
Contents

1	Spike-and-slab meets LASSO: A review of the spike-and-slab LASSO	3
	<i>Ray Bai, Veronika Ročková, and Edward I. George</i>	
1.1	Introduction	4
1.2	Variable selection in high dimensions: Frequentist and Bayesian strategies	5
1.2.1	Penalized likelihood approaches	5
1.2.2	Spike-and-slab priors	5
1.3	The spike-and-slab LASSO	7
1.3.1	Prior specification	7
1.3.2	Selective shrinkage and self-adaptivity to sparsity	8
1.4	Computational details	10
1.4.1	Coordinate-wise optimization	10
1.4.2	Dynamic posterior exploration	11
1.4.3	EM implementation of the spike-and-slab LASSO	12
1.5	Uncertainty quantification	13
1.5.1	Debiasing the posterior mode	13
1.5.2	Posterior sampling for the spike-and-slab LASSO	13
1.6	Illustrations	14
1.6.1	Example on synthetic data	14
1.6.2	Bardet-Beidl syndrome gene expression study	15
1.7	Methodological extensions	16
1.8	Theoretical properties	20
1.9	Discussion	21
	Bibliography	23



1

Spike-and-slab meets LASSO: A review of the spike-and-slab LASSO

Ray Bai

The University of South Carolina (USA)

Veronika Ročková

The University of Chicago (USA)

Edward I. George

The University of Pennsylvania (USA)

CONTENTS

1.1	Introduction	4
1.2	Variable selection in high dimensions: Frequentist and Bayesian strategies	4
	1.2.1 Penalized likelihood approaches	5
	1.2.2 Spike-and-slab priors	5
1.3	The spike-and-slab LASSO	7
	1.3.1 Prior specification	7
	1.3.2 Selective shrinkage and self-adaptivity to sparsity	8
1.4	Computational details	10
	1.4.1 Coordinate-wise optimization	10
	1.4.2 Dynamic posterior exploration	11
	1.4.3 EM implementation of the spike-and-slab LASSO	12
1.5	Uncertainty quantification	12
	1.5.1 Debiasing the posterior mode	13
	1.5.2 Posterior sampling for the spike-and-slab LASSO	13
1.6	Illustrations	14
	1.6.1 Example on synthetic data	14
	1.6.2 Bardet-Beidl syndrome gene expression study	15
1.7	Methodological extensions	16
1.8	Theoretical properties	20
1.9	Discussion	21
	Acknowledgments	21

High-dimensional data sets have become ubiquitous in the past few decades, often with many more covariates than observations. In the frequentist setting, penalized likelihood methods are the most popular approach for variable selection and estimation in high-dimensional data. In the Bayesian framework, spike-and-slab methods are commonly used as probabilistic constructs for high-dimensional modeling. Within the context of linear regression, [48] introduced the spike-and-slab LASSO (SSL), an approach based on a prior which provides

a continuum between the penalized likelihood LASSO and the Bayesian point-mass spike-and-slab formulations. Since its inception, the spike-and-slab LASSO has been extended to a variety of contexts, including generalized linear models, factor analysis, graphical models, and nonparametric regression. The goal of this chapter is to survey the landscape surrounding spike-and-slab LASSO methodology. First we elucidate the attractive properties and the computational tractability of SSL priors in high dimensions. We then review methodological developments of the SSL and outline several theoretical developments. We illustrate the methodology on both simulated and real datasets.

1.1 Introduction

High-dimensional data are now routinely analyzed. In these settings, one often wants to impose a low-dimensional structure such as sparsity. For example, in astronomy and other image processing contexts, there may be thousands of noisy observations of image pixels, but only a small number of these pixels are typically needed to recover the objects of interest [25, 26]. In genetic studies, scientists routinely observe tens of thousands of gene expression data points, but only a few genes may be significantly associated with a phenotype of interest. For example, [59] has confirmed that only seven genes have a non-negligible association with Type I diabetes. Among practitioners, the main objectives in these scenarios are typically: a) identification (or *variable selection*) of the non-negligible variables, and b) *estimation* of their effects.

A well-studied model for sparse recovery in the high-dimensional statistics literature is the normal linear regression model,

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim \mathcal{N}_n(\mathbf{0}, \sigma^2 \mathbf{I}_n), \quad (1.1)$$

where $\mathbf{y} \in \mathbb{R}^n$ is a vector of n responses, $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_p] \in \mathbb{R}^{n \times p}$ is a design matrix of p potential covariates, $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)^T$ is a p -dimensional vector of unknown regression coefficients, and $\boldsymbol{\varepsilon}$ is the noise vector. When $p > n$, we often assume that most of the elements in $\boldsymbol{\beta}$ are zero or negligible. Under this setup, there have been a large number of methods proposed for selecting and estimating the active coefficients in $\boldsymbol{\beta}$. In the frequentist framework, penalized likelihood approaches such as the least absolute shrinkage and selection operator (LASSO) [56] are typically used to achieve sparse recovery for $\boldsymbol{\beta}$. In the Bayesian framework, spike-and-slab priors are a popular approach for sparse modeling of $\boldsymbol{\beta}$.

The spike-and-slab LASSO (SSL), introduced by [48], forms a continuum between these penalized likelihood and spike-and-slab constructs. The spike-and-slab LASSO methodology has experienced rapid development in recent years, and its scope now extends well beyond the normal linear regression model (1.1). The purpose of this chapter is to offer a timely review of the SSL and its many variants. We first provide a basic review of frequentist and Bayesian approaches to high-dimensional variable selection and estimation under the model (1.1). We then review the spike-and-slab LASSO and provide an overview of its attractive properties and techniques to implement it within the context of normal linear regression. Next, we review the methodological developments of the SSL and some of its theoretical developments.

1.2 Variable selection in high dimensions: Frequentist and Bayesian strategies

1.2.1 Penalized likelihood approaches

In the frequentist high-dimensional literature, there have been a variety of penalized likelihood approaches proposed to estimate β in (1.1). A variant of the penalized likelihood approach estimates β with

$$\hat{\beta} = \arg \max_{\beta \in \mathbb{R}^p} -\frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \text{pen}_\lambda(\beta), \quad (1.2)$$

where $\text{pen}_\lambda(\beta)$ is a penalty function indexed by penalty parameter λ . Most of the literature has focused on penalty functions which are separable, i.e. $\text{pen}_\lambda(\beta) = \sum_{j=1}^p \rho_\lambda(\beta_j)$. In particular, the popular least absolute shrinkage and selection operator (LASSO) penalty of [56] uses the function $\rho_\lambda(\beta_j) = -\lambda|\beta_j|$. Besides the LASSO and its many variants [65, 64, 5, 51, 62], other popular choices for $\rho_\lambda(\cdot)$ include non-concave penalty functions, such as the smoothly clipped absolute deviation (SCAD) penalty [14] and the minimax concave penalty (MCP) [61]. All of the aforementioned penalties threshold some coefficients to zero, thus enabling them to perform variable selection and estimation simultaneously. In addition, SCAD and MCP also mitigate the well-known estimation bias of the LASSO.

Any penalized likelihood estimator (1.2) also has a Bayesian interpretation in that it can be seen as a posterior mode under an independent product prior $p(\beta|\lambda) = \prod_{j=1}^p p(\beta_j|\lambda)$, where $\text{pen}_\lambda(\beta) = \log p(\beta|\lambda) = \sum_{j=1}^p \log p(\beta_j|\lambda)$. In particular, the solution to the LASSO is equivalent to the posterior mode under a product of Laplace densities indexed by hyperparameter, λ :

$$p(\beta|\lambda) = \prod_{j=1}^p \frac{\lambda}{2} e^{-\lambda|\beta_j|}. \quad (1.3)$$

This prior, known as the Bayesian lasso, was first introduced by [39]. In [39], both fully Bayes and empirical Bayes procedures were developed to tune the hyperparameter λ in (1.3). The fully Bayes approach of placing a prior on λ , in particular, renders the Bayesian lasso penalty *non*-separable. Thus, the fully Bayesian lasso has the added advantage of being able to share information across different coordinates. Despite this benefit, [42] showed that the fully Bayesian lasso cannot simultaneously adapt to sparsity *and* avoid the estimation bias issue of the original LASSO. Additionally, [11] showed that the posterior under the Bayesian lasso contracts at a suboptimal rate. In Sections 1.3-1.6, we will illustrate how the *spike-and-slab LASSO* mitigates this problem.

1.2.2 Spike-and-slab priors

In the Bayesian framework, variable selection under the linear model (1.1) arises directly from probabilistic considerations and has frequently been carried out through placing spike-and-slab priors on the coefficients of interest. The spike-and-slab prior was first introduced

by [33] and typically has the following form,

$$\begin{aligned} p(\boldsymbol{\beta}|\boldsymbol{\gamma}, \sigma^2) &= \prod_{j=1}^p [(1 - \gamma_j)\delta_0(\beta_j) + \gamma_j p(\beta_j|\sigma^2)], \\ p(\boldsymbol{\gamma}|\theta) &= \prod_{j=1}^p \theta^{\gamma_j} (1 - \theta)^{1-\gamma_j}, \quad \theta \sim p(\theta), \\ \sigma^2 &\sim p(\sigma^2). \end{aligned} \tag{1.4}$$

where δ_0 is a point mass at zero used to model the negligible entries (the “spike”), $p(\beta_j|\sigma^2)$ is a diffuse and/or heavy-tailed density (rescaled by the variance σ^2) to model the non-negligible entries (the “slab”), $\boldsymbol{\gamma}$ is a binary vector that indexes the 2^p possible models, and $\theta \in (0, 1)$ is a mixing proportion. The error variance σ^2 is typically endowed with a conjugate inverse gamma prior or an improper Jeffreys prior, $p(\sigma^2) \propto \sigma^{-2}$. With a well-chosen prior on θ , this prior (1.4) also automatically favors parsimonious models in high dimensions, thus avoiding the curse of dimensionality.

The point-mass spike-and-slab prior (1.4) is often considered “theoretically ideal,” or a “gold standard” for sparse Bayesian problems [45, 40, 9]. In high dimensions, however, exploring the full posterior using point-mass spike-and-slab priors (1.4) can be computationally prohibitive, in large part because of the combinatorial complexity of updating the discrete indicators $\boldsymbol{\gamma}$. There has been some work to mitigate this issue by either using either shotgun stochastic search (SSS) [20, 8] or variational inference (VI) [41] to quickly identify regions of high posterior probability.

As an alternative to the point-mass spike-and-slab prior, fully continuous spike-and-slab models have been developed. In these continuous variants, the point-mass δ_0 in (1.4) is replaced by a continuous density that is heavily concentrated about zero. The first such continuous relaxation was made by [17], who used a normal density with very small variance for the spike and a normal density with very large variance for the slab. Specifically, the prior for $\boldsymbol{\beta}$ in [17] is

$$p(\boldsymbol{\beta}|\boldsymbol{\gamma}, \sigma^2) = \prod_{j=1}^p [(1 - \gamma_j)\mathcal{N}(0, \sigma^2\nu_0) + \gamma_j\mathcal{N}(0, \sigma^2\nu_1)], \tag{1.5}$$

where $0 < \nu_0 \ll \nu_1$. [17] developed a stochastic search variable selection (SSVS) procedure based on posterior sampling with Markov chain Monte Carlo (MCMC) and thresholding the posterior inclusion probabilities, $\Pr(\gamma_j = 1|\mathbf{y}), j = 1, \dots, p$. In practice, the “median thresholding” rule [4], i.e. $\Pr(\gamma_j = 1|\mathbf{y}) > 0.5, j = 1, \dots, p$, is often used to perform variable selection. [21, 37] further extended the model (1.5) by rescaling the variances ν_0 and ν_1 with sample size n in order to better control the amount of shrinkage for each individual coefficient.

To further reduce the computational intensiveness of SSVS, a deterministic procedure called EM variable selection (EMVS) was developed by [46]. The EMVS procedure employs (1.5) as the prior for $\boldsymbol{\beta}^1$ and uses an EM algorithm to target the posterior mode for $(\boldsymbol{\beta}, \theta, \sigma)$. Whereas SSS and SSVS may only visit the highest probability model a small number of times or may never even visit it at all, EMVS has been shown to consistently find the model with the highest posterior probability [46]. Recently, [27] proposed a general algorithmic framework for Bayesian variable selection with graph-structured sparsity

¹[46] also consider continuous spike-and-slab models where the slab, $\mathcal{N}(0, \sigma^2\nu_1)$, is replaced with a polynomial-tailed density, such as a Student’s t or a Cauchy distribution, to prevent overshrinkage of the non-negligible entries in $\boldsymbol{\beta}$.

which subsumes the EMVS algorithm as a special case. Like SSVS, these algorithms also require thresholding the posterior inclusion probabilities to perform variable selection. Letting $(\hat{\beta}, \hat{\theta}, \hat{\sigma})$ denote the posterior mode for (β, θ, σ) , [46] recommend using median thresholding, $\Pr(\gamma_j = 1 | \hat{\beta}, \hat{\theta}, \hat{\sigma}) > 0.5$, for selection.

1.3 The spike-and-slab LASSO

Having reviewed the penalized likelihood and spike-and-slab paradigms for sparse modeling in the normal linear regression model (1.1), we are now in a position to review the spike-and-slab LASSO (SSL) of [48]. The SSL forms a bridge between these two parallel developments, thereby combining the strengths of both approaches into a single procedure. Throughout this section and Section 1.4, we assume that \mathbf{Y} has been centered at zero to avoid the need for an intercept and that the design matrix \mathbf{X} has been centered and standardized so that $\|\mathbf{x}_j\|_2^2 = n$ for all $1 \leq j \leq p$.

1.3.1 Prior specification

The spike-and-slab LASSO prior is specified as

$$\begin{aligned} p(\beta|\theta) &= \prod_{j=1}^p [(1 - \gamma_j)\psi(\beta_j|\lambda_0) + \gamma_j\psi(\beta_j|\lambda_1)], \\ p(\gamma|\theta) &= \prod_{j=1}^p [\theta^{\gamma_j}(1 - \theta)^{1-\gamma_j}], \\ \theta &\sim \text{Beta}(a, b), \end{aligned} \tag{1.6}$$

where $\psi(\beta|\lambda) = (\lambda/2)e^{-\lambda|\beta|}$ denotes the Laplace density with scale parameter λ . Figure 1.1 depicts the Laplace density for two different choices of scale parameter. We see that for large λ ($\lambda = 20$), the density is very peaked around zero, while for small λ ($\lambda = 1$), it is diffuse. Therefore, in our prior (1.6), we typically set $\lambda_0 \gg \lambda_1$, so that $\psi(\cdot|\lambda_0)$ is the “spike” and $\psi(\cdot|\lambda_1)$ is the “slab.”


The original SSL model of [48] assumed known variance $\sigma^2 = 1$. [35] extended the SSL to the unknown variance case. As σ^2 is typically unknown, we consider the hierarchical formulation in [35] in this chapter and place an independent Jeffreys prior on σ^2 ,

$$p(\sigma^2) \propto \sigma^{-2}.$$

Note that unlike the mixture of normals (1.5), we do *not* scale the Laplace priors in $p(\beta|\theta)$ by σ^2 . [35] showed that such scaling severely underestimates the variance σ^2 when β is sparse or when $p > n$, thus making the model prone to overfitting.

By choosing $\lambda_1 = \lambda_0$ in (1.6), we obtain the familiar LASSO ℓ_1 penalty. On the other hand, if $\lambda_0 \rightarrow \infty$, we obtain the “theoretically ideal” point-mass spike-and-slab (1.4) as a limiting case. Thus, a feature of the SSL prior is its ability to induce a nonconcave continuum between the penalized likelihood and (point-mass) spike-and-slab constructs.

Since it is a mixture of two Laplace distributions, the SSL prior (1.6) can be seen as a two-group refinement of the LASSO’s ℓ_1 penalty on the coefficients. Thus, the posterior mode for $p(\beta|\mathbf{y})$ under (1.6) is *exactly* sparse and can be used to perform simultaneous variable selection and parameter estimation. This automatic modal thresholding property offers an advantage over previous spike-and-slab formulations (1.4)-(1.5) which do not give

**FIGURE 1.1**

Plot of the central region for the Laplace density with two different choices of scale parameter.

exactly sparse estimates of the coefficients and which typically require *post hoc* thresholding of the posterior inclusion probabilities for selection.

It is well-known that the original LASSO [56] suffers from estimation bias, wherein coefficients with large magnitude are overshrunk. One may wonder what advantages the SSL (1.6) confers over penalized likelihood approaches such as the adaptive LASSO, SCAD, or MCP penalties [64, 14, 61] which are designed to mitigate the bias problem of the LASSO. In what follows, we discuss two major advantages of the SSL. First, we demonstrate that the SSL mixes two LASSO “bias” terms *adaptively* in such a way that either a very large amount of shrinkage is applied if $|\beta_j|$ is small or a very small amount of shrinkage is applied if $|\beta_j|$ is large. This is in contrast to the adaptive LASSO [64] and similar penalties which assign *fixed* coefficient-specific penalties and thus do not gear the coefficient-specific shrinkage towards these extremes. Second, the prior on θ in (1.6) ultimately renders the coordinates in β *dependent* in the marginal prior $p(\beta)$ and the SSL penalty *non-separable*. This provides the SSL with the additional ability to borrow information across coordinates and adapt to ensemble information about sparsity.

1.3.2 Selective shrinkage and self-adaptivity to sparsity

As noted in Section 1.2.1, any sparsity-inducing Bayesian prior can be recast in the penalized likelihood framework by treating the logarithm of the marginal prior $\log p(\beta)$ as a penalty function. The SSL penalty is defined as

$$\text{pen}(\beta) = \log \left[\frac{p(\beta)}{p(\mathbf{0}_p)} \right], \quad (1.7)$$

where the penalty has been centered so that $\text{pen}(\beta) = 0$ [48]. Using (1.7) and some algebra, the log posterior under the SSL prior (up to an additive constant) can be shown to be

$$L(\beta, \sigma^2) = -\frac{1}{2\sigma^2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 - (n+2) \log \sigma + \sum_{j=1}^p \text{pen}(\beta_j | \theta_j), \quad (1.8)$$

where for $j = 1, \dots, p$,

$$\text{pen}(\beta_j|\theta_j) = -\lambda_1|\beta_j| + \log[p_{\theta_j}^*(0)/p_{\theta_j}^*(\beta_j)], \quad (1.9)$$

with

$$p_{\theta}^*(\beta) = \frac{\theta\psi(\beta|\lambda_1)}{\theta\psi(\beta|\lambda_1) + (1-\theta)\psi(\beta|\lambda_0)}, \quad (1.10)$$

and

$$\theta_j = E[\theta|\beta_{\setminus j}] = \int \theta p(\theta|\beta_{\setminus j})d\theta. \quad (1.11)$$

For large p , [48] note that the conditional expectation $\theta_j = E[\theta|\beta_{\setminus j}]$ is very similar to $E[\theta|\beta]$, so for practical purposes, we treat them as equal and denote $\hat{\theta} = E[\theta|\beta]$. The connection between the SSL and penalized likelihood methods is made clearer when considering the derivative of each singleton penalty $\text{pen}(\beta_j|\hat{\theta})$ in (1.9). This derivative corresponds to an implicit bias term [48] and is given by

$$\frac{\partial \text{pen}(\beta_j|\hat{\theta})}{\partial |\beta_j|} = -\lambda_{\hat{\theta}}^*(\beta_j), \quad (1.12)$$

where

$$\lambda_{\hat{\theta}}^*(\beta_j) = \lambda_1 p_{\hat{\theta}}^*(\beta_j) + \lambda_0 [1 - p_{\hat{\theta}}^*(\beta_j)]. \quad (1.13)$$

The Karush-Kuhn-Tucker (KKT) conditions yield the following necessary condition for the global mode $\hat{\beta}$:

$$\hat{\beta}_j = \frac{1}{n} \left[|z_j| - \sigma^2 \lambda_{\hat{\theta}}^*(\hat{\beta}_j) \right]_+ \text{sign}(z_j), \quad j = 1, \dots, p, \quad (1.14)$$

where $z_j = \mathbf{x}_j^T(\mathbf{y} - \sum_{k \neq j} \mathbf{x}_k \hat{\beta}_k)$. Notice that the condition (1.14) resembles the soft-thresholding operator for the LASSO, except that it contains an adaptive penalty term $\lambda_{\hat{\theta}}^*$ for *each* coefficient. In particular, the quantity (1.13) is a weighted average of the two regularization parameters, λ_1 and λ_0 , and the weight $p_{\hat{\theta}}^*(\beta_j)$. Thus, (1.13)-(1.14) show that the SSL penalty induces an *adaptive* regularization parameter which applies a different amount of shrinkage to each coefficient, unlike the original LASSO which applies the same shrinkage to every coefficient.

It is worth looking at the term $p_{\hat{\theta}}^*(\beta_j)$ more closely. In light of (1.10), this quantity can be viewed as a conditional probability that β_j was drawn from the slab distribution rather than the spike distribution, having seen the regression coefficient β_j . We have $p_{\hat{\theta}}^*(\beta_j) = \Pr(\gamma_j = 1|\beta_j, \hat{\theta})$, where

$$p_{\hat{\theta}}^*(\beta_j) = \frac{1}{1 + \frac{(1-\hat{\theta})}{\hat{\theta}} \frac{\lambda_0}{\lambda_1} \exp[-|\beta_j|(\lambda_0 - \lambda_1)]} \quad (1.15)$$

is an *exponentially increasing* function in $|\beta_j|$. From (1.15), we see that the functional $p_{\hat{\theta}}^*$ has a sudden increase from near-zero to near-one. Therefore, $p_{\hat{\theta}}^*(\beta_j)$ gears $\lambda_{\hat{\theta}}^*$ in (1.13) towards the extreme values λ_1 and λ_0 , depending on the size of $|\beta_j|$. Assuming that λ_1 is sufficiently small (and hence, the slab $\psi(\beta_j|\lambda_1)$ is sufficiently diffuse), this allows the large coefficients to escape the overall shrinkage effect, in sharp contrast to the single Laplace distribution (1.3), where the bias issue remains even if a prior is placed on λ [42].

Apart from its selective shrinkage property, a second key benefit of the SSL model (1.6) is its *self-adaptivity* to the sparsity pattern of the data through the prior on the mixing proportion θ , $p(\theta) \sim \mathcal{B}\text{eta}(a, b)$. As mentioned previously, this prior ultimately renders the SSL penalty *non-separable*. Fully separable penalty functions, such as those described

in Section 1.2.1, are limited by their inability to adapt to common features across model parameters because they treat these parameters independently. In contrast, treating θ (the expected proportion for non-negligible coefficients in β) as random, allows for automatic adaptivity to different levels of sparsity. As shown in (1.10)-(1.11) (and replacing θ_j with $\hat{\theta}$ and $\beta_{\setminus j}$ with β), the mixing weight p_θ^* is obtained by averaging $p_\theta^*(\cdot)$ over $p(\theta|\beta)$, i.e. $p_\theta^*(\beta) = \int_0^1 p_\theta^*(\beta_j)p(\theta|\beta)d\theta$. It is through this averaging that the SSL penalty (1.7) is given an opportunity to borrow information across coordinates and learn about the underlying level of sparsity in β .

For the hyperparameters (a, b) in the beta prior on θ , [48] recommended the default choice of $a = 1, b = p$. By Lemma 4 of [48], this choice ensures that $E[\theta|\hat{\beta}] \sim \hat{q}/p$, where \hat{q} is the number of nonzero coefficients in $\hat{\beta}$. Further, this choice of hyperparameters results in an automatic multiplicity adjustment [50] and ensures that θ is small (or that most of the coefficients belong to the spike) with high probability. Thus, the SSL also favors parsimonious models in high dimensions and avoids the curse of dimensionality.

1.4 Computational details

We now turn our attention to implementation of the SSL model (1.6) under the normal linear regression model (1.1). The method described in Section 1.4.1 is implemented in the publicly available R package `SSLASSO` [44]. However, we also describe an alternative implementation approach in Section 1.4.3, which is amenable to situations outside of the Gaussian likelihood.

1.4.1 Coordinate-wise optimization

As mentioned in Section 1.3 and shown in (1.14), the (global) posterior mode under the SSL prior (1.6) is exactly sparse, while avoiding the excessive bias issue for large coefficients. Therefore, we can obtain estimates for β by targeting the posterior mode.

Marginalizing out γ in (1.6) gives the prior for β (conditional on θ),

$$p(\beta|\theta) = \prod_{j=1}^p [(1 - \theta)\psi(\beta_j|\lambda_0) + \theta\psi(\beta_j|\lambda_1)]. \quad (1.16)$$

Using this reparametrization, [48, 35] developed a highly efficient coordinate ascent algorithm to quickly target the mode for (β, σ^2) .

Since the SSL is a non-convex method, the KKT conditions only give a necessary condition (1.14) for $\hat{\beta}$ to be a global mode, but not a sufficient one. When $p > n$ and $\lambda_0 \gg \lambda_1$, the posterior will typically be multimodal. Nevertheless, it is still possible to obtain a refined characterization of the global mode. Building upon theory developed by [63], [48, 35] gave necessary *and* sufficient conditions for $\hat{\beta}$ to be a *global* mode. By Theorems 3-4 in [48] and Propositions 4-5 in [35], the global mode under the SSL prior (1.6) is a blend of soft-thresholding *and* hard-thresholding, namely

$$\hat{\beta}_j = \frac{1}{n} \left[|z_j| - \sigma^2 \lambda_\theta^*(\hat{\beta}_j) \right]_+ \text{sign}(z_j) \mathbb{I}(|z_j| > \Delta), \quad (1.17)$$

where $z_j = \mathbf{x}_j^T(\mathbf{y} - \sum_{k \neq j} \mathbf{x}_k \hat{\beta}_k)$ and $\Delta \equiv \inf_{t > 0} [nt/2 - \sigma^2 \text{pen}(t\hat{\theta})/t]$. In [35], an approxi-

mation for Δ is given by

$$\Delta = \begin{cases} \sqrt{2n\sigma^2 \log[1/p_\theta^*(0)] + \sigma^2 \lambda_1}, & \text{if } g_{\hat{\theta}}(0) > 0, \\ \sigma^2 \lambda_\theta^*(0) & \text{otherwise,} \end{cases}$$

where $g_\theta(x) = [\lambda_\theta^*(x) - \lambda_1]^2 + (2n/\sigma^2) \log p_\theta^*(x)$. The generalized thresholding operator (1.17) allows us to eliminate many suboptimal local modes from consideration through the threshold Δ . This refined characterization also facilitates a highly efficient coordinate ascent algorithm [30] to find the global mode, which we now detail.

After initializing $(\Delta^{(0)}, \beta^{(0)}, \theta^{(0)}, \sigma^{2(0)})$, the coordinate ascent algorithm iteratively updates these parameters until convergence. The update for the threshold Δ at the t^{th} iteration is

$$\Delta^{(t)} = \begin{cases} \sqrt{2n\sigma^{2(t-1)} \log[1/p_{\hat{\theta}^{(t-1)}}^*(0)] + \sigma^{2(t-1)} \lambda_1} & \text{if } g_{\hat{\theta}^{(t-1)}}(0) > 0, \\ \sigma^{2(t-1)} \lambda_{\hat{\theta}^{(t-1)}}^*(0) & \text{otherwise,} \end{cases}$$

Next, β is updated as

$$\beta_j^{(t)} \leftarrow \frac{1}{n} \left(|z_j| - \lambda_{\hat{\theta}^{(t-1)}}^*(\hat{\beta}_j^{(t-1)}) \right)_+ \text{sign}(z_j) \mathbb{I}(|z_j| > \Delta^{(t)}).$$

Using the approximation for $E[\theta|\hat{\beta}]$ in Lemma 4 of [48], the update for $\hat{\theta}$ is

$$\hat{\theta}^{(t)} \leftarrow \frac{a + \hat{q}^{(t)}}{a + b + p},$$

where $\hat{q}^{(t)}$ is the number of nonzero entries in $\beta^{(t)}$. Finally, the update for σ^2 is

$$\sigma^{2(t)} \leftarrow \frac{\|\mathbf{y} - \mathbf{X}\beta^{(t)}\|_2^2}{n + 2}.$$

This coordinate ascent algorithm specifically appeals to the theoretical framework of [63] to find the global SSL mode $\hat{\beta}$. An alternative way of implementing the SSL model is to use an expectation-maximization (EM) algorithm, in the vein of EMVS [46], which treats the indicators γ in (1.6) as missing data. This EM approach is described in Section 1.4.3 and can be easily modified for other statistical models (such as the methods described in Section 1.7) where the likelihood function differs and the theory of [63] is not applicable.

1.4.2 Dynamic posterior exploration

The performance of the SSL model depends on good choices for the hyperparameters (λ_0, λ_1) in (1.6). To this end, [48] recommend a “dynamic posterior exploration” strategy in which the slab hyperparameter λ_1 is held fixed at a small value and the spike hyperparameter λ_0 is gradually increased along a ladder of increasing values, $\{\lambda_0^1, \dots, \lambda_0^L\}$. The algorithm is not very sensitive to the specific choice of λ_1 , provided that the slab is sufficiently diffuse. For each λ_0^s in the ladder for the spike parameters, we reinitialize $(\Delta^{(0)}, \beta^{(0)}, \theta^{(0)}, \sigma^{2(0)})$ using the MAP estimates for these parameters from the previous spike parameter λ_0^{s-1} as a “warm start.”

This sequential reinitialization strategy allows the SSL to more easily find the global mode. In particular, when $(\lambda_1 - \lambda_0)^2 < 4$ and σ^2 is fixed, the objective (1.8) is convex. The intuition here is to use the solution to the convex problem as a “warm” start for the non-convex problem (when $\lambda_0 \gg \lambda_1$). As we increase λ_0 , the posterior becomes “spikier,”

with the spikes absorbing more and more of the negligible parameters. Meanwhile, keeping λ_1 fixed at a small value allows the larger coefficients to escape the pull of the spike. For large enough λ_0 , the algorithm will eventually stabilize so that further increases in λ_0 do not change the solution. In Section 1.6, we illustrate this with plots of the SSL solution paths.

Additionally, as noted by [35], some care must also be taken when updating σ^2 . When $p > n$ and $\lambda_0 \approx \lambda_1$, the model can become saturated, causing the residual variance to go to zero. To avoid this suboptimal mode at $\sigma^2 = 0$, [35] recommend fixing σ^2 until the λ_0 value in the ladder at which the algorithm starts to converge in less than 100 iterations. Then, β and σ^2 are simultaneously updated for the next largest λ_0 in the sequence. The complete algorithm for coordinate-wise optimization with dynamic posterior exploration is given in Section 4 of the supplementary material in [35]. This algorithm is implemented in the R package `SSLASSO`.

1.4.3 EM implementation of the spike-and-slab LASSO

The coordinate ascent algorithm of Section 1.4.1 specifically appeals to the theoretical framework of [63] to find the global SSL mode $\hat{\beta}$. An alternative approach is to use an EM algorithm in the vein of EMVS [46]. Under this framework, the latent variables γ in (1.6) are treated as “missing” data. In the E-step at the t th iteration, we compute $E[\tau_j | \mathbf{y}, \beta^{(t-1)}, \theta^{(t-1)}, \sigma^{2(t-1)}] = p_{\theta^{(t-1)}}^*(\beta_j^{(t-1)})$, where p_{θ}^* is as in (1.10). The M-step then iterates through the following updates:

$$\begin{aligned} \beta^{(t)} &\leftarrow \arg \max_{\beta \in \mathbb{R}^p} \left\{ -\frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 - \sum_{j=1}^p \sigma^{2(t-1)} \lambda_{\theta^{(t-1)}}^*(\beta_j^{(t-1)}) |\beta_j| \right\}, \\ \theta^{(t)} &\leftarrow \frac{\sum_{j=1}^p p_{\theta^{(t-1)}}^*(\beta_j^{(t)}) + a - 1}{a + b + p - 2}, \\ \sigma^{2(t)} &\leftarrow \frac{\|\mathbf{y} - \mathbf{X}\beta^{(t)}\|_2^2}{n + 2}, \end{aligned}$$

where $\lambda_{\theta}^*(\beta) = \lambda_1 p_{\theta}^*(\beta) + \lambda_0 [1 - p_{\theta}^*(\beta)]$. Note that the update for $\beta^{(t)}$ is an adaptive LASSO regression with weights $\sigma^2 \lambda_{\theta}^*$ and hence can be solved very efficiently using gradient descent algorithms [15]. Like EMVS [46], the dynamic posterior exploration strategy detailed in Section 1.4.2 can be used to find a more optimal mode for (β, σ^2) .

The EM approach detailed here can be easily modified for other statistical models where the SSL prior (1.6) is used but where the theory of [63] is not applicable. Similar to the coordinate ascent algorithm described in Section 1.4.1, this EM algorithm may be sensitive to the initialization of $(\beta^{(0)}, \theta^{(0)}, \sigma^{2(0)})$. The dynamic posterior exploration strategy described earlier can partly help to mitigate this issue, since the posterior starts out relatively flat when $\lambda_0 \approx \lambda_1$ but becomes “spikier” as λ_0 increases. By the time that the spikes have reappeared, the “warm start” solution from the previous λ_0 in the ladder should hopefully be in the basin of dominant mode. Other strategies such as running the algorithm for a wide choice of starting values or deterministic annealing can also aid in adding robustness against poor initializations [31, 57, 46].

1.5 Uncertainty quantification

While the algorithms described in Section 1.4 can be used to rapidly target the modes of the SSL posterior, providing a measure of uncertainty for our estimates is a challenging task. In this section, we outline two possible strategies for the task of uncertainty quantification. The first is based on debiasing the posterior mode. The second involves posterior simulation.

1.5.1 Debiasing the posterior mode

One possible avenue for uncertainty quantification is to use debiasing [62, 58, 22, 3]. Let $\widehat{\Sigma} = \mathbf{X}^T \mathbf{X}/n$ and let $\widehat{\Theta}$ be an approximate inverse of $\widehat{\Sigma}$ (obtainable as an estimate for $\widehat{\Theta}$, for example, using the nodewise regression procedure in [32]). Define the quantity $\widehat{\beta}_d$ as

$$\widehat{\beta}_d = \widehat{\beta} + \widehat{\Theta} \mathbf{X}^T (\mathbf{y} - \mathbf{X} \widehat{\beta})/n. \quad (1.18)$$

where $\widehat{\beta}$ is the MAP estimator of β under the SSL model. By [58], this quantity $\widehat{\beta}_d$ has the following asymptotic distribution:

$$\sqrt{n}(\widehat{\beta}_d - \beta) \sim \mathcal{N}(\mathbf{0}, \sigma^2 \widehat{\Theta} \widehat{\Sigma} \widehat{\Theta}^T). \quad (1.19)$$

For inference, we replace the population variance σ^2 in (1.19) with the modal estimate $\widehat{\sigma}^2$ from the SSL model. Let $\widehat{\beta}_{dj}$ denote the j th coordinate of $\widehat{\beta}_d$. We have from (1.19) that the $100(1 - \alpha)\%$ asymptotic pointwise confidence intervals for $\beta_j, j = 1, \dots, p$, are

$$[\widehat{\beta}_{dj} - c(\alpha, n, \widehat{\sigma}^2), \widehat{\beta}_{dj} + c(\alpha, n, \widehat{\sigma}^2)], \quad (1.20)$$

where $c(\alpha, n, \widehat{\sigma}^2) := \Phi^{-1}(1 - \alpha/2) \sqrt{\widehat{\sigma}^2 (\widehat{\Theta} \widehat{\Sigma} \widehat{\Theta}^T)_{jj}/n}$ and $\Phi(\cdot)$ denotes the cumulative distribution function of $\mathcal{N}(0, 1)$.

Note that the posterior modal estimate $\widehat{\beta}$ under the SSL prior already has much less bias than the LASSO estimator [56]. Therefore, the purpose of the debiasing procedure above is mainly to obtain an estimator with an asymptotically normal distribution from which we can construct asymptotic pointwise confidence intervals. While this procedure is asymptotically valid, [1] showed through numerical studies that constructing confidence intervals based on asymptotic arguments may provide coverage below the nominal level in finite samples, especially small samples. Therefore, it may be more ideal to use the actual SSL posterior $p(\beta|\mathbf{y})$ for inference.

1.5.2 Posterior sampling for the spike-and-slab LASSO

Fully Bayesian inference with the SSL can be carried out via posterior simulation. However, posterior sampling under spike-and-slab priors has continued to pose challenges. One immediate strategy for sampling from the SSL posterior is the SSVS algorithm of [17], described in Section 1.2.2. One can regard the Laplace distribution as a scale mixture of Gaussians with an exponential mixing distribution [39] and perform a variant of SSVS. Recently, several clever computational tricks have been suggested that avoid costly matrix inversions needed by SSVS by using linear solvers [6], low-rank approximations [24], or by disregarding correlations between active and inactive coefficients [36]. These techniques can be suitably adapted for fast posterior sampling of the SSL as well.

Intrigued by the speed of SSL mode detection, [38] explored the possibility of turning

SSL into approximate posterior sampling by performing MAP optimization on many independently perturbed datasets. Building on Bayesian bootstrap ideas, they introduced a method for approximate sampling called Bayesian bootstrap spike-and-slab LASSO (BB-SSL) which scales linearly with both n and p . Beyond its scalability, they show that BB-SSL has strong theoretical support, matching the convergence rate of the original posterior in sparse normal-means and in high-dimensional regression.

1.6 Illustrations

1.6.1 Example on synthetic data

To illustrate the SSL in action, we used the simulation setting in [35]. We set $n = 100$ and $p = 1000$ in (1.1). The design matrix \mathbf{X} was generated from a multivariate Gaussian distribution with mean $\mathbf{0}_p$ and a block-diagonal covariance matrix $\mathbf{\Sigma} = \text{bdiag}(\tilde{\Sigma}, \dots, \tilde{\Sigma})$, where $\tilde{\Sigma} = \{\tilde{\sigma}\}_{i,j=1}^{50}$, with $\tilde{\sigma}_{ij} = 0.9$ if $i \neq j$ and $\tilde{\sigma}_{ii} = 1$. The true vector of regression coefficients β_0 was constructed by assigning regression coefficients $\{-2.5, -2, -1.5, 1.5, 2, 2.5\}$ to $q = 6$ entries located at the indices $\{1, 51, 101, 151, 201, 251\}$ and setting the remaining coefficients equal to zero. Hence, there were 20 independent blocks of 50 highly correlated predictors, where the first six blocks contained only one active predictor. We then generated the response \mathbf{y} using (1.1), where the error variance was set as $\sigma^2 = 3$.

We compared the SSL with the LASSO [56], SCAD [14], and MCP [61]. The SSL method was implemented using the R package `SSLASSO`. The competing methods were implemented using the R package `ncvreg`. We repeated our experiment 500 times with new covariates and responses generated each time. For each experiment, we recorded the mean prediction error (MPE), defined as

$$\text{MPE} = \frac{1}{n} \|\mathbf{X}\hat{\beta} - \mathbf{X}\beta_0\|_2^2.$$

We also kept track of \hat{q} , or the size of the model selected by each of these methods. Finally, we recorded the false discovery rate (FDR), the false negative rate (FNR), and the Matthew's correlation coefficient (MCC), defined respectively as

$$\begin{aligned} \text{FDR} &= \frac{\text{FP}}{\text{TP} + \text{FP}}, & \text{FNR} &= \frac{\text{FN}}{\text{TN} + \text{FN}}, \\ \text{MCC} &= \frac{\text{TP} \times \text{TN} - \text{FP} \times \text{FN}}{\sqrt{(\text{TP} + \text{FP})(\text{TP} + \text{FN})(\text{TN} + \text{FP})(\text{TN} + \text{FN})}}, \end{aligned}$$

where TP, TN, FP, and FN denote the number of true positives, true negatives, false positives, and false negatives respectively. R code to reproduce these experiments is available in the online supplementary material.

Table 1.1 reports our results averaged across the 500 replications. We see that the SSL had the lowest average mean prediction error, in addition to selecting (on average) the most parsimonious model. The LASSO had the second lowest MPE, but it tended to select far more variables than the other methods, leading to the highest FDR. In contrast, SSL had the lowest FDR and the highest MCC, indicating that the SSL had the best overall variable selection performance of all the methods. Our simulation study demonstrates that SSL achieves both parsimony *and* accuracy of estimation and selection.

Figure 1.2 illustrates the benefits of the dynamic posterior exploration approach outlined in Section 1.4.2. Specifically, Figure 1.2 plots the solution paths for the regression coefficients

	MPE	\hat{q}	FDR	FNR	MCC
SSL	0.989	6.03	0.288	0.0017	0.712
LASSO	1.14	34.98	0.841	0.0006	0.370
SCAD	1.35	12.47	0.653	0.0023	0.454
MCP	1.93	10.84	0.760	0.0035	0.308

TABLE 1.1

MPE, estimated model size, FDR, FNR, and MCC for SSL, LASSO, SCAD, and MCP. The results are averaged across 500 replications.

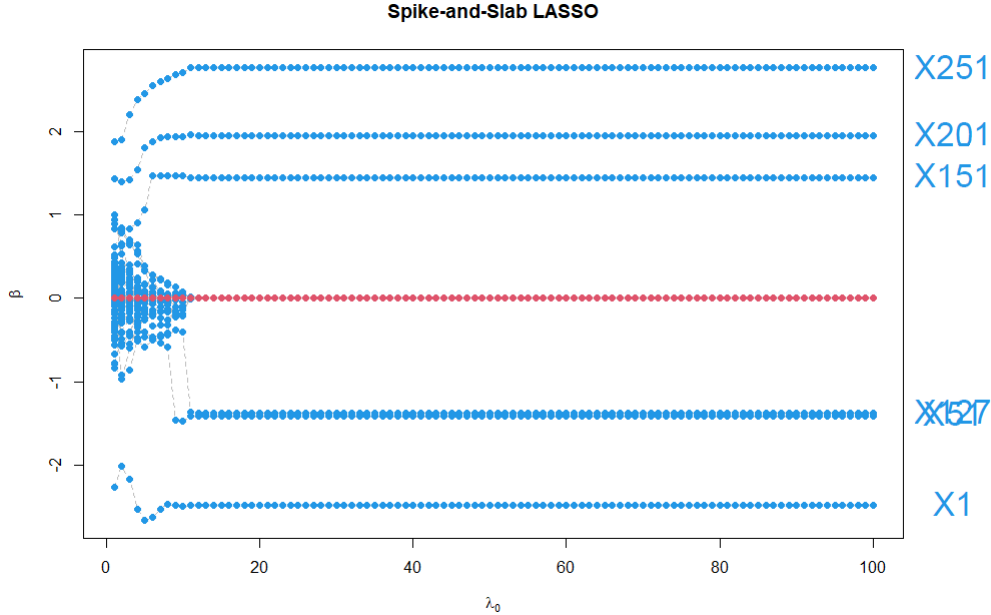


FIGURE 1.2

The solution paths for the SSL along the ladder of spike parameters λ_0 . We see that the solution stabilizes after a certain point, so that further increases in λ_0 do not change the solution.

from one of our experiments as the spike hyperparameter λ_0 increases. We see that the SSL solution stabilizes fairly quickly (when λ_0 is less than 20), so that further increases in λ_0 do not change the solution. This demonstrates that dynamic posterior exploration offers a viable alternative to cross-validation. The R package `SSLASSO` provides the functionality to generate plots of these solution paths.

1.6.2 Bardet-Beidl syndrome gene expression study

We now analyze a microarray data set consisting of gene expression measurements from the eye tissue of 120 laboratory rats². The data was originally studied by [49] to investigate mammalian eye disease. In this data, the goal is to identify genes which are associated with

²Data accessed from the Gene Expression Omnibus www.ncbi.nlm.nih.gov/geo (accession no. GSE5680).

	MPE	Number of selected probe sets
SSL	0.011	28
LASSO	0.012	32
SCAD	0.015	44
MCP	3.699	9

TABLE 1.2

Average MPE and the number of selected probe sets for the Bardet-Beidl Syndrome data analysis.

the gene TRIM32. TRIM32 has previously been shown to cause Bardet-Biedl syndrome [12], a disease affecting multiple organs including the retina.

The original data consists of 31,099 probe sets. For our analysis, we included only the 10,000 probe sets with the largest variances in expression (on the log scale). This resulted in $n = 120$ and $p = 10,000$. We then fit the model (1.1) with an SSL penalty. We compared the SSL approach to LASSO, SCAD, and MCP. To assess predictive accuracy, we randomly split the data set into 90 training observations and 30 test observations. We then fit the models on the training set and used the estimated $\hat{\beta}$ to compute the MPE on the left-out test set. We repeated this process 100 times and took the average MPE.

Table 1.2 shows the results for our analysis, as well as the number of selected probe sets when we fit the different models to the complete data set. SSL had the lowest MPE, indicating the highest predictive power. MCP selected the most parsimonious model, with only nine probe sets out of the 10,000 selected. However, MCP also had a much higher MPE than the other methods. On the other hand, SSL selected 28 probe sets (compared to 32 and 44 for LASSO and SCAD respectively) and still achieved the lowest MPE. Our analysis illustrates that on this particular data set, SSL achieved both the best predictive performance and parsimony.

Of the 28 probe sets selected by SSL as being significantly associated with TRIM32, 14 of them had identifiable gene symbols. These genes were SCGB1A1, CELF1, ASXL3, FGFR2, MOBP, TGM7, SLC39A6, DDX58, TFF2, CLOCK, DUS4L, HTR5B, BIK, and SLC16A6. In particular, according to <https://www.genecards.org>³, SCGB1A1 is known to be an interacting protein for the TRIM32 gene. The other associations that we found may be useful for researchers in studying the genetic factors contributing to Bardet-Biedl syndrome.

1.7 Methodological extensions

While we have focused on the normal linear regression model (1.1) in Sections 1.2-1.6, the spike-and-slab LASSO methodology has now been adopted in a variety of other statistical applications. In this section, we survey some of the extensions of the SSL to models beyond the normal linear regression framework.

Generalized linear models (GLMs). GLMs allow for a flexible generalization of the normal linear regression model (1.1) which can accommodate categorical and count data, in addition to continuous variables. Letting $\mathbf{x}_i \in \mathbb{R}^p$ denote a vector of covariates for the i th observation, GLMs assume that the mean of the response variable is related to the linear

³Accessed from <https://www.genecards.org/cgi-bin/carddisp.pl?gene=TRIM32> on October 11, 2020.

predictor via a link function,

$$E(y_i|\mathbf{x}_i) = h^{-1}(\mathbf{x}_i^T \boldsymbol{\beta}), \quad (1.21)$$

and that the data distribution is expressed as

$$p(\mathbf{y}|\mathbf{X}, \boldsymbol{\beta}, \varphi) = \prod_{i=1}^n p(y_i|\mathbf{x}_i, \boldsymbol{\beta}, \varphi), \quad (1.22)$$

where φ is a dispersion parameter and the distribution $p(y_i|\mathbf{x}_i, \boldsymbol{\beta}, \varphi)$ can take various forms, including normal, binomial and Poisson distributions. Obviously, the normal linear model (1.1) is a special case of (1.21) with the identity link function $h(u) = u$.

[54] extended the SSL (1.3) to GLMs, including binary regression and Poisson regression, by placing the SSL prior (1.6) on the coefficients vector $\boldsymbol{\beta} \in \mathbb{R}^p$ in (1.21) and developing an EM algorithm to perform MAP estimation for $\boldsymbol{\beta}$. For inference with grouped variables in GLMs, [53] further employed group-specific sparsity parameters θ_g for each group of variables, instead of a single θ , as in (1.3).

Survival analysis. The SSL has also proven to be useful for predicting censored survival outcomes and detecting and estimating the effects of relevant covariates. Cox proportional hazards models are the most widely used method for studying the relationship between a censored survival response and an explanatory variable $\mathbf{x}_i \in \mathbb{R}^p$ [28]. This model assumes that the hazard function of survival time t takes the form,

$$h(t|\mathbf{x}_i) = h_0(t) \exp(\mathbf{x}_i^T \boldsymbol{\beta}). \quad (1.23)$$

[55] introduced the spike-and-slab LASSO Cox model which endows the coefficients $\boldsymbol{\beta}$ in (1.23) with the SSL prior (1.6). They developed an EM coordinate ascent algorithm to fit SSL Cox models. [52] further introduced the GssLASSO Cox model which incorporates grouping information by endowing each group of coefficients with a group-specific sparsity parameter θ_g instead of the single θ of (1.6).

Sparse factor analysis and biclustering. Factor models aim to explain the dependence structure among high-dimensional observations through a sparse decomposition of a $p \times p$ covariance matrix $\boldsymbol{\Omega}$ as $\mathbf{B}\mathbf{B}^T + \boldsymbol{\Sigma}$ where \mathbf{B} is a $p \times K$ factor loadings matrix with $K \ll p$ and $\boldsymbol{\Sigma} = \text{diag}(\sigma_1^2, \dots, \sigma_p^2)$. A generic latent factor model is

$$\mathbf{y}_i = \mathbf{B}\boldsymbol{\eta}_i + \boldsymbol{\varepsilon}_i, \quad \boldsymbol{\varepsilon}_i \sim \mathcal{N}_p(\mathbf{0}, \boldsymbol{\Sigma}), \quad (1.24)$$

where \mathbf{y}_i is a p -dimensional continuous response and $\boldsymbol{\eta}_i \sim \mathcal{N}_K(\mathbf{0}, \mathbf{I}_K)$ are unobserved latent factors. Many existing factor analysis approaches entail prespecification of the unknown factor cardinality K and *post hoc* rotations of the original solution to sparsity. For the factor model (1.24), [47] endowed the entries of the loading matrix \mathbf{B} with independent SSL priors,

$$p(\beta_{jk}|\gamma_{jk}, \lambda_{0k}, \lambda_1) = (1 - \gamma_{jk})\psi(\beta_{jk}|\lambda_{0k}) + \gamma_{jk}\psi(\beta_{jk}|\lambda_1).$$

However, instead of endowing each of the indicators γ_{jk} with the usual beta-Bernoulli prior as in (1.6), [47] endowed these with the Indian buffet process (IBP) prior [19], which avoids the need to prespecify K . Further, [47] developed a parameter-expanded EM (PXL-EM) algorithm which employs *intermediate* orthogonal rotations rather than *post hoc* rotations. In addition to obtaining a sparse solution, the PXL-EM algorithm also converges much faster than the vanilla EM algorithm and offers robustness against poor initializations.

For the problem of biclustering, i.e. identifying clusters using only subsets of their associated features, [34] utilized the factor model (1.24) in which *both* the factors $\boldsymbol{\eta} =$

$[\boldsymbol{\eta}_1^T, \dots, \boldsymbol{\eta}_n^T] \in \mathbb{R}^{n \times K}$ and the loadings are sparse. To achieve a doubly sparse representation, [34] placed an SSL prior coupled with an IBP prior on the factors and an SSL prior coupled with a beta-Bernoulli prior on the loadings. An EM algorithm with a variational step was developed to implement spike-and-slab LASSO biclustering.

Graphical models. Suppose we are given data $\mathbf{Y} = (\mathbf{y}_1, \dots, \mathbf{y}_n)^T$, where the \mathbf{y}_i 's are assumed to be iid p -variate random vectors distributed as $\mathcal{N}_p(\mathbf{0}, \boldsymbol{\Omega}^{-1})$ and $p > n$. In this setting, off-diagonal zero entries ω_{ij} encode conditional independence between variables i and j . To obtain a sparse estimate of $\boldsymbol{\Omega} = (\omega_{i,j})_{i,j}$, [16] introduced the following prior on $\boldsymbol{\Omega}$:

$$p(\boldsymbol{\Omega}) = \prod_{i < j} [(1 - \theta)\psi(\omega_{ij}|\lambda_0) + \theta\psi(\omega_{ij}|\lambda_1)] \prod_{i=1}^p [\tau e^{-\tau\omega_{ii}}] \mathbb{I}(\boldsymbol{\Omega} \succ 0) \mathbb{I}(\|\boldsymbol{\Omega}\|_2 \leq B), \quad (1.25)$$

for some $\tau > 0, B > 0$. Here, $\boldsymbol{\Omega} \succ 0$ denotes that $\boldsymbol{\Omega}$ is positive-definite and $\|\boldsymbol{\Omega}\|_2$ denotes the spectral norm of $\boldsymbol{\Omega}$. The prior on $\boldsymbol{\Omega}$ (1.25) entails independent exponential priors on the diagonal entries and SSL priors on the off-diagonal entries. A similar prior formulation was considered in [13], except [13] did not constrain $\boldsymbol{\Omega}$ to lie in the space of $p \times p$ matrices with uniformly bounded spectral norm. [16] showed that constraining the parameter space for $\boldsymbol{\Omega}$ in such a way ensures that a) the corresponding optimization problem for the posterior mode is *strictly* convex, and b) the posterior mode is a symmetric positive definite matrix. [16] developed an EM algorithm to estimate the posterior mode of $p(\boldsymbol{\Omega}|\mathbf{Y})$.

The SSL prior was also extended to perform *joint* estimation of multiple related Gaussian graphical models by [29]. [29] leveraged similarities in the underlying sparse precision matrices and developed an EM algorithm to perform this joint estimation.

Seemingly unrelated regression models. In seemingly unrelated regression models, multiple correlated responses are regressed on multiple predictors. The multivariate linear regression is an important case. Letting $\mathbf{y}_i \in \mathbb{R}^q$ be the vector of q responses and $\mathbf{x}_i \in \mathbb{R}^p$ be the vector of p covariates, this model is

$$\mathbf{y}_i = \mathbf{x}_i^T \mathbf{B} + \varepsilon_i, \quad \varepsilon_i \sim \mathcal{N}_q(\mathbf{0}, \boldsymbol{\Omega}^{-1}), \quad (1.26)$$

[13] introduced the *multivariate spike-and-slab LASSO* (mSSL) to perform joint selection and estimation from the $p \times q$ matrix of regressors \mathbf{B} and the precision matrix $\boldsymbol{\Omega}$. To obtain a sparse estimate of $(\mathbf{B}, \boldsymbol{\Omega})$, [13] placed the SSL prior (1.6) on the individual entries $\beta_{jk}, 1 \leq j \leq p, 1 \leq k \leq q$ in \mathbf{B} and a product prior similar to (1.25) on $\boldsymbol{\Omega}$ (except [13] did not constrain $\boldsymbol{\Omega}$ to have bounded spectral norm). An expectation/conditional maximization (ECM) algorithm was developed to perform this joint estimation.

Causal inference. In observational studies, we are often interested in estimating the causal effect of a treatment T on an outcome y , which requires proper adjustment of a set of potential confounders $\mathbf{x} \in \mathbb{R}^p$. When $p > n$, direct control for all potential confounders is infeasible and standard methods such as propensity scoring [43] often fail. In this case, it is crucial to impose a low-dimensional structure on the confounder space. Given data $(y_i, T_i, \mathbf{x}_i), i = 1, \dots, n$, where T_i is the treatment effect, [1] estimated the (homogeneous) average treatment effect (ATE) $\Delta(t_1, t_2) = E(Y(t_1) - Y(t_2))$ by utilizing the model,

$$y_i | T_i, \mathbf{x}_i, \beta_0, \beta_t, \boldsymbol{\beta}, \sigma^2 \sim \mathcal{N}(0, \beta_0 + \beta_t T_i + \mathbf{x}_i^T \boldsymbol{\beta}, \sigma^2). \quad (1.27)$$

Under (1.27), the ATE is straightforwardly estimated as $\Delta(t_1, t_2) = (t_1 - t_2)\widehat{\beta}_t$. [1] endowed the coefficients of the confounders $\boldsymbol{\beta}$ with the SSL prior (1.6). In addition, [1] also weighted the sparsity parameter θ in (1.6) by raising θ to a power $w_j, j = 1, \dots, p$, for each covariate, in order to better prioritize variables belonging to the slab (i.e. $\gamma_j = 1$) if they are also

associated with the treatment. [1] further extended the model (1.27) to the more general case of heterogeneous treatment effects.

Regression with grouped variables. Group structure arises in many statistical applications. For example, in genomics, genes within the same pathway may form a group and act in tandem to regulate a biological system. For regression with grouped variables, we can model the response \mathbf{y} as

$$\mathbf{y} = \sum_{g=1}^G \mathbf{X}_g \boldsymbol{\beta}_g + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim \mathcal{N}_n(\mathbf{0}, \sigma^2 \mathbf{I}_n), \quad (1.28)$$

where $\boldsymbol{\beta}_g \in \mathbb{R}^{m_g}$ is a coefficients *vector* of length m_g , and \mathbf{X}_g is an $n \times m_g$ covariate matrix corresponding to group $g = 1, \dots, G$. Under model (1.28), it is often of practical interest to select non-negligible groups and estimate their effects. To this end, [3] introduced the *spike-and-slab group lasso* (SSGL). To regularize groups of coefficients, the SSGL replaces the univariate Laplace densities in the univariate SSL (1.6) with *multivariate* Laplace densities. The SSGL prior is

$$p(\boldsymbol{\beta}|\theta) = \prod_{g=1}^G [(1-\theta)\Psi(\boldsymbol{\beta}_g|\lambda_0) + \theta\Psi(\boldsymbol{\beta}_g|\lambda_1)], \quad (1.29)$$

$$\theta \sim \text{Beta}(a, b),$$

where $\Psi(\boldsymbol{\beta}_g|\lambda) \propto \lambda^{m_g} e^{-\lambda\|\boldsymbol{\beta}_g\|_2}$ and $\lambda_0 \gg \lambda_1$. The SSGL (1.29) is a two-group refinement of an ℓ_2 penalty on groups of coefficients. Accordingly, the posterior mode under the SSGL thresholds entire groups of coefficients to zero, while simultaneously estimating the effects of nonzero groups and circumventing the estimation bias of the original group lasso [60]. [3] developed an efficient blockwise-coordinate ascent algorithm to implement the SSGL model.

Nonparameteric additive regression. The advent of the SSGL prior (1.29) paved the way for the spike-and-slab lasso methodology to be extended to nonparametric problems. [3] introduced the *nonparametric spike-and-slab lasso* (NPSSL) for sparse generalized additive models (GAMs). Under this model, the response surface is decomposed into the sum of univariate functions,

$$y_i = \sum_{j=1}^p f_j(x_{ij}) + \varepsilon_i, \quad \varepsilon_i \stackrel{iid}{\sim} \mathcal{N}(0, \sigma^2). \quad (1.30)$$

In [3], each of the f_j 's is approximated using a basis expansion, or a linear combination of basis functions $\mathcal{B}_j = \{g_{j1}, \dots, g_{jd}\}$, i.e.

$$f_j(x_{ij}) \approx \sum_{k=1}^d g_{jk}(x_{ij})\beta_{jk}. \quad (1.31)$$

Under sparsity, most of the f_j 's in (1.30) are assumed to be $f_j = 0$. This is equivalent to assuming that most of the weight vectors $\boldsymbol{\beta}_j = (\beta_{j1}, \dots, \beta_{jd})^T$, $j = 1, \dots, p$, in (1.31) are equal to $\mathbf{0}_d$. The NPSSL is implemented by endowing the basis coefficients $\boldsymbol{\beta} = (\boldsymbol{\beta}_1^T, \dots, \boldsymbol{\beta}_p^T)^T$ with the SSGL prior (1.29) to simultaneously select and estimate nonzero functionals. In addition, [3] also extended the NPSSL to identify and estimate the effects of nonlinear interaction terms $f_{rs}(X_{ir}, X_{is})$, $r \neq s$.

Functional regression. The spike-and-slab lasso methodology has also been extended to functional regression, where the response $y(t)$ is a function that *varies* over some continuum T (often time) A very popular model in this framework is the nonparametric varying

coefficient model,

$$y_i(t) = \sum_{k=1}^p x_{ik}(t)\beta_k(t) + \varepsilon_i(t), \quad t \in T, \quad (1.32)$$

where $y_i(t)$ and $x_{ik}(t)$ are time-varying responses and covariates respectively and $\varepsilon_i(t)$ is a zero-mean stochastic process which captures the within-subject temporal correlations for the i th subject. Under (1.32), the $\beta_k(t)$'s are smooth functions of time (possibly $\beta_k(t) = 0$ for all $t \in T$), and our primary interest is in estimation and variable selection from the $\beta_k(t)$'s, $k = 1, \dots, p$.

[2] introduced the nonparametric varying coefficient spike-and-slab lasso (NVC-SSL) to simultaneously select and estimate the smooth functions $\beta_k(t)$, $k = 1, \dots, p$. Similarly as with the NPSSL, these functions are approximated using basis expansions of smoothing splines, and the basis coefficients are endowed with the SSGL prior (1.29). Unlike GAMs, however, the NVC-SSL model does *not* assume homoscedastic, independent error terms. Instead, the NVC-SSL model accounts for within-subject temporal correlations in its estimation procedure.

False discovery rate control with missing data. Sorted L-One Penalized Estimator (SLOPE) is an elaboration of the LASSO tailored to false discovery control by assigning more penalty to the larger coefficients [7]. SSL, on the other hand, penalizes large coefficients less and its false discovery rate is ultimately determined by a combination of the prior inclusion weight θ and penalties λ_1 and λ_0 . [23] propose a hybrid procedure called adaptive Bayesian SLOPE, which effectively combines the SLOPE method (sorted l_1 regularization) together with the SSL method in the context of variable selection with missing covariate values. As with SSL, the coefficients are regarded as arising from a hierarchical model consisting of two groups: (1) the spike for the inactive and (2) the slab for the active. However, instead of assigning spike priors for each covariate, they propose a joint ‘‘SLOPE’’ spike prior which takes into account ordering of coefficient magnitudes in order to control for false discoveries.

1.8 Theoretical properties

In addition to its computational tractability and its excellent finite-sample performance, the spike-and-slab LASSO has also been shown to provide strong theoretical guarantees. Although this chapter focuses mainly on methodology, we briefly outline a few of the major theoretical developments for spike-and-slab LASSO methods.

A common theme in Bayesian asymptotic theory is the study of the learning rate of posterior point estimates (such as the mean, median or mode) and/or of the full posterior. Working under the frequentist assumption of a ‘‘true’’ underlying model, the aim under the former is to study the *estimation* rate of point estimators under a given risk function, such as expected squared error loss. From a fully Bayes perspective, one may also be interested in the *posterior contraction rate*, or the speed at which the *entire* posterior contracts around the truth. In both cases, the frequentist minimax estimation rate is a useful benchmark, since the posterior cannot contract faster than this rate [18].

In a variety of contexts, including the Gaussian sequence model, sparse linear regression, and graphical models, the SSL global posterior mode has been shown to achieve the minimax estimation rate [45, 42, 48, 16]. From a fully Bayesian perspective, the *entire* posterior under SSL or SSL-type priors has *also* been shown to achieve (near) optimal posterior contraction rates in the contexts of the Gaussian sequence model, linear regression, regres-

sion with grouped variables, nonparametric additive regression, and functional regression [45, 42, 48, 38, 3, 2]. It is not necessarily the case that the posterior mode and the full posterior contract at the same rate [11, 10]. These theoretical results thus show that the SSL is optimal from *both* penalized likelihood *and* fully Bayesian perspectives.

1.9 Discussion

In this chapter, we have reviewed the spike-and-slab LASSO (1.6). The SSL forms a continuum between the penalized likelihood LASSO and the Bayesian point-mass spike-and-slab frameworks, borrowing strength from both constructs while addressing limitations of each. First, the SSL employs a *non*-separable penalty that self-adapts to ensemble information about sparsity and that performs selective shrinkage. Second, the SSL is amenable to fast maximum *a posteriori* finding algorithms which can be implemented in a highly efficient, scalable manner. Third, the posterior mode under the SSL prior can automatically perform both variable selection and estimation. Finally, uncertainty quantification for the SSL can be attained by either debiasing the posterior modal estimate or by utilizing efficient approaches to posterior sampling. Beyond linear regression, the spike-and-slab LASSO methodology is broadly applicable to a wide number of statistical problems, including generalized linear models, factor analysis, graphical models, and nonparametric regression.

Acknowledgments

This work was supported by funding from the University of South Carolina College of Arts & Sciences, the James S. Kemper Foundation Faculty Research Fund at the University of Chicago Booth School of Business, and NSF Grants DMS-1916245, DMS-1944740.



Bibliography

- [1] Joseph Antonelli, Giovanni Parmigiani, and Francesca Dominici. High-dimensional confounding adjustment using continuous spike and slab priors. *Bayesian Analysis*, 14(3):805–828, 09 2019.
- [2] Ray Bai, Mary R. Boland, and Yong Chen. Fast algorithms and theory for high-dimensional Bayesian varying coefficient models. *arXiv pre-print arXiv: 1907.06477*, 2020.
- [3] Ray Bai, Gemma E. Moran, Joseph L. Antonelli, Yong Chen, and Mary R. Boland. Spike-and-slab group lassos for grouped regression and sparse generalized additive models. *Journal of the American Statistical Association (to appear)*, 2020.
- [4] Maria Maddalena Barbieri and James O. Berger. Optimal predictive model selection. *The Annals of Statistics*, 32(3):870–897, 06 2004.
- [5] A. Belloni, V. Chernozhukov, and L. Wang. Square-root lasso: pivotal recovery of sparse signals via conic programming. *Biometrika*, 98(4):791–806, 2011.
- [6] Anirban Bhattacharya, Antik Chakraborty, and Bani K. Mallick. Fast sampling with Gaussian scale mixture priors in high-dimensional regression. *Biometrika*, 103(4):985–991, 2016.
- [7] MaÅgorzata Bogdan, Ewout van den Berg, Chiara Sabatti, Weijie Su, and Emmanuel J. CandÃs. SLOPEÅAdaptive variable selection via convex optimization. *The Annals of Applied Statistics*, 9(3):1103–1140, 2015.
- [8] Leonard Bottolo and Sylvia Richardson. Evolutionary stochastic search for Bayesian model exploration. *Bayesian Analysis*, 5(3):583–618, 09 2010.
- [9] Carlos M. Carvalho, Nicholas G. Polson, and James G. Scott. Handling sparsity via the horseshoe. In David van Dyk and Max Welling, editors, *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, volume 5 of *Proceedings of Machine Learning Research*, pages 73–80, Hilton Clearwater Beach Resort, Clearwater Beach, Florida USA, 2009. PMLR.
- [10] IsmaÃl Castillo and Romain Mismar. Empirical Bayes analysis of spike and slab posterior distributions. *Electronic Journal of Statistics*, 12(2):3953–4001, 2018.
- [11] IsmaÃl Castillo, Johannes Schmidt-Hieber, and Aad van der Vaart. Bayesian linear regression with sparse priors. *The Annals of Statistics*, 43(5):1986–2018, 10 2015.
- [12] Annie P Chiang, John S Beck, Hsan-Jan Yen, Marwan K Tayeh, Todd E Scheetz, Ruth E Swiderski, Darryl Y Nishimura, Terry A Braun, Kwang-Youn A Kim, Jian Huang, et al. Homozygosity mapping with SNP arrays identifies TRIM32, an E3 ubiquitin ligase, as a Bardet–Biedl syndrome gene (BBS11). *Proceedings of the National Academy of Sciences*, 103(16):6287–6292, 2006.

- [13] Sameer K. Deshpande, Veronika Ročková, and Edward I. George. Simultaneous variable and covariance selection with the multivariate spike-and-slab lasso. *Journal of Computational and Graphical Statistics*, 28(4):921–931, 2019.
- [14] Jianqing Fan and Runze Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96(456):1348–1360, 2001.
- [15] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1):1–22, 2010.
- [16] Lingrui Gan, Naveen N. Narisetty, and Feng Liang. Bayesian regularization for graphical models with unequal shrinkage. *Journal of the American Statistical Association*, 114:1218–1231, 2018.
- [17] Edward I. George and Robert E. McCulloch. Variable selection via Gibbs sampling. *Journal of the American Statistical Association*, 88(423):881–889, 09 1993.
- [18] Subhashis Ghosal, Jayanta K. Ghosh, and Aad W. van der Vaart. Convergence rates of posterior distributions. *The Annals of Statistics*, 28(2):500–531, 04 2000.
- [19] Thomas L. Griffiths and Zoubin Ghahramani. The Indian buffet process: An introduction and review. *Journal fo Machine Learning Research*, 12:1185–1224, July 2011.
- [20] Chris Hans, Adrian Dobra, and Mike West. Shotgun stochastic search for ℓ_1 regression. *Journal of the American Statistical Association*, 102(478):507–516, 2007.
- [21] Hemant Ishwaran and J. Sunil Rao. Spike and slab variable selection: Frequentist and bayesian strategies. *The Annals of Statistics*, 33(2):730–773, 04 2005.
- [22] Adel Javanmard and Andrea Montanari. Debiasing the lasso: Optimal sample size for gaussian designs. *The Annals of Statistics*, 46(6A):2593–2622, 2018.
- [23] Wei Jiang, Malgorzata Bogdan, Julie Josse, Blazej Miasojedow, Veronika Rockova, and TraumaBase Group. Adaptive bayesian slope – high-dimensional model selection with missing values. *arXiv pre-print arXiv: 1907.06477*, 2019.
- [24] James Johndrow, Paulo Orenstein, and Anirban Bhattacharya. Scalable approximate MCMC algorithms for the horseshoe prior. *Journal of Machine Learning Research*, 21(73):1–61, 2020.
- [25] Iain M. Johnstone and Bernard W. Silverman. Needles and straw in haystacks: Empirical Bayes estimates of possibly sparse sequences. *The Annals of Statistics*, 32(4):1594–1649, 08 2004.
- [26] Iain M. Johnstone and Bernard W. Silverman. Empirical Bayes selection of wavelet thresholds. *The Annals of Statistics*, 33(4):1700–1752, 08 2005.
- [27] Youngseok Kim and Chao Gao. Bayesian model selection with graph structured sparsity. *arXiv preprint arXiv:1902.03316*, 2019.
- [28] John P. Klein and Melvin L. Moeschberger. *Survival Analysis Techniques for Censored and Truncated Data*. Second edition, 2003.

- [29] Zehang Li, Tyler McCormick, and Samuel Clark. Bayesian joint spike-and-slab graphical lasso. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 3877–3885, Long Beach, California, USA, 09–15 Jun 2019. PMLR.
- [30] Rahul Mazumder, Jerome H. Friedman, and Trevor Hastie. Sparsenet: Coordinate descent with nonconvex penalties. *Journal of the American Statistical Association*, 106(495):1125–1138, 2011. PMID: 25580042.
- [31] G. J. McLachlan and K. E. Basford. *Mixture models: Inference and applications to clustering*. Marcel Dekker, New York, 1988.
- [32] Nicolai Meinshausen and Peter Bühlmann. High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics*, 34(3):1436–1462, 2006.
- [33] T.J. Mitchell and J.J. Beauchamp. Bayesian variable selection in linear regression. *Journal of the American Statistical Association*, 83(404):1023–1032, 1988.
- [34] Gemma E. Moran, Veronika Ročková, and Edward I. George. Spike-and-slab lasso biclustering. *The Annals of Applied Statistics (to appear)*, 2020.
- [35] Gemma E. Moran, Veronika Ročková, and Edward I. George. Variance prior forms for high-dimensional Bayesian variable selection. *Bayesian Analysis*, 14(4):1091–1119, 2019.
- [36] Naveen N. Narisetty, Juan Shen, and Xuming He. Skinny Gibbs: A consistent and scalable Gibbs sampler for model selection. *Journal of the American Statistical Association*, 114(527):1205–1217, 2019.
- [37] Naveen Naidu Narisetty and Xuming He. Bayesian variable selection with shrinking and diffusing priors. *The Annals of Statistics*, 42(2):789–817, 04 2014.
- [38] Lizhen Nie and Veronika Ročková. Fast posterior sampling for the spike-and-slab LASSO. *preprint*, 2020.
- [39] Trevor Park and George Casella. The Bayesian lasso. *Journal of the American Statistical Association*, 103(482):681–686, 2008.
- [40] Nicholas G. Polson and Lei Sun. Bayesian ℓ_0 -regularized least squares. *Applied Stochastic Models in Business and Industry*, 35(3):717–731, 2019.
- [41] Kolyan Ray and Botond Szabó. Variational Bayes for high-dimensional linear regression with sparse priors. *arXiv preprint arXiv:1904.07150*, 2019.
- [42] Veronika Ročková and Edward I George. Bayesian penalty mixing: The case of a non-separable penalty. In *Statistical Analysis for High-Dimensional Data - The Abel Symposium 2014*, pages 233–254. Springer, 2016.
- [43] Paul R. Rosenbaum and Donald B. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 04 1983.
- [44] Veronika Ročková and Gemma E. Moran. *SSLASSO: The Spike-and-Slab LASSO*, 2018. R package version 1.2-1.
- [45] Veronika Ročková. Bayesian estimation of sparse signals with a continuous spike-and-slab prior. *The Annals of Statistics*, 46(1):401–437, 2018.

- [46] Veronika Ročková and Edward I. George. EMVS: The EM approach to Bayesian variable selection. *Journal of the American Statistical Association*, 109(506):828–846, 2014.
- [47] Veronika Ročková and Edward I. George. Fast Bayesian factor analysis via automatic rotations to sparsity. *Journal of the American Statistical Association*, 111(516):1608–1622, 2016.
- [48] Veronika Ročková and Edward I. George. The spike-and-slab LASSO. *Journal of the American Statistical Association*, 113(521):431–444, 2018.
- [49] Todd E Scheetz, Kwang-Youn A Kim, Ruth E Swiderski, Alisdair R Philp, Terry A Braun, Kevin L Knudtson, Anne M Dorrance, Gerald F DiBona, Jian Huang, Thomas L Casavant, 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.
- [50] James G Scott and James O Berger. Bayes and empirical-Bayes multiplicity adjustment in the variable-selection problem. *The Annals of Statistics*, 38(5):2587–2619, 2010.
- [51] Tingni Sun and Cun-Hui Zhang. Scaled sparse linear regression. *Biometrika*, 99(4):879–898, 2012.
- [52] Zaixiang Tang, Shufeng Lei, Xinyan Zhang, Zixuan Yi, Boyi Guo, Jake Y. Chen, Yueping Shen, and Nengjun Yi. Gsslasso Cox: a Bayesian hierarchical model for predicting survival and detecting associating genes by incorporating pathway information. *BMC Bioinformatics*, 20(1), 2019.
- [53] Zaixiang Tang, Yueping Shen, Yan Li, Xinyan Zhang, Jia Wen, Chen Zao Qian, Wenzhuo Zhuang, Xinghua Shi, and Nengjun Yi. Group spike-and-slab lasso generalized linear models for disease prediction and associated genes detection by incorporating pathway information. *Bioinformatics*, 34(6):901–910, 2018.
- [54] Zaixiang Tang, Yueping Shen, Xinyan Zhang, and Nengjun Yi. The spike-and-slab lasso generalized linear models for prediction and associated genes detection. *Genetics*, 205(1):77–88, 2017.
- [55] Zaixiang Tang, Yueping Shen, Xinyan Zhang, and Nengjun Yi. The spike-and-slab lasso Cox model for survival prediction and associated genes detection. *Bioinformatics*, 33(18):2799–2807, 05 2017.
- [56] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 58:267–288, 1996.
- [57] Naonori Ueda and Ryohei Nakano. Deterministic annealing EM algorithm. *Neural Networks*, 11(2):271 – 282, 1998.
- [58] Sara van de Geer, Peter Bühlmann, Yaacov Ritov, and Ruben Dezeure. On asymptotically optimal confidence regions and tests for high-dimensional models. *The Annals of Statistics*, 42(3):1166–1202, 2014.
- [59] Wellcome Trust. Genome-wide association study of 14,000 cases of seven common diseases and 3000 shared controls. *Nature*, 447:661–678, 2007.
- [60] Ming Yuan and Yi Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006.

- [61] Cun-Hui Zhang. Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics*, 38(2):894–942, 04 2010.
- [62] Cun-Hui Zhang and Stephanie S. Zhang. Confidence intervals for low dimensional parameters in high dimensional linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(1):217–242, 2014.
- [63] Cun-Hui Zhang and Tong Zhang. A general theory of concave regularization for high-dimensional sparse estimation problems. *Statistical Science*, 27(4):576–593, 2012.
- [64] Hui Zou. The adaptive lasso and its oracle properties. *Journal of the American Statistical Association*, 101(476):1418–1429, 2006.
- [65] Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.