, 1\}^n$ for integer n , then

$$\Pi_n(B = b) = \int_0^1 P_\alpha(n, \hat{\alpha})\lambda_n(\alpha)d\alpha, \quad \tilde{\Pi}_n(B = b) = \sum_{j=1}^k P_{\alpha_j}(n, \hat{\alpha})\tilde{\Lambda}_n(\alpha_j).$$

There is no reason to restrict the definition of $P_\alpha(n, \hat{\alpha})$ to integer n or to the discrete set of $\hat{\alpha}$ that can be maximum likelihood parameters at sample size n , however, and following [20] we extend the definition to all $\hat{\alpha} \in [0, 1]$ and all real $n > 0$, which will be useful below to handle the Beta($1, n + 1$) prior.

Theorem 2.1. *Take $k = 2(m + 1)\lceil\sqrt{n}\rceil + 1$ for any integer m , and suppose there exists a constant $L \geq 0$ (which is allowed to depend on n) such that*

$$\frac{\sup_{\alpha \in [\alpha_j, \alpha_{j+1}]} \lambda_n(\alpha)\sqrt{\alpha(1 - \alpha)}}{\inf_{\alpha \in [\alpha_j, \alpha_{j+1}]} \lambda_n(\alpha)\sqrt{\alpha(1 - \alpha)}} \leq e^{L\sqrt{n}\delta_k}, \quad \text{for all } j = 0, \dots, k. \tag{5}$$

Then there exists a constant $C_L > 0$ that depends only on L , such that, if $m > C_L$, we have, for $\epsilon = C_L/m$,

$$(1 - \epsilon) \leq \frac{\int_0^1 P_\alpha(n, \hat{\alpha})\lambda_n(\alpha)d\alpha}{\sum_{j=1}^k P_{\alpha_j}(n, \hat{\alpha})\tilde{\Lambda}_n(\alpha_j)} \leq (1 + \epsilon), \quad \text{for all } \hat{\alpha} \in [0, 1], \tag{6}$$

and consequently

$$\frac{1 - \epsilon}{1 + \epsilon} \leq \frac{\Pi_n(\theta | Y)}{\tilde{\Pi}_n(\theta | Y)} \leq \frac{1 + \epsilon}{1 - \epsilon} \quad \text{almost surely.} \tag{7}$$

The result (6) holds even for non-integer n , but (7) implicitly assumes that n is the number of observations in Y and must therefore be integer. The proof is deferred to Appendix C.1 in the supplementary material. We note already that condition (5) is essentially a Lipschitz condition on the log of the density of Λ_n in the β -parametrization (see (5) in Appendix C.1). Under this condition, the theorem shows that, by increasing m , we can approximate $\Pi_n(\theta | Y)$ to any desired accuracy, at the cost of increasing our computation time, which scales linearly with m .

Remark 1. For given m and (integer) n , the tightest possible value of ϵ in (7) may be determined numerically, by maximizing and minimizing the ratio in (6) over $\hat{\alpha} = s/n$ for $s = 0, \dots, n$.

Extension to Arbitrary Beta Priors The Lipschitz condition (5) excludes the important $\text{Beta}(1, n + 1)$ prior, because its density varies too rapidly. We therefore describe an extension that can handle any $\text{Beta}(\kappa, \lambda)$ prior with $\kappa, \lambda \geq 1/2$, even when κ or λ grows linearly with n .

To this end, we interpret $\text{Beta}(\kappa, \lambda)$ as the posterior of a $\text{Beta}(1/2, 1/2)$ prior after observing $\kappa - 1/2$ fake ones and $\lambda - 1/2$ fake zeros. Our effective sample size for the fake observations and the real data together is then $n' = n + \kappa + \lambda - 1$ (which need not be an integer). Since $\text{Beta}(1/2, 1/2)$ is uniform in the β -parametrization, it satisfies (5) with the best possible constant: $L = 0$, so applying Theorem 2.1 we find that (6) holds for sample size n' with ϵ as in the theorem. We then take the discretization points for sample size n' with corresponding discrete prior $\tilde{\Lambda}_{n'}$ defined by (4) (which is actually uniform, with probabilities $1/k$, because $\text{Beta}(1/2, 1/2)$ is uniform in the β -parametrization), and we compute a new prior $\tilde{\Lambda}_n$ on these discretization points as the posterior from $\tilde{\Lambda}_{n'}$ after observing $\kappa - 1/2$ fake ones and $\lambda - 1/2$ fake zeros:

$$\tilde{\Lambda}_n(\alpha_j) = \frac{\frac{1}{k} \alpha_j^{\kappa-1/2} (1 - \alpha_j)^{\lambda-1/2}}{\sum_{j'=1}^k \frac{1}{k} \alpha_{j'}^{\kappa-1/2} (1 - \alpha_{j'})^{\lambda-1/2}} \quad \text{for } j = 1, \dots, k. \tag{8}$$

Corollary 2.2. *For any $\kappa \geq 1/2, \lambda \geq 1/2$ and positive integer n , let $k = 2(m + 1)\lceil\sqrt{n'}\rceil + 1$, where $n' = n + \kappa + \lambda - 1$ and $m > C_0$ is any integer that exceeds the constant C_L from Theorem 2.1 for $L = 0$. Let Λ_n be the $\text{Beta}(\kappa, \lambda)$ prior and let $\tilde{\Lambda}_n$ be as in (8). Then (6) and (7) hold with $\epsilon' = 2\epsilon/(1 - \epsilon)$ instead of $\epsilon = C_0/m$.*

Proof. Since the joint distributions on n' observations satisfy (6), the corresponding posteriors after conditioning these distributions on $\kappa - 1/2$ fake ones and $\lambda - 1/2$ fake zeros must be within factors $\frac{1-\epsilon}{1+\epsilon} = 1 - \frac{2\epsilon}{1+\epsilon} \geq 1 - \epsilon'$ and $\frac{1+\epsilon}{1-\epsilon} = 1 + \epsilon'$. \square

2.3 Which Model Selection Priors Are Spike-and-Slab Priors?

As described in the introduction, it is clear that spike-and-slab priors are a special case of model selection priors. However, to the best of our knowledge, it is not known when a model selection prior has a spike-and-slab representation. One advantage of the spike-and-slab formulation is that we can construct algorithms with $O(n^{3/2})$ run time (Section 2.2), while for a general model selection prior the computational complexity is $O(n^2)$ (Section 2.1). In this section we give sufficient and necessary conditions for when a model selection prior can be expressed in spike-and-slab form.

To characterize the exact relationship between the priors we introduce the following notation. For $\mu = (\mu_0, \mu_1, \dots, \mu_m)$ with $m \geq 2n$, define the $(n + 1) \times (n + 1)$ Hankel matrix $H_n(\mu) = [\mu_{i+j}]_{i,j=0,\dots,n}$ and let $F\mu = (\mu_1, \dots, \mu_m)$ denote the projection that drops the first coordinate. Furthermore, for $A \in \mathbb{R}^{n \times m}$, let $\text{range}(A)$ be the column space of A and let $A \succeq 0$ denote that A is positive semi-definite.

Theorem 2.3. *For odd $n = 2k + 1$, the model selection prior (with factorizing G_S) can be given in the form (2) if and only if there exists a $c_n \in [0, \pi_n(n)]$ such that*

$$H_k(\mu) \succeq 0, \quad H_k(F\mu) \succeq 0, \quad \text{and} \quad (\mu_{k+1}, \mu_{k+2}, \dots, \mu_{2k+1})^\top \in \text{range}(H_k(\mu)),$$

with $\mu = \left(\binom{n}{0}^{-1} \pi_n(0), \dots, \binom{n}{n-1}^{-1} \pi_n(n-1), c_n \right) \in [0, 1]^{n+1}$.

For even $n = 2k$, the model selection prior (with factorizing G_S) can be given in the form (2) if and only if there exists a $c_n \in [0, \pi_n(n)]$ such that

$$H_k(\mu) \succeq 0, \quad H_{k-1}(F\mu) \succeq 0, \quad \text{and} \quad (\mu_{k+1}, \mu_{k+2}, \dots, \mu_{2k})^\top \in \text{range}(H_{k-1}(F\mu)),$$

with the same μ as above.

The proof, which is given in Appendix C.2, shows that establishing this theorem amounts to proving a version of de Finetti's theorem for finite sequences.

Next we give several examples of priors π_n that satisfy (or fail) the conditions of Theorem 2.3, which implies that the model selection prior can (or cannot) be given in spike-and-slab form (2). The proofs for the examples are in Appendix C.3.

First we consider binomial π_n , for which it is already known that there exists a spike-and-slab representation [17, Example 2.1]. Nevertheless, to illustrate the applicability of our results, we show that this choice of π_n satisfies the conditions of Theorem 2.3.

Example 1. The binomial prior $\pi_n(s) \propto \binom{n}{s} p^s (1-p)^{n-s}$, $p \in [0, 1]$, satisfies the conditions of Theorem 2.3 and therefore the corresponding model selection prior can be given in the spike-and-slab form (2) for some appropriate probability measure Λ_n on $[0, 1]$.

The next example treats the Poisson prior as a choice for π_n . To the best of our knowledge there are no results in the literature that establish whether the corresponding model selection prior can be given in the spike-and-slab form (2).

Example 2. For any $\lambda > 0$, the Poisson prior $\pi_n(s) \propto e^{-\lambda} \lambda^s / s!$ restricted to $s \in \{0, 1, \dots, n\}$ satisfies the conditions of Theorem 2.3 and therefore the corresponding model selection prior can be given in the form (2) for some appropriate probability measure Λ_n on $[0, 1]$.

We proceed to give two natural choices for π_n where the corresponding model selection prior cannot be expressed in the form (2). In the first example, π_n has a heavy (polynomial) tail, while in the second it has a light (sub-exponential) tail.

Example 3. Let us consider the prior $\pi_n(0) \propto 1$, $\pi_n(s) \propto s^{-\lambda}$, $s = 1, \dots, n$, for any $\lambda > 1$. For $n > 2^{\lambda-1} / (2^{\lambda-1} - 1)$ this prior does not satisfy the conditions of Theorem 2.3 and therefore the corresponding model selection prior cannot be represented in the form (2).

Example 4. We consider the sub-exponential prior $\pi_n(s) \propto e^{-s^\lambda}$, $s = 0, 1, \dots, n$ for any $\lambda > \log_2(2 + \ln 2)$. For $n > c / (c - 1)$ with $c = e^{2^\lambda - 2} / 2 > 1$ this prior does not satisfy the conditions of Theorem 2.3 and therefore the corresponding model selection prior cannot be represented in the form (2).

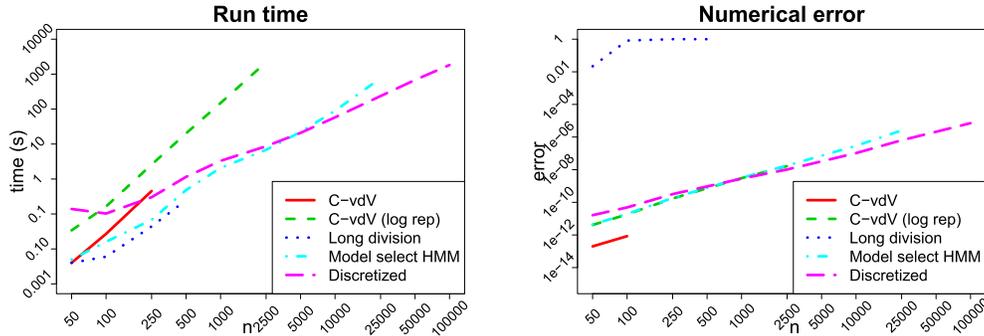


Figure 2: Run time and numerical accuracy for the exact algorithms from Section 2 in calculating $q_{n,i}$ for $i = 1, \dots, n$. A numerical error of 10^{-a} means that the algorithm is able to calculate the mathematically exact answer up to a decimal places.

3 Simulation Study: Reliability of Algorithms

3.1 Comparing the Proposed Algorithms

In this section we investigate the speed and numerical accuracy of our new algorithms to the previously proposed methods for exact computation of the posterior. We consider a sequence of sample sizes $n = 50, 100, 250, 500, 1\,000, 2\,500, \dots, 50\,000, 100\,000$ and construct the true signal θ_0 to have 20% non-zero signal components of value $4\sqrt{2\ln n}$, while the rest of the signal coefficients are set to be zero. For fair comparison we run all algorithms for the spike-and-slab prior with Laplace slab $g(x) = \frac{a}{2}e^{-a|x|}$, with $a = 1$, and mixing hyper-prior $\Lambda_n = \text{Beta}(1, n + 1)$. We have set up the experiments in R, but all algorithms were implemented as subroutines in C++. Since numerical instability is a major concern, we have tracked the numerical accuracy of all methods using interval arithmetic as implemented in the C++ Boost library [8] (with cr-libm as a back-end to compute transcendental functions [21]), which replaces all floating point numbers by intervals that are guaranteed to contain the mathematically exact answer. The lower end-point of each interval corresponds to always rounding down in the calculations, and the upper end-point corresponds to always rounding up. The width of the interval for the final answer therefore measures the numerical error. All experiments were performed on a MacBook Pro laptop with 2.9 GHz Intel Core i5 processor, 8 GB (1867 MHz DDR3) memory, and a solid-state hard drive.

Results The results are summarized in Figure 2, which shows the run time of the algorithms on the left, and their numerical error on the right. The reported numerical error is the maximum numerical error in calculating $q_{n,i}$ over $i = 1, \dots, n$. To avoid overly long computations we have terminated the algorithms if they became numerically unstable or if their run time exceeded half an hour. One can see that the original Castillo-Van der Vaart algorithm was terminated for $n \geq 250$, which was due to numerical inaccuracy. This problem was resolved by applying the logarithmic representation

from Appendix A.2 which made the algorithm numerically stable up to $n \leq 2500$; however, due to the long $O(n^3)$ run time the algorithm was terminated for larger values as it reached the half-hour limit. The natural speed-up idea of applying long division (see Appendix A.3) was not successful for this data as even for small sample sizes the numerical accuracy was poor. We observe that the model selection HMM and the algorithm based on discretization performed superior to the preceding methods: the model selection HMM algorithm has run time $O(n^2)$ and the largest sample size it managed to complete within half an hour was $n = 25\,000$, while the algorithm with discretized mixing parameter in the spike-and-slab prior (initialized according to Corollary 2.2 with parameter $m = 20$) has run time $O(n^{3/2})$ and reached the time limit after sample size $n = 100\,000$. We also note that both algorithms were numerically accurate, giving answers that were reliable up to between 5 and 11 decimal places, depending on n . For sample size $n = 2\,500$, we have further verified empirically that indeed the model selection HMM algorithm computes the same numbers as the Castillo-Van der Vaart algorithm, as was already shown in Section 2.1.

3.2 Approximation Errors for Several Standard Methods

In this section we measure the approximation error of a selection of approximation algorithms by comparing them to the exact model selection HMM algorithm, which serves as a benchmark for the correct answer. We again consider the spike-and-slab prior with $\Lambda_n = \text{Beta}(1, n + 1)$, but for simplicity we use standard Gaussian slabs $g(x) = \frac{1}{\sqrt{2\pi}}e^{-x^2}$, since the approximation methods are typically designed for this choice of slab distribution. Our first approximation method is the discretization algorithm from Section 2.2, which uses a deterministic approximation. The discretization algorithm was again initialized according to Corollary 2.2 with $m = 20$. We further consider a standard Gibbs sampler (with number of iterations $it = 10^3, 10^4, 10^5$, half of which are used as burn-in) and a variational Bayes approximation. We consider the same test data as in the preceding section. The only difference is that we stop at $n = 10\,000$ to limit the run times for the exact HMM algorithm and the Gibbs sampler with $it=10^5$. Both the Gibbs sampler and variational Bayes algorithm were implemented in R. For the latter we used the component-wise variational Bayes algorithm [38, 57, 10, 47]. We measure approximation error by computing $\max_i |q_{n,i} - \tilde{q}_{n,i}|$, where $q_{n,i}$ is the exact slab probability computed by the model selection HMM and $\tilde{q}_{n,i}$ is the slab probability computed by the approximation. We run each non-deterministic approximation method 5 times and report the average approximation error along with the average run time of the algorithms. The results are plotted in Figure 3 and shown numerically in Table 1.

One can see that the discretized version of the algorithm is very accurate, with at least seven decimal places of precision throughout. It approximately loses two decimal places of precision for every ten-fold increase of n , so we can still expect it to be accurate up to five decimal places for $n = 100\,000$. We point out that its approximation error includes both the mathematical approximation from Section 2.2 and the numerical error already studied separately in Figure 2. Since the approximation error in Figure 3 is of the same order as the numerical error in Figure 2, we conclude that the numerical error

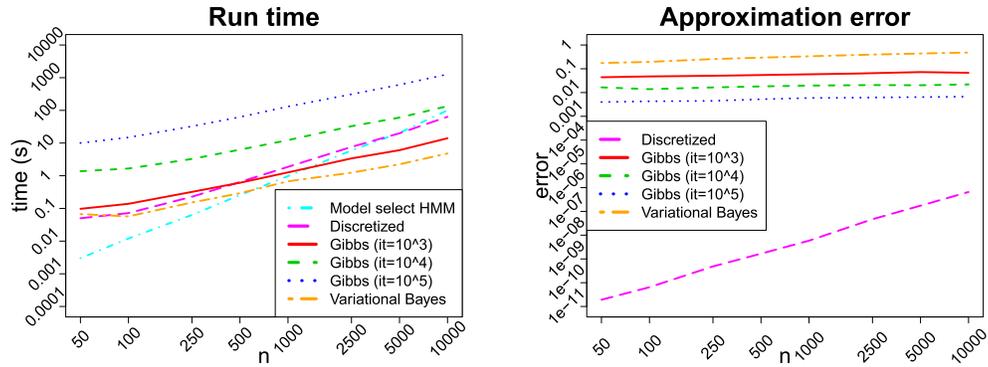


Figure 3: Run times and approximation errors for approximate algorithms. An approximation error of 10^{-a} means that the algorithm is able to calculate the correct answer up to a decimal places.

Method \ n	100	250	500	1 000	2 500	5 000	10 000
Discretized	6.37×10^{-11}	4.89×10^{-10}	1.67×10^{-9}	5.89×10^{-9}	4.69×10^{-8}	1.74×10^{-7}	6.56×10^{-7}
Gibbs (it = 10^3)	4.58×10^{-2}	4.76×10^{-2}	5.03×10^{-2}	5.41×10^{-2}	6.15×10^{-2}	6.28×10^{-2}	6.96×10^{-2}
Gibbs (it = 10^4)	1.23×10^{-2}	1.46×10^{-2}	1.63×10^{-2}	1.75×10^{-2}	2.03×10^{-2}	2.07×10^{-2}	2.25×10^{-2}
Gibbs (it = 10^5)	4.55×10^{-3}	5.05×10^{-3}	5.46×10^{-3}	5.63×10^{-3}	5.71×10^{-3}	6.86×10^{-3}	6.88×10^{-3}
Variational Bayes	1.90×10^{-1}	2.48×10^{-1}	2.93×10^{-1}	3.33×10^{-1}	3.91×10^{-1}	4.40×10^{-1}	4.81×10^{-1}

Table 1: Approximation errors compared to the exact HMM algorithm.

dominates the mathematical approximation error, so the discretization algorithm may be considered an exact method for all practical purposes.

At the same time the Gibbs sampler and the variational Bayes method both provide approximations of the posterior that are far less accurate. Variational Bayes is only accurate up to one decimal place, although in further investigations we did find that it provides a better approximation if we look only at the non-zero coefficients, with an approximation error of order $O(10^{-4})$. For the Gibbs sampler there is no theory that tells us how many iterations we have to take to achieve a certain degree of accuracy. We see here that the precision strongly depends on the number of iterations, ranging from one to three decimal places, but remains approximately constant with increasing n . However, the run time for $it = 10^5$ iterations would become prohibitive for sample sizes much larger than the $n = 10\,000$ we consider.

4 Differential Gene Expression for Ulcerative Colitis and Crohn’s Disease

In this section we compare our methods to various other frequently used Bayesian approaches in the context of differential gene expression data.

Method	Run Time	Nr. of Genes Selected
Variational Bayes (varbvs)	20.95 minutes	166
Spike-and-Slab LASSO	0.01 seconds	557
Horseshoe (10 runs)	1.86 minutes	571–583
Empirical Bayes EBSpars (10 runs)	7.28 seconds	592–604
Discretized: Beta(1, $n + 1$)-binomial prior	24.47 seconds	674
HMM: Beta(1, $n + 1$)-binomial prior	2.06 minutes	674
Empirical Bayes JS	0.03 seconds	3168
HMM: Beta(1, 1)-binomial prior	2.00 minutes	3169

Table 2: Run time and number of selected genes on gene expression data. The reported run times for Empirical Bayes EBSpars and the Horseshoe are the averages over their runs.

Data We consider a data set from Burczynski et al. [9] containing the gene expression levels of $n = 22\,283$ genes in peripheral blood mononuclear cells, with the raw data provided by the National Center for Biotechnology Information.¹ This is an observational study, with microarray gene expression data on 26 subjects who suffered from ulcerative colitis and 59 subjects with Crohn’s disease. We calculate Z-scores to identify differences in average gene expression levels between the two disease groups following the standard approach described by Quackenbush [45], which consists of dividing the difference of the average log-transformed, normalized gene expressions for the two groups by the standard error. More specifically, let us denote by $U_{i,j}$ and $C_{i,j}$ the measured intensities of the i -th gene and j -th person with ulcerative colitis and Crohn’s disease, respectively. As a first step we normalize the intensities for each patient, i.e. we take $U'_{i,j} = U_{i,j} / \sum_i U_{i,j}$ and $C'_{i,j} = C_{i,j} / \sum_i C_{i,j}$ for each gene i and patient j . Then the Z-score for the i -th gene is computed as

$$Z_i = \frac{\overline{\log U'_i} - \overline{\log C'_i}}{\sqrt{\sigma_{U',i}^2/26 + \sigma_{C',i}^2/59}} \quad \text{for } i = 1, \dots, 22\,283,$$

where

$$\overline{\log U'_i} = \frac{1}{26} \sum_{j=1}^{26} \log U'_{i,j}, \quad \sigma_{U',i}^2 = \frac{1}{25} \sum_{j=1}^{26} (\log U'_{i,j} - \overline{\log U'_i})^2,$$

and $\overline{\log C'_i}$ and $\sigma_{C',i}^2$ are defined accordingly. Since it is assumed that the number of genes with a different expression level between the two groups is small compared to the total number of genes n , the data fit into the sparse normal sequence model with $n = 22\,283$.

Methods We compare the run times and the selected genes for the eight procedures listed in Table 2. We consider the model selection HMM algorithm for the Beta(1, $n + 1$)-binomial prior with Laplace slab (with hyper-parameter $a = 0.5$), and the discretization algorithm from Corollary 2.2 with $m = 20$, which is a faster way to compute exactly the same results. Genes i with marginal posterior probability $q_{n,i} \geq 1/2$ are selected.

¹Via the Gene Expression Omnibus (GEO) website under dataset record number GDS1615. See www.ncbi.nlm.nih.gov/sites/GDSbrowser?acc=GDS1615.

For comparison, we also consider the model selection HMM for the Beta(1, 1)-binomial prior, which corresponds to using a uniform prior Λ_n on the mixing parameter α . In this section, we used the implementations of our algorithms from our R package [62], which is approximately 5 times faster than the implementation from Section 3, because it does not incur the overhead of tracking numerical accuracy using interval arithmetic.

We compare to the empirical Bayes method of Johnstone and Silverman [31], which uses a spike-and-slab prior, but estimates the mixing parameter α using empirical Bayes. The method does not explicitly include a prior on α , but we may interpret it as using a uniform prior Λ_n . We again use a Laplace slab (with the default parameter $a = 0.5$) and select genes by hard thresholding at marginal posterior probability 1/2, as implemented in the R package [55].

We also include EBSparse, which is a fractional empirical Bayes procedure proposed by Martin and Walker [40]. It can be interpreted as using a spike-and-slab prior with $\Lambda_n = \text{Beta}(1, \gamma n)$, but with Gaussian slabs $G_i = \mathcal{N}(Y_i, \tau^2)$ whose means depend on the data. Furthermore, in the formula for the posterior the likelihood is tempered by raising it to the power κ . We use the authors' R implementation [41], with the recommended hyper-parameter settings $\kappa = 0.99$, $\gamma = 0.25$, $\tau^2 = 100$, and $M = 1000$ Monte Carlo samples. As the sampler is randomized, we run the algorithm 10 times.

We further consider the Spike-and-Slab LASSO of Ročková [48], which computes the maximum a posteriori parameters using Laplace distributions both for the spikes and for the slabs. As in [48, Section 6], we take the slab scale parameter to be $\lambda_1 = 0.1$, and estimate the spike scale parameter λ_0 via the two-step procedure described there, for the Beta(1, $n + 1$) hyper-prior on the mixing parameter. An R implementation called SSLasso was provided by Ročková [49].

We also add the Horseshoe estimator [11] with the Cauchy hyper-prior on its hyper-parameter τ , truncated to the interval $[1/n, 1]$, as recommended by Van der Pas et al. [60]. We use the R package [58], with its default Markov Chain Monte Carlo sampler settings of 1000 iterations burn-in and 5000 iterations after burn-in. Genes are selected if their credible sets exclude zero [60]. As the sampler is randomized, we run the algorithm 10 times.

Finally, we compare with the variational Bayes algorithm (varbvs R-package) described in [10]. Notably, this method uses Gaussian slabs. The hyper-parameters (e.g. the variance of the prior and the noise) are automatically fitted to the data. We set the tolerance and maximum number of iterations to be 10^{-4} and 1000, respectively.

Results Results are reported in Table 2 and Figure 4. Although we list run times to illustrate computational feasibility, it is important to keep in mind that the methods in this section compute different quantities, so their most important difference lies in which genes they select. On this point, the main conclusion is that the alternative methods give very different results from using the exact Bayesian posterior for the model selection prior.

All methods except the Horseshoe and EBSparse select genes in decreasing order of the absolute values of Z_i . Genes are generally selected by the Horseshoe and EBSparse in

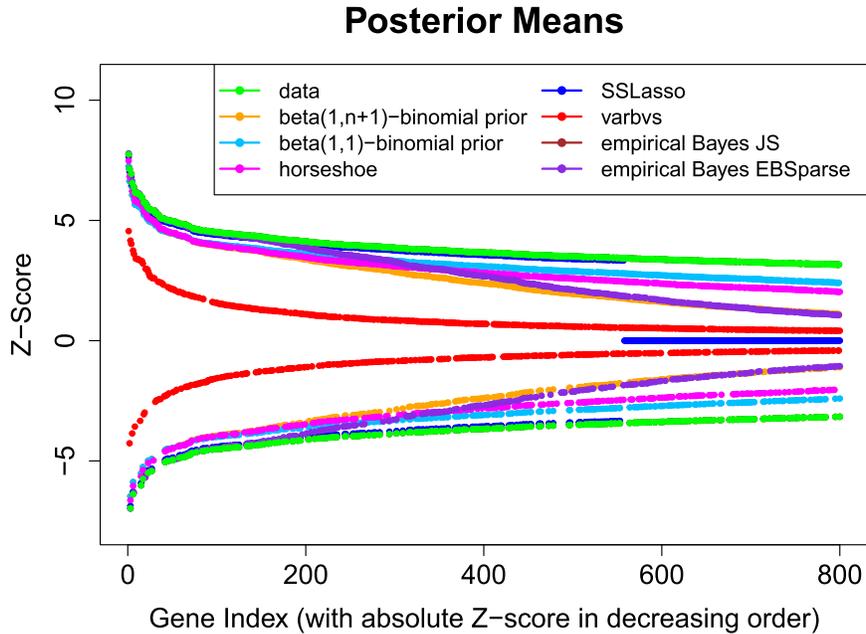


Figure 4: Posterior means/MAP estimates for the 800 genes in the gene expression data with largest Z-scores (in absolute value).

decreasing order of absolute value of Z_i as well, but with some swaps for genes for which the absolute values are close to each other, so it appears that for all sampling-based methods the sampler is suffering from limited precision, as we also observed for the Gibbs samplers in Section 3.2. The methods can be divided into three main categories based on the number of genes they select: on one extreme is the variational Bayes (varbvs) method, which provides the sparsest solution; then the majority of methods select a number of genes between 557 and 674; and finally at the other extreme are the Empirical Bayes JS procedure and the Beta(1,1)-binomial prior, which are both based on the same prior and both select a very large number of genes, making these two methods the most conservative. The lack of sparsity induced by the Beta(1,1)-binomial prior is perhaps not surprising, given that it does not satisfy the exponential decrease condition of [17]. We study this further in Section 5, where we compare different choices for the hyper-parameters of the beta-binomial prior in simulations.

We further see that the Spike-and-Slab LASSO and the empirical Bayes JS procedures finish almost instantly. The EBSpars method takes several seconds to run, as does the discretization algorithm. The Horseshoe and our exact model selection HMM take approximately two minutes to run, while the variational Bayes varbvs method requires a little over 20 minutes. Nevertheless, all methods are feasible even for practitioners who would like to perform multiple similar experiments, for example with different variations of the prior or slab distributions. By contrast, we do not include the Castillo-Van

der Vaart algorithm with logarithmic representation, because based on extrapolation of Figure 2 we expect it to take around 20 days.

In Figure 4 we also plot the posterior means (or, in case of the Spike-and-Slab LASSO, the MAP estimator) and the 800 largest Z-scores in absolute value. Since the posterior means for the model selection HMM and the discretization algorithm are the same, we label both as Beta($1, n + 1$)-binomial in reference to the prior that was used. We further note that the empirical Bayes JS estimates are invisible behind the data points. We observe that the varbvs method induces the heaviest shrinkage, followed first by the Beta($1, n + 1$)-binomial prior and the Empirical Bayes EBSparse method, and then by the Horseshoe and the Beta($1, 1$)-binomial prior. The least shrinkage is applied by the empirical Bayes JS method, which does not shrink the observed Z-scores very much (if at all). The Spike-and-Slab LASSO is in a category of its own, because it is a MAP estimator. It applies no shrinkage to the coefficients that are selected, and sets all other coefficients to zero.

5 Asymptotics of Spike-and-Slab Priors

The choice of the prior Λ_n on the mixing hyper-parameter α in spike-and-slab priors is considered to be highly relevant for the behavior of the posterior. Castillo and Van der Vaart [17] recommend to use $\Lambda_n = \text{Beta}(\kappa, \lambda)$ with parameters $\kappa = 1$ and $\lambda = n + 1$. This prior induces heavy penalization for dense models (models with large sparsity parameter s) and was shown to have optimal theoretical properties. However, it is unknown whether such heavy penalization is indeed necessary and whether even heavier penalization will result in suboptimal behavior.

In this section we investigate the asymptotic behavior of the posterior for different choices of the hyper-parameters κ and λ using our new exact algorithms, which can scale up to large sample sizes. We consider: i) the uniform prior with $\kappa = 1$ and $\lambda = 1$, which is often considered a natural choice [54]; ii) mild shrinkage, $\kappa = 1$ and $\lambda = \sqrt{n}$; iii) the choice $\kappa = 1$ and $\lambda = n + 1$ recommended by Castillo and Van der Vaart; iv) heavy shrinkage, $\kappa = 1$ and $\lambda = n^2$; and finally v) a sparsity-discouraging choice, $\kappa = n$ and $\lambda = 1$. We consider two experiments: A_1 and A_2 . In both cases the sample sizes range from $n = 50$ to $n = 20\,000$. In Experiment A_1 we set the true sparsity level to $s = 10$ and consider uniformly distributed non-zero signal coefficients between 1 and 10, i.e. $\theta_i \sim U(1, 10)$ for $i \in \mathcal{S}$. In Experiment A_2 the true sparsity level is taken to be $s = \lceil n^{1/3} \rceil$ and the non-zero signal coefficients are set to $\theta_i = 2\sqrt{2 \log n}$ for $i \in \mathcal{S}$, which is a factor of 2 above the detection threshold. In Appendix B of the supplementary material we consider an additional experiment A_3 that is similar to A_1 but with $s = 25$ and $\theta_i \sim U(5, 10)$ for $i \in \mathcal{S}$, which gives similar results as Experiment A_1 .

We repeat each experiment 20 times and report the average ℓ_2 -error between the posterior mean for θ and the true signal θ in Table 3 and Table 6 for Experiments A_1 and A_2 , respectively. In Tables 4, 5, 7 and 8 we also report the average false discovery rates and the average true positive rates. Standard deviations are provided in parentheses in all cases.

Method \ n	50	100	200	500	1 000	2 000	5 000	10 000	20 000
i) $\kappa = 1, \lambda = 1$	4.52 (0.64)	4.65 (0.68)	4.84 (0.75)	5.30 (1.01)	5.18 (0.74)	5.63 (0.90)	5.95 (0.77)	6.65 (1.26)	6.38 (0.98)
ii) $\kappa = 1, \lambda = \sqrt{n}$	4.28 (0.63)	4.46 (0.69)	4.69 (0.75)	5.20 (1.01)	5.13 (0.73)	5.57 (0.88)	5.92 (0.78)	6.62 (1.27)	6.36 (0.98)
iii) $\kappa=1, \lambda = n+1$	4.09 (0.68)	4.20 (0.82)	4.43 (0.88)	4.95 (1.11)	5.14 (0.89)	5.45 (0.99)	6.18 (0.94)	6.42 (1.67)	6.43 (1.06)
iv) $\kappa = 1, \lambda = n^2$	5.04 (1.05)	5.21 (1.42)	6.55 (1.48)	7.09 (1.93)	7.87 (1.21)	8.15 (1.42)	9.45 (1.55)	10.35 (2.24)	10.43 (2.16)
v) $\kappa = n, \lambda = 1$	5.71 (0.59)	7.60 (0.62)	10.63 (0.54)	16.89 (0.61)	23.39 (0.61)	33.43 (0.78)	52.93 (0.52)	74.79 (0.57)	105.84 (0.55)

Table 3: ℓ_2 distance of the posterior mean from the true signal in Experiment A_1 .

Method \ n	50	100	200	500	1 000	2 000	5 000	10 000	20 000
i) $\kappa = 1, \lambda = 1$	0.36 (0.14)	0.23 (0.13)	0.19 (0.11)	0.13 (0.12)	0.05 (0.09)	0.09 (0.12)	0.08 (0.09)	0.12 (0.10)	0.08 (0.09)
ii) $\kappa = 1, \lambda = \sqrt{n}$	0.16 (0.10)	0.16 (0.11)	0.17 (0.11)	0.11 (0.11)	0.06 (0.09)	0.09 (0.12)	0.08 (0.08)	0.11 (0.10)	0.08 (0.09)
iii) $\kappa = 1, \lambda = n + 1$	0.05 (0.06)	0.03 (0.05)	0.02 (0.04)	0.02 (0.04)	0.01 (0.03)	0.01 (0.02)	0.00 (0.00)	0.01 (0.03)	0.00 (0.00)
iv) $\kappa = 1, \lambda = n^2$	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
v) $\kappa = n, \lambda = 1$	0.80 (0.00)	0.90 (0.00)	0.95 (0.00)	0.98 (0.00)	0.99 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)

Table 4: False discovery rate in Experiment A_1 .

Method \ n	50	100	200	500	1 000	2 000	5 000	10 000	20 000
i) $\kappa = 1, \lambda = 1$	0.91 (0.11)	0.83 (0.15)	0.86 (0.08)	0.80 (0.17)	0.76 (0.13)	0.73 (0.17)	0.73 (0.13)	0.71 (0.14)	0.66 (0.16)
ii) $\kappa = 1, \lambda = \sqrt{n}$	0.89 (0.11)	0.83 (0.14)	0.85 (0.09)	0.78 (0.17)	0.75 (0.13)	0.73 (0.17)	0.72 (0.13)	0.71 (0.15)	0.65 (0.16)
iii) $\kappa = 1, \lambda = n + 1$	0.84 (0.12)	0.76 (0.11)	0.80 (0.10)	0.74 (0.18)	0.70 (0.12)	0.68 (0.18)	0.64 (0.12)	0.64 (0.19)	0.59 (0.15)
iv) $\kappa = 1, \lambda = n^2$	0.71 (0.14)	0.62 (0.14)	0.59 (0.16)	0.57 (0.17)	0.52 (0.13)	0.48 (0.20)	0.44 (0.15)	0.40 (0.17)	0.38 (0.16)
v) $\kappa = n, \lambda = 1$	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)

Table 5: True positive rate in Experiment A_1 .

Method \ n	50	100	200	500	1 000	2 000	5 000	10 000	20 000
i) $\kappa = 1, \lambda = 1$	3.03 (0.98)	3.77 (0.70)	3.9 (0.58)	4.48 (0.96)	4.68 (0.51)	5.65 (0.84)	7.13 (0.81)	7.70 (1.21)	8.68 (0.82)
ii) $\kappa = 1, \lambda = \sqrt{n}$	2.79 (0.91)	3.56 (0.68)	3.80 (0.59)	4.32 (0.93)	4.55 (0.51)	5.53 (0.83)	7.02 (0.81)	7.61 (2.15)	8.59 (1.21)
iii) $\kappa = 1, \lambda = n + 1$	2.41 (0.88)	3.01 (0.68)	3.07 (0.74)	3.48 (0.79)	3.64 (0.62)	4.27 (0.63)	5.33 (0.78)	5.68 (0.94)	6.32 (0.77)
iv) $\kappa = 1, \lambda = n^2$	3.43 (1.82)	3.63 (2.06)	3.89 (1.86)	5.13 (2.51)	4.49 (1.90)	5.08 (2.27)	5.62 (1.42)	6.27 (2.15)	6.42 (1.70)
v) $\kappa = n, \lambda = 1$	5.31 (0.84)	7.69 (0.48)	10.67 (0.52)	16.81 (0.71)	23.56 (0.56)	33.59 (0.50)	53.00 (0.44)	74.92 (0.53)	105.8 (0.54)

Table 6: ℓ_2 distance of the posterior mean from the true signal in Experiment A_2 .

Method \ n	50	100	200	500	1 000	2 000	5 000	10 000	20 000
i) $\kappa = 1, \lambda = 1$	0.25 (0.22)	0.22 (0.17)	0.15 (0.12)	0.12 (0.09)	0.10 (0.08)	0.11 (0.08)	0.11 (0.06)	0.12 (0.06)	0.10 (0.03)
ii) $\kappa = 1, \lambda = \sqrt{n}$	0.15 (0.16)	0.18 (0.14)	0.14 (0.11)	0.11 (0.09)	0.08 (0.08)	0.10 (0.08)	0.10 (0.06)	0.11 (0.06)	0.10 (0.03)
iii) $\kappa = 1, \lambda = n + 1$	0.01 (0.04)	0.06 (0.08)	0.03 (0.07)	0.04 (0.06)	0.01 (0.03)	0.02 (0.05)	0.02 (0.03)	0.02 (0.02)	0.01 (0.02)
iv) $\kappa = 1, \lambda = n^2$	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
v) $\kappa = n, \lambda = 1$	0.92 (0.00)	0.95 (0.00)	0.97 (0.00)	0.98 (0.00)	0.99 (0.00)	0.99 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)

Table 7: False discovery rate in Experiment A_2 .

Method \ n	50	100	200	500	1 000	2 000	5 000	10 000	20 000
i) $\kappa = 1, \lambda = 1$	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)
ii) $\kappa = 1, \lambda = \sqrt{n}$	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)
iii) $\kappa = 1, \lambda = n + 1$	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)
iv) $\kappa = 1, \lambda = n^2$	0.89 (0.19)	0.94 (0.09)	0.94 (0.08)	0.95 (0.09)	0.98 (0.05)	0.99 (0.03)	0.99 (0.02)	0.99 (0.02)	1.00 (0.01)
v) $\kappa = n, \lambda = 1$	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)

Table 8: True positive rate in Experiment A_2 .

In Experiment A_1 we see that the ℓ_2 -error is not very sensitive to the choice of hyperparameters: the uniform prior i), the mild shrinkage ii), and Castillo and Van der Vaart's recommendation iii) all perform comparably. Only the heavy shrinkage iv) is introducing too high penalization, especially for large models. Unsurprisingly, the choice of hyper-parameters v) is also substantially worse than the others, because it expresses exactly the wrong type of prior assumptions by heavily penalizing sparse models. In Experiment A_2 we see that the best hyper-parameters are Castillo and Van der Vaart's recommendation iii) and the heavy shrinkage iv), with the latter having a large variability in performance. Hyper-parameter choices i) and ii) are introducing no or only mild penalization for large models and indeed are also observed to have somewhat worse performance than choices iii) and iv), with the difference getting more pronounced for larger sample sizes. Finally, as in Experiment A_1 , the hyper-parameter setting v) is the worst by far.

We also study the false discovery rate (FDR) and true positive rate (TPR) of the spike-and-slab priors (relatedly, see [14] for the theoretical underpinning of FDR control with empirical Bayes spike-and-slab priors). Unsurprisingly, the FDR is smallest in both experiments in case of heavy shrinkage v), but almost equally good rates are obtained for the recommended choice iii). Mild ii) or no i) shrinkage result in somewhat worse FDR, while the sparsity discouraging setting v) essentially selects all the noise. In Experiment A_1 the best TPR is obtained, not surprisingly, by setting v), which conservatively selects everything. Hyper-parameter choices i) and ii) perform comparably well, closely followed by iii), while the heavy shrinkage method iv) is substantially worse. In Experiment A_2 all hyper-parameter settings perform equally well, except for the heavy shrinkage iv), which is slightly worse. The good performance of the methods is due to the relatively high value ($2\sqrt{2\log n}$) for the non-zero signal coefficients, which lies above the detection threshold $\sqrt{2\log n}$.

We conclude that, overall, the recommended choice iii) indeed appears to have an advantage over the alternatives, and that even heavier penalization as in choice iv) is harmful.

The above simulation study is just one example of how our exact algorithms can be used to study asymptotic properties of model selection priors, and more specifically spike-and-slab priors. Another possible application not considered here would, for instance, be to study the accuracy of Bayesian uncertainty quantification (see [16] for frequentist coverage of Bayesian credible sets resulting from empirical Bayes spike-and-slab priors).

6 Discussion

We have proposed fast and exact algorithms for computing the Bayesian posterior distribution corresponding to model selection priors (including spike-and-slab priors as a special case) in the sparse normal sequence model. Since the normal sequence model corresponds to linear regression with identity design, the question arises whether the derived algorithms can be extended to sparse linear regression with more general designs or other more complex models. We first note that all methods are agnostic about

where the conditional densities of the spikes $p(Y_i | B_i = 0) = \phi(Y_i)$ and the slabs $p(Y_i | B_i = 1) = \psi(Y_i)$ come from. It is therefore trivial to extend them to any model that replaces the distribution of Y_i given \mathcal{S} by

$$Y_i | \mathcal{S} \sim \begin{cases} \psi_i & \text{if } i \in \mathcal{S}, \\ \phi_i & \text{otherwise,} \end{cases}$$

for any densities ψ_i and ϕ_i . (In fact, this is already supported by our R package [62].) Such extensions make it possible to easily handle other noise models for ε_i or general diagonal designs; and, as pointed out by a referee, it also allows incorporating a non-atomic prior on θ_i in case $i \notin \mathcal{S}$. We further anticipate that extensions to general sparse design matrices may be possible by generalizing the HMM from Section 2.1 to more general Bayesian networks and applying a corresponding inference algorithm to compute marginal posterior probabilities. However, for non-sparse design matrices the extension would be very challenging, if possible at all, because the Bayesian network of the hidden states could become fully connected. An interesting intermediate case is studied by Papaspiliopoulos and Rossell [43], who consider best-subset selection for block-diagonal designs. For the normal sequence model, their assumptions amount to the requirement that Λ_n is a point-mass on a single α , and they point out that in this case “best-subset selection becomes trivial.” For non-diagonal designs their results are non-trivial, because they are able to integrate over a continuous hyper-prior on the variance σ^2 of the noise ε_i . In contrast, we assume fixed σ^2 , which we then take to be $\sigma^2 = 1$ without loss of generality. Our methods can be used to calculate the marginal likelihood $p(Y | \sigma^2)$ without further computational overhead, so it would be possible to run them multiple times to incorporate a discrete prior on a grid of values for σ^2 , but it is not obvious if our results can be extended to continuous priors over σ^2 . Exploration of these directions is left for future work.

Even without extending our methods to full linear regression or continuous priors on σ^2 , we believe that they are already very useful as a benchmark procedure: any approximation technique for general linear regression may be applied to the special case of sparse normal sequences and its approximation error computed as in Section 3.2. If a method does not work well in this special case, then certainly we cannot trust it for more general regression. The existence of such a benchmark method is very important, since, for instance, there are no available diagnostics to determine whether Markov Chain Monte Carlo samplers have converged to their stationary distribution or if they have explored a sufficient proportion of the models in the model space.

We have also explored the exact connection between general model selection priors and the more specific spike-and-slab priors. Since for spike-and-slab priors one can construct faster algorithms, it is useful to know which model selection priors can be represented in this form. The proof of our result amounts to a finite sample version of de Finetti’s theorem for a particular subclass of exchangeable distributions, which may be of interest in its own right.

Supplementary Material

Supplement: Fast Exact Bayesian Inference for Sparse Signals in the Normal Sequence Model (DOI: [10.1214/20-BA1227SUPP](https://doi.org/10.1214/20-BA1227SUPP); .pdf). The supplement contains a review of the exact algorithm by Castillo and Van der Vaart and a discussion on how to perform all computations in a logarithmic representation. It further includes an additional variation on Experiment A_1 from Section 5. Finally, the proofs for all theorems and the examples from Section 2.3 are also given in the supplement.

References

- [1] Abramovich, F., Benjamini, Y., Donoho, D. L., and Johnstone, I. M. (2006). “Adapting to unknown sparsity by controlling the false discovery rate.” *Annals of Statistics*, 34(2): 584–653. MR2281879. doi: <https://doi.org/10.1214/009053606000000074>. 933, 934
- [2] Altman, E. I. (1968). “Financial ratios, discriminant analysis and the prediction of corporate bankruptcy.” *The Journal of Finance*, 23(4): 589–609. 933
- [3] Altman, E. I., Iwanicz-Drozdzowska, M., Laitinen, E. K., and Suvas, A. (2014). “Distressed firm and bankruptcy prediction in an international context: A review and empirical analysis of Altman’s Z-score model.” Available at SSRN 2536340. 933
- [4] Anscombe, F. J. (1948). “The transformation of Poisson, binomial and negative-binomial data.” *Biometrika*, 35: 246–254. MR0028556. doi: <https://doi.org/10.1093/biomet/35.3-4.246>. 941
- [5] Bayarri, M. J., Berger, J. O., Forte, A., García-Donato, G., et al. (2012). “Criteria for Bayesian model choice with application to variable selection.” *The Annals of Statistics*, 40(3): 1550–1577. MR3015035. doi: <https://doi.org/10.1214/12-AOS1013>. 934
- [6] Belitser, E. and Nurushev, N. (2020). “Needles and straw in a haystack: Robust confidence for possibly sparse sequences.” *Bernoulli*, 26(1): 191–225. MR4036032. doi: <https://doi.org/10.3150/19-BEJ1122>. 935
- [7] Birgé, L. and Massart, P. (2001). “Gaussian model selection.” *Journal of the European Mathematical Society*, 3(3): 203–268. MR1848946. doi: <https://doi.org/10.1007/s100970100031>. 934
- [8] Boost open source contributors (2018). *BOOST C++ Libraries*. Version 1.67. URL <https://www.boost.org> 945
- [9] Burczynski, M. E., Peterson, R. L., Twine, N. C., Zuberek, K. A., Brodeur, B. J., Casciotti, L., Maganti, V., Reddy, P. S., Strahs, A., Immermann, F., Spinelli, W., Schwertschlag, U., Slager, A. M., Cotreau, M. M., and Dorner, A. J. (2006). “Molecular classification of Crohn’s disease and ulcerative colitis patients using transcriptional profiles in peripheral blood mononuclear cells.” *The Journal of Molecular Diagnostics*, 8(1): 51–61. 948

- [10] Carbonetto, P., Stephens, M., et al. (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>. 946, 949
- [11] Carvalho, C. M., Polson, N. G., and Scott, J. G. (2010). “The horseshoe estimator for sparse signals.” *Biometrika*, 97(2): 465–480. MR2650751. doi: <https://doi.org/10.1093/biomet/asq017>. 935, 949
- [12] Castillo, I. (2017). “Personal communication.” 936
- [13] Castillo, I. and Mismar, R. (2018). “Empirical Bayes analysis of spike and slab posterior distributions.” *ArXiv:1801.01696 preprint*. MR3885271. doi: <https://doi.org/10.1214/18-EJS1494>. 940
- [14] Castillo, I. and Roquain, E. (2018). “On spike and slab empirical Bayes multiple testing.” 953
- [15] Castillo, I., Schmidt-Hieber, J., and van der Vaart, A. (2015). “Bayesian linear regression with sparse priors.” *Annals of Statistics*. MR3375874. doi: <https://doi.org/10.1214/15-AOS1334>. 933
- [16] Castillo, I. and Szabo, B. (2018). “Spike and slab empirical Bayes sparse credible sets.” *ArXiv:1808.07721 preprint*. MR4036030. doi: <https://doi.org/10.3150/19-BEJ1119>. 940, 953
- [17] Castillo, I. and van der Vaart, A. (2012). “Needles and Straw in a Haystack: Posterior concentration for possibly sparse sequences.” *Annals of Statistics*, 40(4): 2069–2101. MR3059077. doi: <https://doi.org/10.1214/12-AOS1029>. 934, 936, 937, 940, 944, 950, 951, 960
- [18] Clements, N., Sarkar, S. K., and Guo, W. (2012). “Astronomical transient detection controlling the false discovery rate.” In Feigelson, E. D. and Babu, G. J. (eds.), *Statistical Challenges in Modern Astronomy V*, 383–396. New York, NY: Springer New York. MR3220203. 933
- [19] Cook, J. D. (2011). “Basic properties of the soft maximum.” Working Paper Series, Working Paper 70, UT MD Anderson Cancer Center Department of Biostatistics. Available from www.johndcook.com/blog/articles/.
- [20] de Rooij, S. and van Erven, T. (2009). “Learning the switching rate by discretising Bernoulli sources online.” In *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics (AISTATS)*, 432–439. 935, 941, 942
- [21] Defour, D., Daramy, C., de Dinechin, F., Gallet, M., Gast, N., Lauter, C., and Muller, J.-M. (2016). *CRLibm C++ Library*. Version 1.0beta5. URL <https://github.com/taschini/crlibm> 945
- [22] Diaconis, P. and Freedman, D. (1980). “Finite exchangeable sequences.” *The Annals of Probability*, 8(4): 745–764. MR0577313.
- [23] Efron, B. (2012). *Large-Scale Inference: Empirical Bayes Methods for Estimation*,

- Testing, and Prediction*. Cambridge University Press. MR2724758. doi: <https://doi.org/10.1017/CB09780511761362>. 933
- [24] Fernandez, C., Ley, E., and Steel, M. F. (2001). “Benchmark priors for Bayesian model averaging.” *Journal of Econometrics*, 100(2): 381–427. MR1820410. doi: [https://doi.org/10.1016/S0304-4076\(00\)00076-2](https://doi.org/10.1016/S0304-4076(00)00076-2). 934
- [25] Freeman, M. and Tukey, J. (1950). “Transformations related to the Angular and the Square Root.” *Annals of Mathematical Statistics*, 21: 607–611. MR0038028. doi: <https://doi.org/10.1214/aoms/1177729756>. 941
- [26] George, E. I. and McCulloch, R. E. (1997). “Approaches for Bayesian variable selection.” *Statistica Sinica*, 339–373. 934, 935
- [27] Golubev, G. K. (2002). “Reconstruction of sparse vectors in white Gaussian noise.” *Problems of Information Transmission*, 38(1): 65–79. MR2101314. doi: <https://doi.org/10.1023/A:1020098307781>. 934
- [28] Guglielmetti, F., Fischer, R., and Dose, V. (2009). “Background–source separation in astronomical images with Bayesian probability theory – I. The method.” *Monthly Notices of the Royal Astronomical Society*, 396(1): 165–190. doi: <https://doi.org/10.1111/j.1365-2966.2009.14739.x>. 933
- [29] Higham, N. J. (2002). *Accuracy and Stability of Numerical Algorithms*. SIAM, 2nd edition. MR1927606. doi: <https://doi.org/10.1137/1.9780898718027>.
- [30] Johndrow, J. E. and Orenstein, P. (2017). “Scalable MCMC for Bayes Shrinkage Priors.” *ArXiv:1705.00841v2 preprint*. 935
- [31] Johnstone, I. M. and Silverman, B. W. (2004). “Needles and straw in haystacks: Empirical Bayes estimates of possibly sparse sequences.” *Annals of Statistics*, 32(4): 1594–1649. MR2089135. doi: <https://doi.org/10.1214/009053604000000030>. 933, 935, 940, 949
- [32] Karatsuba, A. A. and Ofman, Y. P. (1962). “Multiplication of many-digital numbers by automatic computers.” *Doklady Akademii Nauk*, 145(2): 293–294.
- [33] Kerns, G. J. and Székely, G. J. (2006). “Definetti’s theorem for abstract finite exchangeable sequences.” *Journal of Theoretical Probability*, 19(3): 589–608. MR2280511. doi: <https://doi.org/10.1007/s10959-006-0028-z>.
- [34] Knuth, D. E. (1997). *The Art of Computer Programming: Sorting and Searching*, volume 3. Pearson Education. MR3077154.
- [35] Koolen, W. M. and de Rooij, S. (2013). “Universal codes from switching strategies.” *IEEE Transactions on Information Theory*, 59(11): 7168–7185. MR3124635. doi: <https://doi.org/10.1109/TIT.2013.2273353>. 939
- [36] Krämer, A., Green, J., Pollard, J., Jack, and Tugendreich, S. (2013). “Causal analysis approaches in Ingenuity Pathway Analysis.” *Bioinformatics*, 30(4): 523–530. doi: <https://doi.org/10.1093/bioinformatics/btt703>. 933
- [37] Liang, F., Paulo, R., Molina, G., Clyde, M. A., and Berger, J. O. (2008). “Mix-

- tures of g-priors for Bayesian variable selection.” *Journal of the American Statistical Association*, 103(481): 410–423. MR2420243. doi: <https://doi.org/10.1198/016214507000001337>. 934
- [38] Logsdon, B. A., Hoffman, G. E., and Mezey, J. G. (2010). “A variational Bayes algorithm for fast and accurate multiple locus genome-wide association analysis.” *BMC Bioinformatics*, 11(1): 58. 946
- [39] Malone, B. M., Tan, F., Bridges, S. M., and Peng, Z. (2011). “Comparison of four ChIP-Seq analytical algorithms using rice endosperm H3K27 trimethylation profiling data.” *PLoS ONE*, 6(9): 1–12. doi: <https://doi.org/10.1371/journal.pone.0025260>. 933
- [40] Martin, R. and Walker, S. G. (2014). “Asymptotically minimax empirical Bayes estimation of a sparse normal mean vector.” *Electronic Journal of Statistics*, 8(2): 2188–2206. MR3273623. doi: <https://doi.org/10.1214/14-EJS949>. 935, 949
- [41] Martin, R. and Walker, S. G. (2020). *EBSparse R code for the sparse normal sequence model*. 3rd version (09/24/2014), retrieved April, 2020 from <https://www4.stat.ncsu.edu/~rmartin/research.html>. 949
- [42] Mitchell, T. J. and Beauchamp, J. J. (1988). “Bayesian variable selection in linear regression.” *Journal of the American Statistical Association*, 83(404): 1023–1032. MR0997578. 934
- [43] Papaspiliopoulos, O. and Rossell, D. (2017). “Bayesian block-diagonal variable selection and model averaging.” *Biometrika*, 104(2): 343–359. MR3698258. doi: <https://doi.org/10.1093/biomet/asx019>. 954
- [44] Pereyra, M. (2016). “Proximal Markov chain Monte Carlo algorithms.” *Statistics and Computing*, 26(4): 745–760. MR3515019. doi: <https://doi.org/10.1007/s11222-015-9567-4>. 935
- [45] Quackenbush, J. (2002). “Microarray data normalization and transformation.” *Nature Genetics*, 32(4s): 496. 933, 948
- [46] Rabiner, L. R. (1989). “A tutorial on hidden Markov models and selected applications in speech recognition.” *Proceedings of the IEEE*, 77(2): 257–285. 935, 938
- [47] Ray, K. and Szabo, B. (2019). “Variational Bayes for high-dimensional linear regression with sparse priors.” 946
- [48] Ročková, V. (2018). “Bayesian estimation of sparse signals with a continuous spike-and-slab prior.” *Annals of Statistics*, 46(1): 401–437. MR3766957. doi: <https://doi.org/10.1214/17-AOS1554>. 935, 949
- [49] Ročková, V. (2018). *Spike-and-Slab LASSO R code for the sparse normal sequence model*. Personal communication. 949, 960
- [50] Ročková, V. and George, E. I. (2014). “EMVS: The EM approach to Bayesian variable selection.” *Journal of the American Statistical Association*, 109(506): 828–846. MR3223753. doi: <https://doi.org/10.1080/01621459.2013.869223>. 935

- [51] Schmüdgen, K. (2017). *The Moment Problem*, volume 277. Springer. [MR3729411](#).
- [52] Schönhage, A. and Strassen, V. (1971). “Schnelle Multiplikation großer Zahlen.” *Computing*, 7(3-4): 281–292. [MR0292344](#). doi: <https://doi.org/10.1007/bf02242355>.
- [53] Scott, J. G. and Berger, J. O. (2006). “An exploration of aspects of Bayesian multiple testing.” *Journal of Statistical Planning and Inference*, 136(7): 2144–2162. [MR2235051](#). doi: <https://doi.org/10.1016/j.jspi.2005.08.031>. 934, 935, 940
- [54] Scott, J. G. and Berger, J. O. (2010). “Bayes and empirical-Bayes multiplicity adjustment in the variable-selection problem.” *Annals of Statistics*, 2587–2619. [MR2722450](#). doi: <https://doi.org/10.1214/10-AOS792>. 934, 951
- [55] Silverman, B. W., Evers, L., Xu, K., Carbonetto, P., and Stephens, M. (2017). *EbayesThresh: Empirical Bayes Thresholding and Related Methods*. R package version 1.4-12. URL <https://CRAN.R-project.org/package=EbayesThresh> [MR2364426](#). 949
- [56] Thomas, J. G., Olson, J. M., Tapscott, S. J., and Zhao, L. P. (2001). “An efficient and robust statistical modeling approach to discover differentially expressed genes using genomic expression profiles.” *Genome Research*, 11(7): 1227–1236. 933
- [57] Titsias, M. K. and Lázaro-Gredilla, M. (2011). “Spike and slab variational inference for multi-task and multiple kernel learning.” In *Advances in Neural Information Processing Systems*, 2339–2347. 946
- [58] van der Pas, S., Scott, J., Chakraborty, A., and Bhattacharya, A. (2016). *Horseshoe: Implementation of the Horseshoe Prior*. R package version 0.1.0. URL <https://CRAN.R-project.org/package=horseshoe> 935, 949
- [59] van der Pas, S., Szabó, B., and van der Vaart, A. (2017). “Adaptive posterior contraction rates for the horseshoe.” *Electronic Journal of Statistics*, 11(2). [MR3705450](#). doi: <https://doi.org/10.1214/17-EJS1316>. 935
- [60] van der Pas, S., Szabó, B., and van der Vaart, A. (2017). “Uncertainty quantification for the horseshoe (with discussion).” *Bayesian Analysis*, 12(4): 1221–1274. [MR3724985](#). doi: <https://doi.org/10.1214/17-BA1065>. 935, 949
- [61] van der Pas, S. L., Kleijn, B. J. K., and van der Vaart, A. W. (2014). “The horseshoe estimator: Posterior concentration around nearly black vectors.” *Electronic Journal of Statistics*, 8(2): 2585–2618. [MR3285877](#). doi: <https://doi.org/10.1214/14-EJS962>. 935
- [62] van Erven, T., de Rooij, S., and Szabo, B. (2019). *SequenceSpikeSlab: Exact Bayesian Model Selection Methods for the Sparse Normal Sequence Model*. R package version 0.1. URL <https://CRAN.R-project.org/package=SequenceSpikeSlab> 936, 949, 954
- [63] van Erven, T. and Szabo, B. (2020). “Supplementary Material of “Fast Exact

Bayesian Inference for Sparse Signals in the Normal Sequence Model”.” *Bayesian Analysis*. doi: <https://doi.org/10.1214/20-BA1227SUPP>. 936

- [64] Volf, P. and Willems, F. (1998). “Switching between two universal source coding algorithms.” In *Proceedings of the Data Compression Conference, Snowbird, Utah*, 491–500. 935, 937, 939
- [65] Zhou, M., Chen, H., Ren, L., Sapiro, G., Carin, L., and Paisley, J. W. (2009). “Non-parametric Bayesian dictionary learning for sparse image representations.” In *Advances in Neural Information Processing Systems 22 (NIPS)*, 2295–2303. 935

Acknowledgments

We would like to thank Steven de Rooij for performing the experiments with Karatsuba’s algorithm reported in Appendix A.3 in the supplemental material, for suggesting the use of interval arithmetic in the experiments, and for detailed discussions underlying Theorem 2.1. We further thank Ismaël Castillo for providing the R code for the experiments in [17] and Veronika Ročková for providing an R implementation of the Spike-and-Slab LASSO for the normal sequence model [49].