

Sparse Horseshoe Estimation via Expectation-Maximisation

Shu Yu Tew¹[0000–0002–6560–3131], Daniel F. Schmidt¹(✉)[0000–0002–1788–2375],
and Enes Makalic²[0000–0003–3017–0871]

¹ Department of Data Science and AI, Monash University, Clayton, Australia
{shu.tew, daniel.schmidt}@monash.edu

² Centre for Epidemiology and Biostatistics, The University of Melbourne, Australia
emakalic@unimelb.edu.au

Abstract. The horseshoe prior is known to possess many desirable properties for Bayesian estimation of sparse parameter vectors, yet its density function lacks an analytic form. As such, it is challenging to find a closed-form solution for the posterior mode. Conventional horseshoe estimators use the posterior mean to estimate the parameters, but these estimates are not sparse. We propose a novel expectation-maximisation (EM) procedure for computing the MAP estimates of the parameters in the case of the standard linear model. A particular strength of our approach is that the M-step depends only on the form of the prior and it is independent of the form of the likelihood. We introduce several simple modifications of this EM procedure that allow for straightforward extension to generalised linear models. In experiments performed on simulated and real data, our approach performs comparable, or superior to, state-of-the-art sparse estimation methods in terms of statistical performance and computational cost.

Keywords: Horseshoe Regression · Sparse Regression · Non-convex Penalised Regression · Maximum a Posteriori Estimation · Expectation-Maximisation.

1 Introduction

Sparse modelling has become increasingly important in statistical learning due to the growing demand for the analysis of high dimensional data with parameters whose dimension exceeds the sample size. Inference under such a setting often involves the assumption that the parameter vector is sparse. This implies that most of the components in the vector are insignificant and should be removed to achieve a good estimate. The lasso (i.e., ℓ_1 -penalised regression) [27] is a popular sparsity inducing technique well-known for its ability to simultaneously perform variable selection and coefficient estimation. The lasso is a convex penalisation technique, and is known to suffer from over-shrinkage and consistency issues [8]; to combat these, authors introduced non-convex penalty to replace the usual ℓ_1 -norm used in the lasso. Well-known non-convex penalties include the smoothly clipped absolute deviation (SCAD) [12] and the minimax concave penalty (MCP) [30].

However, these approaches are all frequentist in nature, usually relying on an additional principle such as cross-validation to select the degree of regularisation, and do not have immediate access to reliable measures of uncertainty to complement the point estimates they produce [8]. Bayesian inference, on the other hand, naturally quantifies uncertainty directly through the posterior distribution. Bayesian penalised regression offers readily available uncertainty estimates, automatic estimation of the penalty parameters, and more flexibility in terms of penalties that can be considered [28]. In this paper, we consider the standard Gaussian linear regression model

$$\mathbf{y} = \beta_0 \mathbf{1}_n + \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim N(\mathbf{0}, \sigma^2 \mathbf{I}_n) \quad (1)$$

where $\mathbf{y} = (y_1, \dots, y_n)^T \in \mathbb{R}^n$ is a vector of outcome variable, $\mathbf{X} \in \mathbb{R}^{n \times p}$ is a matrix of predictor variables, $\beta_0 \in \mathbb{R}$ is the intercept, and $\boldsymbol{\epsilon}$ is a vector of i.i.d. normally distributed random error with mean zero and unknown variance σ^2 . Given \mathbf{y} and \mathbf{X} , our goal is to accurately identify and estimate the non-zero components of the unknown regression coefficients $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p) \in \mathbb{R}^p$.

1.1 Bayesian Penalised Regression

Most penalised regression methods can be solved in the Bayesian framework by interpreting the estimates of $\boldsymbol{\beta}$ as the posterior point estimate of choice under an appropriate prior distribution [28]. For example, the lasso estimate can be interpreted as posterior mode estimates when the regression parameters $\boldsymbol{\beta}$ follow an independent and identical double exponential (Laplace) prior [27]. Motivated by this relationship, much work has been done in proposing different Bayesian approaches to the sparse estimation problem via a variety of sparsity inducing priors. Some popular examples include the Bayesian Lasso [19], the horseshoe estimator [7] and the normal-gamma estimator [6]. These estimators are frequently implemented using Monte Carlo Markov Chain (MCMC) algorithms to sample from the posterior distribution, and the resulting Bayesian estimates are usually posterior means or medians.

Interestingly, however, while these priors may be sparsity promoting, in the sense that they concentrate probability mass around sparse coefficient vectors, the resulting posterior means or medians will never themselves be sparse. Sparse estimates can be achieved by instead considering the posterior mode, i.e., maximum a posteriori (MAP) estimation [13, 1, 4]. Unfortunately, most implementations of Bayesian MAP-based sparse estimators are not flexible in the sense that they only work for linear models when paired with a specified prior. These methods do not easily generalise to other regression models and are difficult to apply to priors that lack an analytic form (e.g., horseshoe prior, Strawderman-Berger prior and normal-Gamma prior). Alternatively, simple thresholding (“sparsification”) rules for the posterior mean or median of β_j s can be used to produce sparse estimates, but these tend to lack theoretical justification, and inference can be highly sensitive to the choice of the threshold [16, 7]. To address these limitations:

1. we propose a novel expectation-maximisation (EM) procedure to solve for the *exact* posterior mode of a regression model under the horseshoe prior;

2. we introduce several simple modifications of our base EM procedure that allow for straightforward extension to models beyond the usual Gaussian linear model, e.g., generalised linear models.

Experiments on simulated data and real datasets in Section 6 demonstrate that our proposed Bayesian EM sparse horseshoe estimator can give comparable, if not better, performance when compared to state-of-the-art non-convex solvers in terms of statistical performance and computational cost. As far as the authors are aware, this is the first work that explores the posterior mode under the exact representation of the horseshoe prior, instead of using an approximated density function in place of the horseshoe prior.

2 Background and Related Work

The EM algorithm [10] is one of the most widely used methods for MAP estimation of sparse Bayesian linear models, with [13] and [15] being classic papers on this approach. A particular strength, in comparison to approximate methods such as variational Bayes [29], is that it is guaranteed to converge to exact posterior modes (a stationary point in the likelihood) whenever it can be applied [10]. By introducing appropriate latent variables, the hierarchical decomposition of the Laplace prior allows [13] for the derivation of an efficient EM algorithm to recover the lasso estimates. Following this work, many authors have attempted similar procedure to achieve sparse estimation using various priors and likelihood models that admit a scale mixture of normals (SMN) representation. This includes the lasso prior [19], the generalised double Pareto prior [1] and the horseshoe-like prior [4].

In this paper, we focus on the class of global-local shrinkage priors [21], and more specifically the horseshoe prior, which has been recognised as a good default prior choice for Bayesian sparse estimation [3, 7]. The horseshoe prior has a pole at $\beta_j = 0$, and heavy, Cauchy-like tails. These properties are desirable in sparse estimation because they allow small coefficients to be heavily shrunk towards zero while ensuring large coefficients are not over-shrunk. This is in contrast to the popular Bayesian lasso and Bayesian ridge hierarchies that apply the same amount of shrinkage to all coefficients and potentially over-shrink coefficients that are far from zero. The Gaussian linear regression model (1) corresponds to the following likelihood

$$\mathbf{y} | \mathbf{X}, \boldsymbol{\beta}, \sigma^2 \sim N_n(\mathbf{X}\boldsymbol{\beta}, \sigma^2 \mathbf{I}_n) \quad (2)$$

where $N_k(\cdot, \cdot)$ is the k -variate Gaussian distribution. In the class of global-local shrinkage priors, each β_j is assigned a continuous shrinkage prior centered at $\beta_j = 0$ that can be represented in an SMN form as

$$\begin{aligned} \beta_j | \tau^2, \lambda_j^2, \sigma^2 &\sim N(0, \tau^2 \lambda_j^2 \sigma^2) \\ \lambda_j^2 &\sim \pi(\lambda_j^2) d\lambda^2 \\ \tau^2 &\sim \pi(\tau^2) d\tau^2 \end{aligned} \quad (3)$$

where τ is the *global* shrinkage parameter that controls the overall degree of shrinkage, λ_j is the *local* shrinkage parameter associated with the j -th predictor and it controls the shrinkage for individual coefficients, and $\pi(\cdot)$ is an appropriate prior distribution of choice assigned to the shrinkage parameters.

Given the (unnormalised) joint posterior distribution of the hierarchy (2)-(3),

$$p(\boldsymbol{\beta}, \tau, \boldsymbol{\lambda} | \mathbf{y}) \propto p(\mathbf{y} | \boldsymbol{\beta}) \cdot \pi(\boldsymbol{\beta} | \boldsymbol{\lambda}, \tau) \cdot \pi(\boldsymbol{\lambda}) \cdot \pi(\tau), \quad (4)$$

the conventional EM approach to find the posterior mode estimates treats the hyperparameters (i.e., the shrinkage parameters) $\boldsymbol{\lambda}$ as “missing data” (i.e., the latent variables). This approach iteratively finds the expected values of the latent variables, and solves for the maximisation problem

$$\operatorname{argmax}_{\boldsymbol{\beta}} \text{E} [\log p(\boldsymbol{\beta}, \tau, \boldsymbol{\lambda} | \mathbf{y}) | \boldsymbol{\beta}]$$

to produce a sequence of estimates for the regression coefficients. This approach is effective because the scale-mixture form of the local-global prior means that conditional on $\boldsymbol{\lambda}$, the posterior for $\boldsymbol{\beta}$ is Gaussian and maximisation is (relatively) straightforward. In the case that the likelihood is non-Gaussian, but itself admits a scale-mixture-of-normals representation, such as the Laplace or logistic regression model, this approach can be adapted appropriately.

This approach is not easily applicable if the prior placed on $\boldsymbol{\beta}$ lacks a simple analytic form; in this case, the expected value of the shrinkage parameter will not have a closed form, and may be difficult or impossible to compute. One such prior is the horseshoe discussed previously. To attempt to address this problem [4] introduced a “horseshoe-like” prior that mimics the behavior of the horseshoe density proposed by [7]. This horseshoe-like prior has a closed form density function and a scale-mixture representation that allows for implementation of an EM algorithm to obtain the “horseshoe-like” MAP estimates.

3 Bayesian EM Sparse Linear Regression

In this section, we present a novel, general EM procedure to compute the MAP estimate for the linear model. The key innovation underlying our approach is to treat the coefficients, $\boldsymbol{\beta}$, as latent variables; this is in contrast to the usual application of the EM algorithm [13] that treats the shrinkage hyperparameters, $\boldsymbol{\lambda}$, as missing data.

As is standard in penalised regression, and without any loss of generality, we assume that the predictors are standardised to have mean zero and standard deviation one, and the target has a mean of zero, i.e., the estimate of the intercept is simply $\hat{\beta}_0 = (1/n) \sum y_i$. This means we can simply ignore the intercept when estimating the remaining coefficients $\boldsymbol{\beta}$.

3.1 The Basic EM Algorithm

Here we consider the linear regression model define in hierarchy (2)-(3). The proposed EM algorithm solves for the posterior mode estimates by iterating

through the following two steps until convergence is achieved:

E-step. We take the expectation of the complete negative log-posterior (with respect to the missing variable β), conditional on the current values of λ , τ^2 and σ^2 , and the observed data, \mathbf{y} ; the resulting quantity is called the ‘‘Q-function’’:

$$\begin{aligned}
 Q(\lambda, \tau, \sigma^2 | \hat{\lambda}^{(t)}, \hat{\tau}^{(t)}, \hat{\sigma}^{2(t)}) &= \mathbb{E} \left[-\log p(\beta, \lambda, \tau, \sigma^2 | \mathbf{y}) | \hat{\lambda}^{(t)}, \hat{\tau}^{(t)}, \hat{\sigma}^{2(t)}, \mathbf{y} \right] \\
 &= \left(\frac{n+p}{2} \right) \log \sigma^2 + \frac{\mathbb{E} \left[\|\mathbf{y} - \mathbf{X}\beta\|^2 | \hat{\lambda}^{(t)}, \hat{\tau}^{(t)}, \hat{\sigma}^{2(t)} \right]}{2\sigma^2} + \frac{p}{2} \log \tau^2 \\
 &\quad + \frac{1}{2} \sum_{j=1}^p \log \lambda_j^2 + \frac{1}{2\sigma^2\tau^2} \sum_{j=1}^p \frac{\mathbb{E} \left[\beta_j^2 | \hat{\lambda}^{(t)}, \hat{\tau}^{(t)}, \hat{\sigma}^{2(t)} \right]}{\lambda_j^2} - \log \pi(\lambda, \tau) \quad (5)
 \end{aligned}$$

where $\pi(\lambda, \tau)$ is the joint prior distribution for the hyperparameters. For notational simplicity we use $\mathbb{E} [\beta_j^2]$ and $\mathbb{E} [\|\mathbf{y} - \mathbf{X}\beta\|^2]$ for the conditional expectations of β_j^2 and the sum of squared residuals, respectively, throughout the sequel.

M-step. Update the parameter estimates by minimising the Q-function with respect to the shrinkage hyperparameters and noise variance, i.e.,

$$\{ \hat{\lambda}^{(t+1)}, \hat{\tau}^{(t+1)}, \hat{\sigma}^{2(t+1)} \} = \arg \min_{\lambda, \tau, \sigma^2} \left\{ Q \left(\lambda, \tau, \sigma^2 | \hat{\lambda}^{(t)}, \hat{\tau}^{(t)}, \hat{\sigma}^{2(t)} \right) \right\} \quad (6)$$

Implementation of this EM algorithm requires only knowledge of the negative log-prior of choice $-\log \pi(\lambda, \tau)$, the conditional expectations $\mathbb{E} [\beta^2]$ and $\mathbb{E} [\|\mathbf{y} - \mathbf{X}\beta\|^2]$. In Section 3.2, we discuss several different approaches to compute the conditional expectations for the E-step. This algorithm is quite general, with only the M-step depending on the choice of prior for the coefficients β_j . Application to the specific case of the horseshoe is discussed in Section 4.

The overall algorithm iterates the E-step and M-step until a convergence criterion is satisfied. Once convergence is achieved, we use the mode of the posterior distribution of β , conditional on the final values of $\hat{\lambda}^{(t)}$, $\hat{\tau}^{(t)}$ and $\hat{\sigma}^{2(t)}$ as our point estimate. Given that this conditional distribution is Gaussian, the mode is just the mean of the normal distribution (7) given by (9). Finally we set components of the final conditional posterior mode that are very small (smaller in absolute value than a small fraction of the standard error) to exactly zero. This step is technically not required, as given sufficient iterations the posterior mode will converge to be exactly sparse, but substantially reduces run time. In this paper we used $|\beta_j^{(t)}| < (5\sqrt{n})^{-1}$ as the threshold for all experiments.

3.2 Computing the conditional expectations

The conditional expected values $E[\beta^2]$ and $E[\|y - \mathbf{X}\beta\|^2]$ depend on the conditional posterior distribution of the regression coefficients $\beta \in \mathbb{R}^p$ [17]:

$$\begin{aligned} \beta | \lambda, \tau, \sigma, \mathbf{y} &\sim N_p(\mathbf{A}^{-1}\mathbf{X}^T\mathbf{y}, \sigma^2\mathbf{A}^{-1}) \\ \mathbf{A} &= (\mathbf{X}^T\mathbf{X} + \tau^{-2}\mathbf{A}^{-1}) \end{aligned} \quad (7)$$

where $\mathbf{A} = \text{diag}(\lambda_1^2, \dots, \lambda_p^2)$. One important point to note here is that the posterior distribution of β does not depend on the choice of prior applied to λ_j . This implies that regardless of the (marginal) prior assigned to β , as long as it has a SMN representation, both $p(\mathbf{y}|\beta)$ and $\pi(\beta|\lambda, \tau)$ will always be Gaussian densities, e.g., it does not matter whether the prior assigned to β is a Laplace prior (to solve for the Lasso) or a horseshoe prior, the E-step will remain the same. Changing the marginal priors on the regression coefficients only requires straightforward modification to the M-step for updates on the shrinkage parameters (see Section 4). This is an interesting advantage of our EM approach as the E-step is frequently very difficult to implement, particularly when dealing with conditional distributions that lack a standard density.

Exact expectations. The expected value of β_j^2 can be solved for by taking the sum of the variance and the square of the expected value of β_j :

$$E[\beta_j^2 | \lambda^{(t)}, \tau^{(t)}] = \text{Var}[\beta_j] + E[\beta_j]^2. \quad (8)$$

Due to the properties of Gaussian distributions, $\text{Var}[\beta_j] = (\text{Cov}[\beta])_{j,j}$, and

$$\begin{aligned} \text{Cov}[\beta] &= \sigma^2\mathbf{A}^{-1}, \\ E[\beta_j] &= (\mathbf{A}^{-1}\mathbf{X}^T\mathbf{y})_j. \end{aligned} \quad (9)$$

Direct computation of this expectation value is potentially computationally expensive and numerically unstable because it involves inverting the $p \times p$ matrix, \mathbf{A} . The efficient algorithm proposed in [24] avoids explicitly computing the inverse of \mathbf{A} when sampling from multivariate normal distributions of the form (7), and the approach in [5] is similar but utilises the matrix inversion lemma for improved computational complexity when $p > n$. Both of these algorithms can be trivially modified to compute the exact mean and covariance matrix of the conditional posterior distribution of the hierarchy (7).

Similarly, finding $E[\|y - \mathbf{X}\beta\|^2]$ involves solving for the expected value of a quadratic form of the regression coefficients, which in turns requires inverting the same matrix, \mathbf{A} . From standard results involving expectations of quadratic forms we have:

$$E[\|y - \mathbf{X}\beta\|^2] = \|y - \mathbf{X}E[\beta]\|^2 + \text{tr}(\mathbf{X}^T\mathbf{X} \cdot \text{Cov}[\beta]) \quad (10)$$

where $\text{tr}(\cdot)$ is the usual trace operator.

Approximate expectations. The two main components required to compute the expectations in the E-step for this linear model are the conditional posterior mean and variance of the regression coefficients β under the distribution (7). The conditional posterior mean can be found using the Rue’s [24] or Bhattacharya’s [5] algorithm without having to explicitly solve for the inverse of \mathbf{A} , with a simple extension of the same procedures allowing us to solve for the exact covariance matrix $\sigma^2 \mathbf{A}^{-1}$. However, the quantities that we need are the conditional variances (i.e., the diagonal elements of the conditional covariance matrix). Therefore, instead of inverting the full covariance matrix \mathbf{A} , the conditional variances may be approximately found by inverting only the diagonal elements of \mathbf{A} , i.e.,

$$\begin{aligned} \text{Var}[\beta_j] &\approx \sigma^2 \frac{1}{A_{j,j}} \\ &= \sigma^2 \left(\|\mathbf{x}_j\|^2 + \frac{1}{\tau^2 \lambda_j^2} \right)^{-1}. \end{aligned} \tag{11}$$

where \mathbf{x}_j is the j -th column of \mathbf{X} . This approximation takes only $O(p)$ operations. In the specific case that the predictors $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_p)$ are orthogonal, $\mathbf{X}^T \mathbf{X} = c \mathbf{I}_p$ and the approximation (11) will recover the exact variance of β_j . This approximation also applies to (10), i.e., instead of computing the matrix multiplication of the components in the trace function, one may instead consider only the diagonal elements of the matrix, yielding the approximation

$$\text{tr}(\mathbf{X}^T \mathbf{X} \cdot \text{Var}[\beta]) \approx \sigma^2 \sum_{j=1}^p \left(\|\mathbf{x}_j\|^2 + \frac{1}{\tau^2 \lambda_j^2} \right)^{-1} \|\mathbf{x}_j\|^2. \tag{12}$$

In comparison to the $O(p^2)$ operations required to compute the trace of the product of matrices in (10), this approximation requires only $O(p)$ operations.

4 Application to the Horseshoe Estimator

We now demonstrate the application of the EM procedure described in Section 3.1 to the horseshoe prior and provide the exact EM updates for each of the shrinkage parameters. Here, we assume no prior knowledge on the sparsity of the regression coefficient and assign the recommended default prior for the global variance parameter [14, 22]:

$$\tau \sim C^+(0, 1), \quad \tau \in (0, 1) \tag{13}$$

where $C^+(0, 1)$ denotes a standard half-Cauchy distribution with probability density function of

$$p(z) = \frac{2}{\pi(1+z^2)}, \quad z > 0. \tag{14}$$

We limit the range of τ to $(0, 1)$ following [20]. The degree of sparsity applied to individual regression coefficients now depends on the choice of the prior

distribution assigned to the local variance (shrinkage) component. Carvalho et al. [7] introduced the use of the horseshoe prior in sparse regression and demonstrated its robustness at handling sparsity with large signals. This is achieved by placing a half-Cauchy prior distribution over both the local and global shrinkage parameter. It was subsequently suggested that the horseshoe prior can be generalised by assigning an inverted-beta prior on $\lambda_j^2 \sim \beta'(a, b)$ with probability density function [22]

$$p(\lambda_j^2) = \frac{(\lambda_j^2)^{a-1}(1 + \lambda_j^2)^{-a-b}}{B(a, b)} \quad (15)$$

where $B(a, b)$ is the beta function, and $a > 0, b > 0$. We refer readers to [26] for further details about the effect of these hyperparameter a and b on the origin and tail properties of the generalised horseshoe prior. Our interest is in the particular case that $a = b = 1/2$ as this corresponds to a usual half Cauchy prior for λ_j . Substituting the negative logarithm of this prior distribution into (5) – (6) and holding τ^2 and σ^2 fixed yields the M-step update for each λ_j^2 :

$$\begin{aligned} \hat{\lambda}_j^2 &= \underset{\lambda_j^2}{\operatorname{argmin}} \left\{ \log \lambda_j^2 + \frac{W_j}{\lambda_j^2} + \log(1 + \lambda_j^2) \right\} \\ &= \frac{1}{4} \left(\sqrt{1 + 6W_j + W_j^2} + W_j - 1 \right) \end{aligned} \quad (16)$$

$$(\hat{\tau}^2, \hat{\sigma}^2) = \underset{\tau^2, \sigma^2}{\operatorname{argmin}} Q(\tau^2, \sigma^2; \hat{\lambda}^2) \quad (17)$$

where $W_j = E[\beta_j^2]/(2\sigma^2\tau^2)$. For convenience, the negative log-prior and the M-step update for the horseshoe prior are summarised in Table 1. Given the $\hat{\lambda}^2$ updates (16), the $\hat{\tau}$ and $\hat{\sigma}^2$ estimates can be found using numerical optimisation. This two-dimensional optimisation problem (17) can be further reduced to the following one-dimensional optimisation problem by (approximately) estimating σ^2 using the expected residual sum-of-squares given in Equation (10) or (11):

$$\begin{aligned} \hat{\sigma}^2 &= E[||\mathbf{y} - \mathbf{X}\boldsymbol{\beta}||^2] / n, \\ \hat{\tau}^2 &= \underset{\tau^2}{\operatorname{argmin}} Q(\tau^2; \hat{\lambda}^2, \hat{\sigma}^2). \end{aligned} \quad (18)$$

This approximate update for σ^2 is within $O(n^{-1})$ of the exact solution, and allows us to substantially reduce the complexity of the optimisation problem. Figure 1 compares the shrinkage profiles for the Lasso, horseshoe and horseshoe-like estimator. The plots demonstrate how the three estimators shrink the least-square (unpenalised) estimate $\hat{\beta}_{\text{LS}}$ towards zero. All three procedures shrink small least-squares estimates to zero. The lasso affects all values of $\hat{\beta}_{\text{LS}}$ by translating them towards zero by the same amount, while the horseshoe and horseshoe-like leave large values of $\hat{\beta}_{\text{LS}}$ largely unshrunk. The horseshoe-like prior exerts more of a hard thresholding-like behaviour as $\hat{\beta}_{\text{LS}}$ approaches zero, while the horseshoe prior mimics firm thresholding.

Table 1: Negative log-prior distribution and EM updates for the proposed horseshoe estimator, where $W_j = E[\beta_j^2]/(2\sigma^2\tau^2)$.

$-\log \pi(\boldsymbol{\lambda}, \tau)$	M-step
$\sum_{j=1}^p \left[\log(1 + \lambda_j^2) + \frac{\log \lambda_j^2}{2} \right] + \log(1 + \tau^2)$	$\hat{\lambda}_j^{2(t+1)} = \frac{\sqrt{1 + 6W_j + W_j^2} + W_j - 1}{4}$

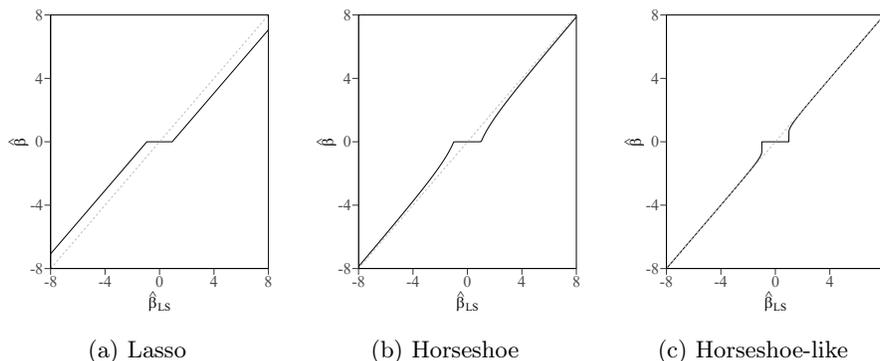


Fig. 1: The posterior mode estimates $\hat{\beta}$ versus $\hat{\beta}_{LS}$ for the (a) Lasso, (b) Horseshoe and (c) Horseshoe-like estimator. For illustration purposes, τ is chosen such that all three estimators give nearly identical shrinkage within approximately 1 unit of the origin.

5 Extension to Generalised Linear Models

We now consider an extension of the EM approach proposed in Section 3 to non-normal data models through the framework of generalised linear models (GLMs) [18]. A GLM models the conditional mean of the target by an appropriate (potentially) non-linear transformation of the linear predictor; well known examples include binomial logistic regression and Poisson log-linear regression. In general, when the targets are non-Gaussian, the conditional distribution of the coefficients will not have a standard form, and finding the exact expectations $E[\beta_j^2]$ is difficult. However, it is frequently the case that the conditional distribution can be approximated by a heteroskedastic Gaussian distribution

$$\begin{aligned} \boldsymbol{\beta} | \boldsymbol{\lambda}, \tau, \boldsymbol{\omega}, \mathbf{z} &\sim N_p(\mathbf{A}_\omega^{-1} \mathbf{X}^T \boldsymbol{\Omega} \mathbf{z}, \mathbf{A}_\omega^{-1}), \\ \mathbf{A}_\omega &= \mathbf{X}^T \boldsymbol{\Omega} \mathbf{X} + \tau^{-2} \mathbf{A}^{-1} \end{aligned} \tag{19}$$

where $\boldsymbol{\Omega} = \text{diag}(\boldsymbol{\omega})$, $\boldsymbol{\omega} = (\omega_1, \dots, \omega_n)$ is a vector of weights, and \mathbf{z} is an adjusted version of the targets, \mathbf{y} . Via standard central-limit theorem arguments, the accuracy of this approximation increases as the sample size n grows. The weights

can be obtained via a linearisation argument (i.e., the well known IRLS algorithm) or, preferably, via a scale-mixture of normals representation of the likelihood when available. For example, the logistic regression implementation used in this paper utilises the well-known Polyá-gamma representation of logistic regression [23]; under this scheme, the adjusted targets and weights are

$$z_i = (y - 1/2)/\omega_i, \quad \omega_i = \left(\frac{1}{2\eta_i} \right) \tanh \left(\frac{\eta_i}{2} \right),$$

where $\boldsymbol{\eta} = \mathbf{X}\boldsymbol{\beta} + \beta_0 \mathbf{1}_n$ is the linear predictor. Given appropriate weights and adjusted targets, one may simply approximate the conditional posterior-covariance of the coefficients by $\text{Cov}[\boldsymbol{\beta}] \approx \mathbf{A}\bar{\boldsymbol{\omega}}^{-1}$ and use either (9) or (11) to obtain approximate expressions for $\text{E}[\beta_j^2]$. Alternatively, one could potentially utilise a stochastic variant of the EM algorithm [9] to compute the required expectations, but due to space constraints we do not consider this further in this paper.

6 Experimental Results

We compare the performance of our proposed EM-based horseshoe estimator (HS-EM, using the exact expectations, and HS-apx, using the approximate expectations) against several state-of-the-art sparse methods including non-convex estimators (SCAD, MCP, HS-like estimator [4]), as well as the well known lasso and ridge estimator. We analyse the performance of the estimators in terms of variables selected and prediction accuracy on both simulated and real data. The experiments in this section are designed for variable selection under high dimensional settings, and we use the following metrics:

- **MSE**. Mean squared prediction error; given by $(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})' \boldsymbol{\Sigma} (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})$ for simulated data experiments, and $(1/n) \|\hat{\mathbf{y}} - \mathbf{y}\|_2^2$ for real data experiments.
- **Time**. Computation time (in seconds).
- **No.V**. Number of variables included in the model.
- **TNZ**. True non-zeros. Number of non-zero coefficients correctly identified.
- **FNZ**. False non-zeros. Number of zero coefficients incorrectly identified as non-zero.

Our proposed estimator terminates when it satisfies the convergence criterion: $\sum_{j=1}^p (|\beta_j^{(t)} - \beta_j^{(t+1)}|) / (1 + \sum_{j=1}^p (|\beta_j^{(t+1)}|)) < \omega$, where $\beta^{(t)}$ denotes the coefficient estimates at iteration t and ω is the tolerance parameter which we set it to 10^{-5} . All experiments are performed in the R statistical platform. The proposed EM horseshoe estimator is implemented in the **bayesreg**³ Bayesian regression package. We use the **glmnet** package to implement the lasso and ridge, while the **ncvreg** package implements both MCP and SCAD. The hyperparameters of the non-Bayesian techniques are tuned using ten-fold cross-validation. All function arguments are set to the default value unless mentioned otherwise. Datasets and code for the experimental results in this section are publicly available⁴.

³ Available at <https://cran.r-project.org/web/packages/bayesreg/index.html>

⁴ Available at <https://github.com/shuyu-tew/Sparse-Horseshoe-EM.git>

6.1 Simulated Data

We considered $n = 70$, $p = 350$ and simulated 100 different data sets from the linear model, $\mathbf{y} \sim N_n(\mathbf{X}\boldsymbol{\beta}, \sigma^2 \mathbf{I}_n)$ with $\sigma^2 \in \{1, 9\}$. The predictor matrix, $\mathbf{X} \in \mathbb{R}^{n \times p}$ was generated from a correlated multivariate normal distribution $N_p(0, \boldsymbol{\Sigma})$ with $\Sigma_{ij} = \rho^{|i-j|}$ and $\rho \in \{0, 0.8\}$.

The first ten entries of $\boldsymbol{\beta}$ were set to 3, the next ten were set to -3 , and the remaining entries were zero. This setup is similar to the experiments presented in [4]. The results are shown in Table 2. Overall, the HS-like estimator performed the best, obtaining the lowest MSE at the expense of substantial overfitting. The good performance in MSE under this setting is not unexpected, as this problem setting favors estimators that overfit (i.e., are less conservative) to a certain extent. Given a sparse vector with only 20 non-zero coefficients, this setting in which all non-zero coefficients have the same magnitude is one of the hardest settings for a sparse method, as excluding any one of the true coefficients is equally damaging. The prediction risk of this problem is $E[\|\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}\|^2]$; if a method incorrectly includes an irrelevant predictor in the model, the risk will approximately be the variance of the estimate associated with this variable, i.e., $E[\hat{\beta}_j^2] \approx \sigma^2/n$; in contrast, incorrectly excluding non-zero predictors will incur an error of 3^2 , which is much larger than $\sigma^2/70$ even for the setting of greater noise, i.e., $\sigma^2 = 9$. In general, approximately half of the coefficients included in the HS-like model are false non-zeros (FNZ); while most of the coefficients included in our HS models are true non-zeros with consistent small difference between TNZ and No.V.

The HS-EM, SCAD and MCP methods all generally select much fewer variables than the HS-like estimator, with the HS-EM method overfitting less than HS-like and underfitting more than MCP or SCAD. In terms of MSE, our HS-EM method is competitive with, or superior to SCAD and MCP for all settings, despite producing substantially sparser estimates. It appears that both SCAD and MCP are sensitive to the correlation structure of the data; their performance is comparable to our HS-EM approach when there is no correlation, but when correlation is introduced to the data, both methods performed substantially worse in terms of MSE. Such behaviour is not observed in our proposed HS estimator. Of the two proposed HS estimators, HS-apx appears to be approximately three times faster than HS-EM. Otherwise, both estimators give roughly similar performance.

6.2 Real data

We further analyse the performance of our proposed method on six real-world datasets. All datasets are available for download from the UCI machine learning repository [2] unless mentioned otherwise. Each dataset is randomly split such that the training data has a sample size of n with the remaining $N - n$ datapoints used as testing samples. This procedure is repeated 50 times and the averaged summary statistics for the performance measures are presented in Tables 4 and 5 for the linear regression and logistic regression experiments, respectively.

Table 2: Performance of different sparse estimators (with associated standard error in parentheses) for the simulated data experiments.

	HS-EM	HS-apx	HS-like	Lasso	MCP	SCAD	Ridge
$(\rho = 0, \sigma^2 = 1)$							
MSE	162.7(3.5)	162.8(2.9)	119.8(4.7)	139.1(3.4)	159.6(1.9)	149.8(1.8)	176.6(0.5)
Time	5.93(0.19)	1.37(0.04)	0.27(0.02)	0.14(0.01)	0.21(0.01)	0.32(0.01)	0.64(0.01)
No.V	6.11(0.35)	5.83(0.36)	43.4(0.53)	21.4(1.59)	8.61(0.52)	18.7(0.72)	350(0.00)
TNZ	4.37(0.29)	4.27(0.28)	14.2(0.32)	8.77(0.61)	4.89(0.29)	8.31(0.32)	20.0(0.00)
FNZ	1.63(0.16)	1.56(0.16)	29.2(0.72)	12.6(1.06)	3.72(0.28)	10.4(0.51)	330(0.00)
$(\rho = 0, \sigma^2 = 9)$							
MSE	171.6(3.1)	170.2(2.9)	141.5(3.5)	151.9(2.7)	164.2(1.9)	153.6(1.9)	177.1(0.5)
Time	5.89(0.22)	1.37(0.04)	0.27(0.01)	0.15(0.01)	0.21(0.01)	0.33(0.01)	0.64(0.01)
No.V	5.38(0.34)	5.19(0.34)	46.3(0.29)	15.9(1.47)	7.83(0.49)	18.4(0.77)	350(0.00)
TNZ	3.69(0.27)	3.68(0.27)	13.2(0.31)	6.61(0.54)	4.31(0.28)	7.58(0.31)	20.0(0.00)
FNZ	1.69(0.17)	1.51(0.15)	33.1(0.36)	9.35(0.99)	3.52(0.29)	10.8(0.56)	330(0.00)
$(\rho = 0.7, \sigma^2 = 1)$							
MSE	12.7(1.29)	15.1(1.57)	3.79(0.57)	6.48(0.47)	78.1(3.23)	74.3(3.01)	426.6(3.81)
Time	2.83(0.21)	0.97(0.05)	0.22(0.01)	0.11(0.01)	0.14(0.01)	0.16(0.01)	0.65(0.01)
No.V	16.9(0.25)	16.4(0.29)	25.1(0.46)	33.4(0.49)	14.1(0.33)	22.4(0.49)	350(0.00)
TNZ	16.9(0.25)	16.4(0.29)	19.7(0.05)	19.7(0.06)	8.21(0.16)	9.64(0.16)	20.0(0.00)
FNZ	0.04(0.03)	0.04(0.03)	5.27(0.51)	13.7(0.48)	5.89(0.31)	12.8(0.49)	330(0.00)
$(\rho = 0.7, \sigma^2 = 9)$							
MSE	34.8(1.79)	38.1(1.91)	36.3(1.38)	27.1(1.17)	85.9(3.28)	84.8(3.21)	428.7(3.89)
Time	4.41(0.26)	1.17(0.04)	0.24(0.01)	0.12(0.01)	0.14(0.01)	0.16(0.01)	0.64(0.01)
No.V	12.9(0.27)	12.3(0.28)	40.9(0.28)	28.6(0.62)	13.3(0.31)	21.1(0.45)	350(0.00)
TNZ	12.6(0.27)	12.1(0.27)	18.7(0.11)	18.4(0.08)	7.66(0.13)	8.85(0.16)	20.0(0.00)
FNZ	0.24(0.05)	0.24(0.05)	22.2(0.29)	10.3(0.59)	5.61(0.28)	12.2(0.42)	330(0.00)

Linear Regression Table 3 presents the posterior mean estimates and their 95% credible intervals (CI) for the diabetes data [11] ($N = 442, P = 10$). For comparison, we included the posterior mode estimates computed from the proposed HS-EM estimator. The HS-EM posterior mode estimates are very similar to the corresponding posterior means, and all posterior mode estimates are within the corresponding 95% credible intervals; however, the HS-EM posterior mode

Table 3: The posterior mode and mean estimates of the predictors for the diabetes data using the horseshoe estimator. The posterior mean and 95% credible intervals are computed using the `bayesreg` package in R.

	AGE	SEX	BMI	BP	S1	S2	S3	S4	S5	S6
Mean	-0.009	-18.68	5.769	1.034	-0.223	0.013	-0.592	2.419	48.84	0.179
2.50%	-0.341	-30.93	4.371	0.571	-0.937	-0.342	-1.415	-3.462	32.24	-0.225
97.5%	0.326	-5.144	7.109	1.457	0.098	0.656	0.189	11.36	70.14	0.734
Mode	.	-17.54	5.741	1.021	.	.	-0.909	.	43.58	.

provides a sparse point estimate as it excludes all variables (except S3) with 95% CIs that include zero.

In addition to the diabetes data, we also analysed the benchmark Boston housing data ($N = 506, P = 14$), the concrete compressive strength dataset ($N = 1030, P = 9$) and the eye data ($N = 120, P = 200$) [25]. To make the problem more difficult, for each dataset we added a number of noise variables generated using the same procedure described in Section 6.1; in all cases we added $p = 15$ additional noise variables, with $\rho = 0.8$. Model fitting is done using all the $P + 15$ (original plus noise) predictors as well as all interactions and possible transformations of the variables (logs, squares and cubics).

Overall, our proposed HS estimator performs the best, attaining the lowest MSE on all real data, while generally selecting the simplest models (included the lowest number of variables). Similar to the results observed in Section 6.1, the results for HS-apx are virtually identical to the exact HS-EM procedure, while being substantially faster. SCAD and MCP exhibit comparable performance to HS in terms of prediction error for the diabetes data and eye data, but have relatively poor performance for Boston and concrete data. Interestingly, while HS-like dominated the other non-convex estimators in terms of MSE on the simulated experiments, it performs quite poorly on the real data analysis, particularly on the diabetes and Boston housing data. The HS-like estimator tended to include the highest number of variables, suggesting it is potentially prone to overfitting as discussed in Section 6.1. Ridge regression generally performs worse than the sparse estimators, as would be expected given the experimental design.

Logistic Regression We also tested the HS-EM estimator on two binary classification problems: the Pima indians data ($N = 768, P = 8$) and the heart disease data ($N = 302, P = 14$). Similar to the linear regression analysis, we augment the data with 10 noise variables and model the predictors together with their interactions and transformations. For the heart disease data, we only included the interactions between the variables and not the transformations, as

Table 4: Performance of different sparse estimators (with associated standard error in parentheses) on the 4 datasets for linear regression model.

	HS-EM	HS-apx	HS-like	Lasso	MCP	SCAD	Ridge
Diabetes ($n = 100, p = 385$)							
MSE	3383.3 (30.5)	3407.2(32.1)	13068(491.1)	3654.4(47.9)	3624.4(66.2)	3667.6(68.5)	4181.6(42.8)
Time	3.01 (0.02)	1.02 (0.02)	0.32 (0.01)	0.29 (0.01)	1.13 (0.01)	1.47 (0.02)	0.74 (0.01)
No.V	1.62 (0.09)	1.54 (0.09)	95.3 (0.29)	4.14 (0.44)	3.68 (0.41)	6.86 (0.74)	385 (0.00)
Boston Housing ($n = 100, p = 473$)							
MSE	26.76(0.71)	26.72 (0.71)	58.01(3.05)	31.41(0.89)	293.1(24.3)	323.9(29.7)	49.19(1.19)
Time	4.74(0.34)	1.86(0.03)	0.67(0.03)	0.20(0.01)	1.21(0.01)	1.39(0.01)	1.06(0.01)
No.V	2.82(0.16)	2.84(0.16)	47.9(0.78)	4.60(0.52)	4.52(0.41)	10.9(0.87)	473(0.00)
Concrete ($n = 100, p = 327$)							
MSE	73.71 (3.67)	73.76 (3.66)	222.4 (8.38)	81.67 (1.98)	113.9 (15.9)	108.3 (16.1)	176.2 (2.49)
Time	2.39 (0.17)	0.86 (0.03)	0.36 (0.02)	0.23 (0.01)	0.98 (0.01)	1.14 (0.01)	0.55 (0.01)
No.V	5.46 (0.13)	5.42 (0.14)	67.8 (0.57)	10.7 (0.61)	6.60 (0.44)	13.8 (0.85)	327 (0.00)
Eye ($n = 100, p = 200$)							
MSE	0.79 (0.06)	0.79 (0.06)	1.50 (0.23)	0.84 (0.14)	0.79 (0.06)	0.86 (0.09)	0.85 (0.09)
Time	0.52 (0.02)	0.24 (0.01)	0.15 (0.01)	0.20 (0.01)	0.13 (0.01)	0.19 (0.01)	0.23 (0.01)
No.V	3.84 (0.11)	3.72 (0.09)	0.00 (0.00)	18.0 (0.55)	5.20 (0.27)	10.5 (0.42)	200 (0.00)

there were a substantial number of categorical predictors. For this experiment we varied the number of training samples $n \in \{100, 200\}$.

The HS-EM estimator obtains the best classification accuracy and negative log-loss on the Pima data for both training sample sizes. Ridge regression, on the other hand, performs the best in terms of classification accuracy for the heart data. This suggest that many of the predictors and interactions between the variables in the heart data are potentially correlated with the response variable. Therefore, a sparse estimator might not be suitable for modelling this data. Nonetheless, our HS methods gives comparable results to the other sparse methods in terms of log-loss and in general, as the sample size increases, the difference in performance between all estimators becomes less significant.

Table 5: Performance of different sparse estimators on the Pima and heart disease data. CA is the classification accuracy and NLL is the negative log-loss.

Dataset	HS-EM	HS-apx	Lasso	MCP	SCAD	Ridge
$(n = 100, p = 214)$						
CA	0.743 (0.01)	0.741(0.01)	0.704(0.01)	0.741(0.01)	0.735(0.01)	0.674(0.01)
NLL	358.1 (2.42)	358.8(2.55)	381.5(3.33)	360.9(4.12)	362.8(3.90)	400.7(2.63)
Time	0.622(0.02)	0.533(0.01)	0.298(0.01)	1.475(0.07)	2.056(0.07)	0.251(0.01)
No.V	1.760(0.09)	1.700(0.09)	3.780(0.25)	3.900(0.27)	8.480(0.59)	214.0(0.00)
Pima						
$(n = 200, p = 214)$						
CA	0.756 (0.01)	0.756 (0.01)	0.736(0.01)	0.754(0.01)	0.756 (0.01)	0.706(0.01)
NLL	290.8 (1.97)	291.1(1.96)	306.4(1.52)	292.7(2.32)	291.9(1.69)	320.6(1.01)
Time	0.838(0.03)	0.733(0.02)	0.835(0.02)	3.150(0.16)	4.075(0.19)	0.424(0.01)
No.V	3.020(0.15)	3.060(0.15)	4.740(0.25)	4.760(0.32)	8.980(0.59)	214.0(0.00)
$(n = 100, p = 400)$						
CA	0.758(0.01)	0.751(0.01)	0.792(0.01)	0.786(0.01)	0.797(0.01)	0.799 (0.01)
NLL	107.4(1.73)	110.3(1.66)	99.89(0.95)	96.20(1.28)	94.36 (0.93)	103.2(1.03)
Time	3.943(0.09)	3.623(0.11)	0.237(0.01)	1.208(0.05)	1.483(0.06)	0.383(0.01)
No.V	2.820(0.11)	2.800(0.12)	10.10(0.49)	7.080(0.28)	13.76(0.51)	400.0(0.00)
Heart						
$(n = 200, p = 400)$						
CA	0.798(0.01)	0.788(0.01)	0.813(0.01)	0.802(0.01)	0.805(0.01)	0.828 (0.01)
NLL	47.86(0.95)	49.53(0.94)	46.31(0.41)	44.90 (0.71)	45.05(0.73)	48.25(0.32)
Time	5.629(0.14)	5.022(0.14)	0.570(0.01)	2.755(0.11)	3.517(0.12)	0.621(0.01)
No.V	5.520(0.15)	5.480(0.14)	14.28(0.67)	9.340(0.37)	17.38(0.58)	400.0(0.00)

7 Discussion

In this paper we introduced a novel EM algorithm that allows us to efficiently find, for the first time, the exact, sparse posterior mode of the horseshoe estimator. The experimental results suggest that in comparison with state-of-the-art non-convex sparse estimators, the HS-EM estimator is quite robust to the underlying structure of the problem. Both the MCP and SCAD algorithm appear sensitive to correlation in the predictors, and the HS-like estimator, based on an approximation to the horseshoe prior, appears sensitive to the signal-to-noise ratio of the problem. In contrast, the HS-EM algorithm, even when not performing the best, appears to always remain competitive with the best performing method

while selecting highly sparse models. Comparing the shrinkage profiles of the HS and HS-like estimator (see Figure 1), the mixed performance of the HS-like procedure, is possibly due to the fact that HS-like approach more aggressively zeros out coefficients, making it more sensitive to misspecification of the global shrinkage parameter. We also conjecture that the conventional EM framework, in which the shrinkage hyperparameters λ are treated as missing data, provides a poorer basis for estimation of the global hyperparameter, in comparison to our approach of treating the coefficients as missing data, which potentially integrates the uncertainty of the parameters more effectively into the Q-function. This is a topic for further investigation.

Due to its prior preference for dense models, ridge regression can perform much worse than sparse alternatives if the underlying problem is sparse. However, it remained competitive or superior to other sparse methods on several of the real datasets. This strongly suggests that a method that is adaptive to the unknown degree of sparsity would be beneficial; we anticipate extending our algorithm to the generalised horseshoe discussed in Section 4, and automatically tuning our prior to produce a technique that will adapt to the sparsity of the problem.

While our HS-EM algorithm is able to locate, exactly, a sparse posterior mode of the horseshoe estimator, there is an obvious question as to what mode we are finding. Posterior modes (and means) are not invariant under reparameterisation, which suggests that maximising for λ_j (with appropriate transformation of prior distribution) in place of λ_j^2 , for example, will produce a new shrinkage procedure with potentially different properties. This is a focus of future research. The EM algorithm in this paper can also be easily extended to other popular priors such as the Bayesian lasso. Finally, given the relative simplicity of the M-step updates in this framework, it is of interest to consider extensions to non-linear models, such as neural networks. In this case, the required conditional expectations are unlikely to be available in closed form, but given access to a Gibbs sampler (or approximate Gibbs sampler) the procedure can easily be extended to utilise a standard stochastic EM implementation. This could potentially provide a simple, and (in expectation) exact, alternative to variational Bayes for these problems.

References

1. Armagan, A., Dunson, D., Lee, J.: Bayesian generalized double Pareto shrinkage. *Biometrika* (2010)
2. Asuncion, A., Newman, D.: UCI machine learning repository (2007)
3. Bhadra, A., Datta, J., Polson, N.G., Willard, B.: Default Bayesian analysis with global-local shrinkage priors. *Biometrika* **103**(4), 955–969 (2016)
4. Bhadra, A., Datta, J., Polson, N.G., Willard, B.T.: The horseshoe-like regularization for feature subset selection. *Sankhya B* pp. 1–30 (2019)
5. Bhattacharya, A., Chakraborty, A., Mallick, B.K.: Fast sampling with Gaussian scale mixture priors in high-dimensional regression. *Biometrika* p. asw042 (2016)
6. Brown, P.J., Griffin, J.E.: Inference with normal-gamma prior distributions in regression problems. *Bayesian analysis* **5**(1), 171–188 (2010)
7. Carvalho, C.M., Polson, N.G., Scott, J.G.: The horseshoe estimator for sparse signals. *Biometrika* **97**(2), 465–480 (2010)

8. Casella, G., Ghosh, M., Gill, J., Kyung, M.: Penalized regression, standard errors, and Bayesian lassos. *Bayesian analysis* **5**(2), 369–411 (2010)
9. Celeux, G.: The SEM algorithm: a probabilistic teacher algorithm derived from the EM algorithm for the mixture problem. *Comput Stat quarterly* **2**, 73–82 (1985)
10. Dempster, A.P., Laird, N.M., Rubin, D.B.: Maximum likelihood from incomplete data via the EM algorithm. *J R Stat Soc Series B Stat Methodol* **39**(1), 1–22 (1977)
11. Efron, B., Hastie, T., Johnstone, I., Tibshirani, R.: Least angle regression. *The Annals of statistics* **32**(2), 407–499 (2004)
12. Fan, J., Li, R.: Variable selection via nonconcave penalized likelihood and its oracle properties. *J Am Stat Assoc* **96**(456), 1348–1360 (2001)
13. Figueiredo, M.A.: Adaptive sparseness for supervised learning. *IEEE transactions on pattern analysis and machine intelligence* **25**(9), 1150–1159 (2003)
14. Gelman, A.: Prior distributions for variance parameters in hierarchical models. *Bayesian Analysis* **1**(3), 515–533 (2006)
15. Kiiveri, H.T.: A Bayesian approach to variable selection when the number of variables is very large. *Lecture Notes-Monograph Series* pp. 127–143 (2003)
16. Li, H., Pati, D.: Variable selection using shrinkage priors. *Computational Statistics & Data Analysis* **107**, 107–119 (2017)
17. Makalic, E., Schmidt, D.F.: A simple sampler for the horseshoe estimator. *IEEE Signal Processing Letters* **23**(1), 179–182 (2016)
18. Nelder, J.A., Wedderburn, R.W.: Generalized linear models. *Journal of the Royal Statistical Society: Series A (General)* **135**(3), 370–384 (1972)
19. Park, T., Casella, G.: The Bayesian lasso. *J Am Stat Assoc* **103**(482), 681–686 (2008)
20. van der Pas, S., Szabó, B., van der Vaart, A.: Adaptive posterior contraction rates for the horseshoe. *Electronic Journal of Statistics* **11**(2), 3196–3225 (2017)
21. Polson, N.G., Scott, J.G.: Shrink globally, act locally: Sparse Bayesian regularization and prediction. *Bayesian statistics* **9**, 501–538 (2010)
22. Polson, N.G., Scott, J.G.: On the half-Cauchy prior for a global scale parameter. *Bayesian Analysis* **7**(4), 887–902 (2012)
23. Polson, N.G., Scott, J.G., Windle, J.: Bayesian inference for logistic models using Pólya-gamma latent variables. *Journal of the American statistical Association* **108**(504), 1339–1349 (2013)
24. Rue, H.: Fast sampling of Gaussian Markov random fields. *J R Stat Soc Series B Stat Methodol* **63**(2), 325–338 (2001)
25. Scheetz, T.E., Kim, K.Y.A., Swiderski, R.E., Philp, A.R., Braun, T.A., Knudtson, K.L., Dorrance, A.M., DiBona, G.F., Huang, J., Casavant, T.L., et al.: Regulation of gene expression in the mammalian eye and its relevance to eye disease. *Proceedings of the National Academy of Sciences* **103**(39), 14429–14434 (2006)
26. Schmidt, D.F., Makalic, E.: Bayesian generalized horseshoe estimation of generalized linear models. In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. pp. 598–613. Springer (2019)
27. Tibshirani, R.: Regression shrinkage and selection via the lasso. *J R Stat Soc Series B Stat Methodol* **58**(1), 267–288 (1996)
28. Van Erp, S., Oberski, D.L., Mulder, J.: Shrinkage priors for Bayesian penalized regression. *Journal of Mathematical Psychology* **89**, 31–50 (2019)
29. Wand, M.P., Ormerod, J.T., Padoan, S.A., Frühwirth, R.: Mean field variational Bayes for elaborate distributions. *Bayesian Analysis* **6**(4), 847–900 (2011)
30. Zhang, C.H., et al.: Nearly unbiased variable selection under minimax concave penalty. *The Annals of statistics* **38**(2), 894–942 (2010)