

Group Spike and Slab Variational Bayes

Michael Komodromos*, Kolyan Ray, Marina Evangelou and Sarah Filippi
Department of Mathematics, Imperial College London

November 15, 2023

Abstract

We introduce Group Spike-and-slab Variational Bayes (GSVB), a scalable method for group sparse regression. A fast co-ordinate ascent variational inference (CAVI) algorithm is developed for several common model families including Gaussian, Binomial and Poisson. Theoretical guarantees for our proposed approach are provided by deriving contraction rates for the variational posterior in grouped linear regression. Through extensive numerical studies, we demonstrate that GSVB provides state-of-the-art performance, offering a computationally inexpensive substitute to MCMC, whilst performing comparably or better than existing MAP methods. Additionally, we analyze three real world datasets wherein we highlight the practical utility of our method, demonstrating that GSVB provides parsimonious models with excellent predictive performance, variable selection and uncertainty quantification.

Keywords: High-dimensional regression, Group sparsity, Generalized Linear Models, Sparse priors, Bayesian Methods

*M.K. gratefully acknowledge EPSRC's StatML CDT, Imperial's CRUK center and Imperial's Experimental Cancer Medicine center.

For the purpose of open access, the author has applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising.

1 Introduction

Group structures arise in various applications, such as genetics (Wang et al., 2007; Breheny and Huang, 2009), imaging (Lee and Cao, 2021), multi-factor analysis of variance (Meier et al., 2008), non-parametric regression (Huang et al., 2010), and multi-task learning, among others. In these settings, p -dimensional feature vectors, $x_i = (x_{i1}, \dots, x_{ip})^\top \in \mathbb{R}^p$, for $i = 1, \dots, n$ observations can be partitioned into groups of features. Formally this means that we can construct sub-vectors, $x_{G_k} = \{x_j : j \in G_k\}$ for $k = 1, \dots, M$, where the groups $G_k = \{G_{k,1}, \dots, G_{k,m_k}\}$ are disjoint sets of indices satisfying $\bigcup_{k=1}^M G_k = \{1, \dots, p\}$. Typically, these group structures are known beforehand, for example in genetics where biological pathways (gene sets) are known, or they are constructed artificially, for example, through a basis expansion in non-parametric additive models.

In the regression setting incorporating these group structures is crucial, as disregarding them can result in sub-optimal models (Huang and Zhang, 2010; Lounici et al., 2011). In this manuscript, we focus on the general linear regression model (GLM) where, for each observation $i = 1, \dots, n$, the response, Y_i can be modelled by

$$\mathbb{E}[Y_i|x_i, \beta] = f\left(\sum_{k=1}^M x_{i,G_k}^\top \beta_{G_k}\right), \quad i = 1, \dots, n \quad (1)$$

where $f : \mathbb{R} \rightarrow \mathbb{R}$ represents the link function, $\beta = (\beta_1, \dots, \beta_p)^\top \in \mathbb{R}^p$ denotes the model coefficient vector with $\beta_{G_k} = \{\beta_j : j \in G_k\}$. Beyond incorporating the group structure, it is often of practical importance to identify the groups of features that are associated with the response. This holds particularly true when there are a large number of them. To address this, various methods have been proposed over the years, with one of the most popular being the group LASSO (Yuan and Lin, 2006), which applies an $\ell_{2,1}$ norm to groups of coefficients. Following Yuan and Lin (2006) there have been numerous extensions, including the group SCAD (Wang et al., 2007), the group LASSO for logistic regression (Meier et al., 2008),

the group bridge (Huang et al., 2009), group LASSO with overlapping groups (Jacob et al., 2009), and the sparse group LASSO (Breheny and Huang, 2009; Simon et al., 2013), among others (see Huang et al. (2012) for a detailed review of frequentist methods).

In a similar vein, Bayesian group selection methods have arisen. The earliest of which being the Bayesian group LASSO (Raman et al., 2009; Kyung et al., 2010), which uses a multivariate double exponential distribution prior to impose shrinkage on groups of coefficients. Notably, the maximum a posteriori (MAP) estimate under this prior coincides with the estimate under the group LASSO. Other methods include Bayesian sparse group selection (Xu and Ghosh, 2015; Chen et al., 2016) and the group spike-and-slab LASSO (Bai et al., 2020) which approach the problem via stochastic search variable selection (Mitchell and Beauchamp, 1988; Chipman, 1996). Formally, these methods utilize a group spike-and-slab prior, a mixture distribution of a multivariate Dirac mass on zero and a continuous distribution over \mathbb{R}^{m_k} , where m_k is the size of the k th group. Such priors have been shown to work excellently for variable selection as they are able to set coefficients exactly to zero, avoiding the use of shrinkage to enforce sparsity. For a comprehensive review see Lai and Chen (2021) and Jreich et al. (2022).

However, a serious drawback of these methods is that most use Markov Chain Monte Carlo (MCMC) to sample from the posterior (Raman et al., 2009; Xu and Ghosh, 2015; Chen et al., 2016). In general, MCMC approaches are known to be computationally expensive when there are a large number of variables, and within the group setting they can result in poor mixing when group sizes are large (Chen et al., 2016). To circumvent these issues, some authors proposed computing MAP estimates (Kyung et al., 2010; Bai et al., 2020), by relaxing the form of the prior, replacing the multivariate Dirac mass with a continuous distribution concentrated at zero (Ročková and George, 2018). Although these

algorithms are fast to compute, they come at the sacrifice of interpretability, as posterior inclusion probabilities no longer guarantee the coefficient is zero but rather concentrated at zero. Beyond this, these algorithms only return a point estimate for β and therefore do not provide uncertainty quantification – a task at the heart of Bayesian inference.

To bridge the gap between scalability and uncertainty quantification several authors have turned to variational inference (VI). An approach to inference wherein the posterior distribution is approximated by a tractable family of distributions known as the variational family (see [Zhang et al. \(2019\)](#) for a review). In the context of high-dimensional Bayesian inference, VI has proven particularly successful, and has been employed in linear regression ([Carbonetto and Stephens, 2012](#); [Ormerod et al., 2017](#); [Ray and Szabó, 2022](#)), logistic regression ([Ray et al., 2020](#)) and survival analysis ([Komodromos et al., 2022](#)) to name a few. Within the context of Bayesian group regression, to our knowledge only the Bayesian group LASSO has seen a variational counterpart ([Babacan et al., 2014](#)).

In this manuscript we examine the variational Bayes (VB) posterior arising from the group spike-and-slab prior with multivariate double exponential slab and Dirac spike, referring to our method as Group Spike-and-slab Variational Bayes (GSVB). We provide scalable variational approximations to three common classes of generalized linear models: the Gaussian with identity link function, Binomial with logistic link function, and Poisson with exponential link function. We outline a general scheme for computing the variational posterior via co-ordinate ascent variational inference. We further show that for specific cases, the variational family can be re-parameterized to allow for more efficient updates.

Through extensive numerical experiments we demonstrate that GSVB achieves state-of-the-art performance while significantly reducing computation time by several orders of magnitude compared to MCMC. Moreover, through our comparison with MCMC, we high-

light that the variational posterior provides excellent uncertainty quantification through marginal credible sets with impressive coverage. Additionally, the proposed method is compared against the spike-and-slab group LASSO (Bai et al., 2020), a state-of-the-art Bayesian group selection algorithm that returns MAP estimates. Within this comparison our method demonstrates comparable or better performance in terms of group selection and effect estimation, whilst carrying the added benefit of providing uncertainty quantification, a feature not available by other scalable methods in the literature.

To highlight the practical utility of our method, we analyse three real datasets, demonstrating that our method provides excellent predictive accuracy, while also achieving parsimonious models. Additionally, we illustrate the usefulness of the VB posterior by showing its ability to provide posterior predictive intervals, a feature not available to methods that provide MAP estimates, and computationally prohibitive to compute via MCMC.

Theoretical guarantees for our method are provided in the form of posterior contraction rates, which quantify how far the posterior places most of its probability from the ‘ground truth’ generating the data for large sample sizes. This approach builds on the work of Castillo et al. (2015) and Ray and Szabó (2022), and extends previous Bayesian contraction rate results for true posteriors in the group sparse setting (Ning et al., 2020; Bai et al., 2020) to their variational approximation.

Concurrent to our work, Lin et al. (2023) proposed a similar spike-and-slab model for group variable selection in linear regression, which appeared on arXiv shortly after our preprint.

Notation. Let $y = (y_1, \dots, y_n)^\top \in \mathbb{R}^n$ denote the realization of the random vector $Y = (Y_1, \dots, Y_n)$. Further, let $X = (x_1, \dots, x_n)^\top \in \mathbb{R}^{n \times p}$ denote the design matrix, where for a group G_k let $X_{G_k} = (x_{1,G_k}, \dots, x_{n,G_k})^\top \in \mathbb{R}^{n \times m_k}$. Similarly, for $G_k^c = \{1, \dots, p\} \setminus G_k$,

denote $X_{G_k^c} = (x_{1,G_k^c}, \dots, x_{n,G_k^c})^\top \in \mathbb{R}^{n \times (p-m_k)}$ where $x_{i,G_k^c} = \{x_{ij} : j \in G_k^c\}$, and $\beta_{G_k^c} = \{\beta_j : j \in G_k^c\}$. Wherein, without loss of generality we assume that the elements of the groups are ordered such that $1 = G_{1,1} < G_{1,2} < \dots < G_{M,m_M} = p$. Finally, the Kullback-Leibler divergence is defined as $D_{\text{KL}} = D_{\text{KL}}(Q||P) = \int_{\mathcal{X}} \log\left(\frac{dQ}{dP}\right) dQ$, where Q and P are probability measures on \mathcal{X} , such that Q is absolutely continuous with respect to P .

2 Prior and Variational family

To model the coefficients β , we consider a group spike-and-slab prior. For each group G_k , the prior over β_{G_k} consists of a mixture distribution of a multivariate Dirac mass on zero and a multivariate double exponential distribution, $\Psi(\beta_{G_k})$, whose density is given by,

$$\psi(\beta_{G_k}; \lambda) = C_k \lambda^{m_k} \exp(-\lambda \|\beta_{G_k}\|)$$

where $C_k = [2^{m_k} \pi^{(m_k-1)/2} \Gamma((m_k+1)/2)]^{-1}$ and $\|\cdot\|$ is the ℓ_2 -norm. In the context of sparse Bayesian group regression the multivariate double exponential has been previously considered by [Raman et al. \(2009\)](#) and [Kyung et al. \(2010\)](#) as part of the Bayesian group LASSO and by [Xu and Ghosh \(2015\)](#) within a group spike-and-slab prior.

Formally the prior, which we consider throughout, is given by $\Pi(\beta) = \bigotimes_{k=1}^M \Pi_k(\beta_{G_k})$, where each $\Pi_k(\beta_{G_k})$ has the hierarchical representation,

$$\begin{aligned} \beta_{G_k} | z_k &\stackrel{\text{ind}}{\sim} z_k \Psi(\beta_{G_k}; \lambda) + (1 - z_k) \delta_0(\beta_{G_k}) \\ z_k | \theta_k &\stackrel{\text{ind}}{\sim} \text{Bernoulli}(\theta_k) \\ \theta_k &\stackrel{\text{iid}}{\sim} \text{Beta}(a_0, b_0) \end{aligned} \tag{2}$$

for $k = 1, \dots, M$, where $\delta_0(\beta_{G_k})$ is the multivariate Dirac mass on zero with dimension $m_k = \dim(\beta_{G_k})$. In conjunction with the log-likelihood, $\ell(\mathcal{D}; \beta)$ for a dataset $\mathcal{D} = \{(y_i, x_i)\}_{i=1}^n$,

we write the posterior density under the given prior as,

$$d\Pi(\beta|\mathcal{D}) = \Pi_D^{-1} e^{\ell(\mathcal{D};\beta)} d\Pi(\beta) \quad (3)$$

where $\Pi_D = \int_{\mathbb{R}^p} e^{\ell(\mathcal{D};\beta)} d\Pi(\beta)$ is a normalization constant known as the model evidence.

The posterior arising from the prior (2) and the dataset \mathcal{D} assigns probability mass to all 2^M possible sub-models, i.e. each subset $\mathcal{S} \subseteq \{1, \dots, M\}$, such that $z_k = 1, k \in \mathcal{S}$ and $z_k = 0$ otherwise. As using MCMC procedures to sample from this complex posterior distribution is computationally prohibitive, even for a moderate number of groups, we resort to variational inference to approximate it. This approximation is known as the variational posterior which is an element of the variational family \mathcal{Q} given by

$$\tilde{\Pi} = \underset{Q \in \mathcal{Q}}{\operatorname{argmin}} D_{\text{KL}}(Q \parallel \Pi(\cdot|\mathcal{D})) . \quad (4)$$

For our purposes the variational family we consider is a mean-field variational family,

$$\mathcal{Q} = \left\{ Q(\mu, \sigma, \gamma) = \bigotimes_{k=1}^M Q_k(\mu_{G_k}, \sigma_{G_k}, \gamma_k) := \bigotimes_{k=1}^M [\gamma_k N_k(\mu_{G_k}, \operatorname{diag}(\sigma_{G_k}^2)) + (1 - \gamma_k)\delta_0] \right\} \quad (5)$$

where $\mu \in \mathbb{R}^p$ with $\mu_{G_k} = \{\mu_j : j \in G_k\}$, $\sigma^2 \in \mathbb{R}_+^p$ with $\sigma_{G_k}^2 = \{\sigma_j^2 : j \in G_k\}$, $\gamma = (\gamma_1, \dots, \gamma_M)^\top \in [0, 1]^M$. $N_k(\mu, \Sigma)$ denotes the multivariate Normal distribution with mean parameter μ and covariance Σ . Notably, under $Q \in \mathcal{Q}$, the vector of coefficients for each group G_k is a spike-and-slab distribution where the slab consists of the product of independent Normal distributions,

$$\beta_{G_k} \stackrel{\text{ind}}{\sim} \gamma_k \left[\bigotimes_{j \in G_k} N(\mu_j, \sigma_j^2) \right] + (1 - \gamma_k)\delta_0,$$

meaning the structure (correlations) between elements within the same group are not captured. To mitigate this, a second variational family is introduced, where the covariance within groups is unrestricted. Formally,

$$\mathcal{Q}' = \left\{ Q'(\mu, \Sigma, \gamma) = \bigotimes_{k=1}^M Q'_k(\mu_{G_k}, \Sigma_{G_k}, \gamma_k) := \bigotimes_{k=1}^M [\gamma_k N(\mu_{G_k}, \Sigma_{G_k}) + (1 - \gamma_k)\delta_0] \right\} \quad (6)$$

where $\Sigma \in \mathbb{R}^{p \times p}$ is a covariance matrix for which $\Sigma_{ij} = 0$, for $i \in G_k, j \in G_l, k \neq l$ and $\Sigma_{G_k} = (\Sigma_{ij})_{i,j \in G_k} \in \mathbb{R}^{m_k \times m_k}$ denotes the covariance matrix of the k th group. Notably, \mathcal{Q}' should provide greater flexibility when approximating the posterior, because unlike $\mathcal{Q} \subset \mathcal{Q}'$, it is able to capture the dependence between coefficients in the same group. The importance of which is highlighted empirically in Section 5 wherein the two families are compared.

Note that the posterior does not take the form Q or Q' , as the use of these variational families replaces the 2^M model weights by M VB group inclusion probabilities, γ_k , thereby introducing substantial additional independence. For example, information as to whether two groups of variables are selected together or not is lost. However, the form of the VB approximation retains many of the interpretable features of the original posterior such as the inclusion probabilities of particular groups.

3 Computing the Variational Posterior

Computing the variational posterior defined in (4) relies on optimizing a lower bound on the model evidence, $\Pi_{\mathcal{D}}$, known as the evidence lower bound (ELBO). The ELBO follows from the non-negativity of the Kullback-Leibler divergence and is given by $\mathbb{E}_Q \left[\ell(\mathcal{D}; \beta) - \log \frac{dQ}{d\Pi} \right]$. Intuitively, the first term evaluates the model’s fit to the data, while the second term acts as a regularizer, ensuring the variational posterior is “close” to the prior distribution.

Various strategies exist to tackle the minimization of the negative ELBO,

$$\mathcal{L}_Q(\mathcal{D}) := \mathbb{E}_Q \left[\log \frac{dQ}{d\Pi} - \ell(\mathcal{D}; \beta) \right] \quad (7)$$

with respect to $Q \in \mathcal{Q}$ (Zhang et al., 2019). One popular approach is co-ordinate ascent variational inference (CAVI), where sets of parameters of the variational family are optimized in turn while the remainder are kept fixed. Although this strategy does not (in

general) guarantee the global optimum, it is easy to implement and often leads to scalable inference algorithms. In the following subsections we detail how this algorithm is constructed, noting that derivations are presented under the variational family \mathcal{Q}' , as the equations for $\mathcal{Q} \subset \mathcal{Q}'$ follow by restricting Σ_{G_k} to be a diagonal matrix.

3.1 Computing the Evidence lower bound

To compute the negative ELBO, note that the KL divergence between the prior and the variational distribution, $D_{\text{KL}}(Q' \parallel \Pi) = \mathbb{E}_{Q'} [\log(dQ'/d\Pi)]$ is constant regardless of the form of the likelihood. To evaluate an expression for this quantity the group independence structure between the prior and variational distribution is exploited, allowing for the log Radon-Nikodym, $\log(dQ'/d\Pi)$, to be expressed as,

$$\log \frac{dQ'}{d\Pi}(\beta) = \sum_{k=1}^M \log \frac{dQ'_k}{d\Pi_k}(\beta_{G_k}) = \sum_{k=1}^M \mathbb{I}_{z_k=1} \log \frac{\gamma_k dN_k}{\bar{w} d\Psi_k}(\beta_{G_k}) + \mathbb{I}_{z_k=0} \log \frac{(1-\gamma_k)d\delta_0}{(1-\bar{w})d\delta_0}(\beta_{G_k}).$$

where $\bar{w} = \alpha_0/(\alpha_0 + b_0)$. Subsequently, it follows that,

$$\begin{aligned} D_{\text{KL}}(Q' \parallel \Pi) &= \sum_{k=1}^M \left(\gamma_k \log \frac{\gamma_k}{\bar{w}} - \frac{\gamma_k}{2} \log(\det(2\pi\Sigma_{G_k})) - \frac{\gamma_k m_k}{2} - \gamma_k \log(C_k) \right. \\ &\quad \left. - \gamma_k m_k \log(\lambda) + \mathbb{E}_{Q'} [\mathbb{I}_{z_k=1} \lambda \|\beta_{G_k}\|] + (1-\gamma_k) \log \frac{1-\gamma_k}{1-\bar{w}} \right) \end{aligned} \quad (8)$$

where we have used the fact that $\mathbb{E}_{\beta_{G_k} \sim N(\mu_{G_k}, \Sigma_{G_k})} [(\beta_{G_k} - \mu_{G_k})^\top \Sigma_{G_k}^{-1} (\beta_{G_k} - \mu_{G_k})] = m_k$. Notably, there is no closed form for $\mathbb{E}_{Q'} [\mathbb{I}_{z_k=1} \lambda \|\beta_{G_k}\|]$, meaning $\mathcal{L}_{Q'}(\mathcal{D})$ is not tractable and optimization over this quantity would require costly Monte Carlo methods.

To circumvent this issue, a surrogate objective is constructed and used instead of the negative ELBO. This follows by applying Jensen's inequality to $\mathbb{E}_{Q'} [\mathbb{I}_{z_k=1} \lambda \|\beta_{G_k}\|]$, giving,

$$\mathbb{E}_{Q'} [\mathbb{I}_{z_k=1} \lambda \|\beta_{G_k}\|] = \gamma_k \mathbb{E}_{N_k} [\lambda \|\beta_{G_k}\|] \leq \gamma_k \lambda \left(\sum_{i \in G_k} \Sigma_{ii} + \mu_i^2 \right)^{1/2}. \quad (9)$$

Thus, $D_{\text{KL}}(Q' \parallel \Pi)$ can be upper-bounded by,

$$\begin{aligned} \varrho(\mu, \Sigma, \gamma) := & \sum_{k=1}^M \left(\gamma_k \log \frac{\gamma_k}{\bar{w}} - \frac{\gamma_k}{2} \log(\det(2\pi\Sigma_{G_k})) - \frac{\gamma_k m_k}{2} \right. \\ & \left. - \gamma_k \log(C_k) - \gamma_k m_k \log(\lambda) + \gamma_k \lambda \left(\sum_{i \in G_k} \Sigma_{ii} + \mu_i^2 \right)^{1/2} + (1 - \gamma_k) \log \frac{1 - \gamma_k}{1 - \bar{w}} \right) \end{aligned} \quad (10)$$

and in turn, $\mathcal{L}_{Q'}(\mathcal{D}) \leq \varrho(\mu, \Sigma, \gamma) - \mathbb{E}_{Q'}[\ell(\mathcal{D}; \beta)]$. Via this upper bound, we are able to construct a tractable surrogate objective. Formally, this objective is denoted by, $\mathcal{F}(\mu, \Sigma, \gamma, \vartheta)$, where we introduce ϑ to parameterize additional hyperparameters, specifically those introduced to model the variance term under the Gaussian family, and those introduced to bound the Binomial likelihood. Under this parametrization we define,

$$\mathcal{F}(\mu, \Sigma, \gamma, \vartheta) := \varrho(\mu, \Sigma, \gamma) + \tilde{\varrho}(\vartheta) + \Lambda(\mu, \Sigma, \gamma, \vartheta) \geq \mathcal{L}_{Q'}(\mathcal{D}), \quad (11)$$

where $\tilde{\varrho}(\vartheta) : \Theta \rightarrow \mathbb{R}$ and $\Lambda(\mu, \Sigma, \gamma, \vartheta) \geq \mathbb{E}_{Q'}[-\ell(\mathcal{D}; \beta)]$ (which is an inequality under non-tractable likelihoods). To this end, what remains to be computed is an expression for the expected negative log-likelihood and the term $\tilde{\varrho}(\vartheta)$ where appropriate. In the following three subsections we provide derivations of the objective function $\mathcal{F}(\mu, \Sigma, \gamma, \theta)$ for the Gaussian, Binomial and Poisson likelihoods taking the canonical link functions for each.

3.1.1 Gaussian

Under the Gaussian family with identity link function, $Y_i \stackrel{\text{iid}}{\sim} N(x_i^\top \beta, \tau^2)$ for all $i = 1, \dots, n$, the log-likelihood is given by $\ell(\mathcal{D}; \beta, \tau^2) = -\frac{n}{2} \log(2\pi\tau^2) - \frac{1}{2\tau^2} \|y - X\beta\|^2$. To model τ^2 an inverse Gamma prior is considered due to its popularity among practitioners, i.e. $\tau^2 \stackrel{\text{ind}}{\sim} \Gamma^{-1}(a, b)$ which has density $\frac{b^a}{\Gamma(a)} x^{-a-1} \exp(-\frac{b}{x})$ where $a, b > 0$. The prior is therefore given by $\Pi(\beta, \tau^2) = \Pi(\beta)\Gamma^{-1}(\tau^2)$, and the posterior density by $d\Pi(\beta, \tau^2 | \mathcal{D}) = \Pi_D^{-1} e^{\ell(\mathcal{D}; \beta, \tau^2)} d\Pi(\beta, \tau^2)$, where $\Pi_D = \int_{\mathbb{R}^p \times \mathbb{R}_+} e^{\ell(\mathcal{D}; \beta, \tau^2)} d\Pi(\beta, \tau^2)$ and $\mathcal{D} = \{(y_i, x_i)\}_{i=1}^n$ is the observed data.

To include this term within the inference procedure, the variational families \mathcal{Q} and \mathcal{Q}' are extended by $\mathcal{Q}_\tau = \mathcal{Q} \times \{\Gamma^{-1}(a', b') : a', b' > 0\}$ and $\mathcal{Q}'_\tau = \mathcal{Q}' \times \{\Gamma^{-1}(a', b') : a', b > 0\}$ respectively. Under the extended variational family \mathcal{Q}'_τ , the negative ELBO is given by

$$\mathbb{E}_{\mathcal{Q}'_\tau} \left[\log \frac{d\mathcal{Q}'}{d\Pi}(\beta) + \log \frac{d\Gamma^{-1}(a', b')}{d\Gamma^{-1}(a, b)}(\tau^2) - \ell(\mathcal{D}; \beta, \tau^2) \right],$$

where the second term follows from the independence of τ^2 and β in the prior and variational family and is given by,

$$\tilde{\varrho}(\vartheta) := \mathbb{E}_{\mathcal{Q}'_\tau} \left[\log \left(\frac{d\Gamma^{-1}(a', b')}{d\Gamma^{-1}(a, b)} \right) \right] = (a' - a)\kappa(a') + a \log \frac{b'}{b} + \log \frac{\Gamma(a)}{\Gamma(a')} + \frac{(b - b')a'}{b'} \quad (12)$$

where $\vartheta = \{(a', b')\}$. The expectation of the negative log-likelihood is given by,

$$\begin{aligned} \Lambda(\mu, \Sigma, \gamma, \vartheta) := \mathbb{E}_{\mathcal{Q}'_\tau} [-\ell(\mathcal{D}; \beta)] &= \frac{n}{2}(\log(2\pi) + \log(b') - \kappa(a')) \\ &+ \frac{a'}{2b'} \left(\|y\|^2 + \left(\sum_{i=1}^p \sum_{j=1}^p (X^\top X)_{ij} \mathbb{E}_{\mathcal{Q}'_\tau} [\beta_i \beta_j] \right) - 2 \sum_{k=1}^M \gamma_k \langle y, X_{G_k} \mu_{G_k} \rangle \right) \end{aligned} \quad (13)$$

where,

$$\mathbb{E}_{\mathcal{Q}'_\tau} [\beta_i \beta_j] = \begin{cases} \gamma_k (\Sigma_{ij} + \mu_i \mu_j) & i, j \in G_k \\ \gamma_k \gamma_h \mu_i \mu_j & i \in G_k, j \in G_h, h \neq k \end{cases} \quad (14)$$

Substituting (12) and (13) into (11) gives $\mathcal{F}(\mu, \Sigma, \gamma, \theta)$ under the Gaussian family.

3.1.2 Binomial

Under the Binomial family with logistic link, $Y_i \stackrel{\text{iid}}{\sim} \text{Bernoulli}(p_i)$ for all $i = 1, \dots, n$ where $p_i = \mathbb{P}(Y_i = 1 | x_i) = \exp(x_i^\top \beta) / (1 + \exp(x_i^\top \beta))$. The log-likelihood is given by $\ell(\mathcal{D}, \beta) = \sum_{i=1}^n y_i (x_i^\top \beta) - \log(1 + \exp(x_i^\top \beta))$ where $\mathcal{D} = \{(y_i, x_i)\}_{i=1}^n$ with $y_i \in \{0, 1\}$.

Unlike the Gaussian family, variational inference in this setting is challenging because of the intractability of the expected log-likelihood under the variational family. To overcome this issue several authors have proposed bounds or approximations to maintain tractability

(see [Depraetere and Vandebroek \(2017\)](#) for a review). Here we employ a bound introduced by [Jaakkola and Jordan \(1996\)](#), given as

$$s(x) \geq s(t) \exp \left\{ \frac{x-t}{2} - \frac{a(t)}{2}(x^2 - t^2) \right\} \quad (15)$$

where $s(x) = (1 + \exp(-x))^{-1}$ and $a(t) = \frac{s(t)-1/2}{t}$, $x \in \mathbb{R}$ and $t \in \mathbb{R}$ is an additional parameter that must be optimized to ensure tightness of the bound.

Using (15) allows for the negative log-likelihood to be bounded by,

$$-\ell(\mathcal{D}; \beta) \leq \sum_{i=1}^n -y_i x_i^\top \beta - \log s(t_i) + \frac{x_i^\top \beta + t_i}{2} + \frac{a(t_i)}{2} ((x_i^\top \beta)^2 - t_i^2) \quad (16)$$

where $t_i \in \mathbb{R}$ is a hyper-parameter for each observation. Taking the expectation of (16) with respect to the variational family gives

$$\begin{aligned} \mathbb{E}_{Q'} [-\ell(\mathcal{D}; \beta)] \leq \Lambda(\mu, \Sigma, \gamma, \vartheta) := & \sum_{i=1}^n \left(\sum_{k=1}^M \gamma_k (1/2 - y_i) x_{i,G_k}^\top \mu_{G_k} \right) + \frac{t_i}{2} - \log s(t_i) \\ & + \frac{a(t_i)}{2} \left(\left(\sum_{j=1}^p \sum_{l=1}^p (x_{ij} x_{il} \mathbb{E}_{Q'} [\beta_j \beta_l]) \right) - t_i^2 \right) \end{aligned} \quad (17)$$

where $\vartheta = \{t_1, \dots, t_n\}$. In turn substituting (17) and $\tilde{\varrho}(\vartheta) = 0$ into (11) gives the objective under the Binomial family.

3.1.3 Poisson

Finally, under the Poisson family with exponential link function, $Y_i \stackrel{\text{iid}}{\sim} \text{Poisson}(\lambda_i)$ for all $i = 1, \dots, n$ with $\lambda_i = \exp(x_i^\top \beta) > 0$. The log-likelihood is given by, $\ell(\mathcal{D}; \beta) = \sum_{i=1}^n y_i x_i^\top \beta - \exp(x_i^\top \beta) - \log(y!)$, whose (negative) expectation is tractable and given by,

$$\mathbb{E}_{Q'} [-\ell(\mathcal{D}; \beta)] = \Lambda(\mu, \Sigma, \gamma, \vartheta) := \sum_{i=1}^n \log(y!) - M_{Q'}(x_i) + \left(\sum_{k=1}^M \gamma_k y_i x_{i,G_k}^\top \mu_{G_k} \right) \quad (18)$$

where $M_{Q'}(x_i) = \prod_{k=1}^M M_{Q_k}(x_{i,G_k})$ is the moment generating function under the variational family, with $M_{Q_k}(x_{i,G_k}) := \gamma_k M_{N_k}(x_{i,G_k}) + (1 - \gamma_k)$ being the moment generating function

for the k th group and $M_{N_k}(x_{i,G_k}) = \exp \left\{ x_{i,G_k}^\top \mu_{G_k} + \frac{1}{2} x_{i,G_k}^\top \Sigma_{G_k} x_{i,G_k} \right\}$. Unlike the previous two families the Poisson family, does not require any additional variational parameters, therefore $\vartheta = \{\}$ and $\tilde{\rho}(\vartheta) = 0$.

3.2 Coordinate ascent algorithm

Recall the aim is to approximate the posterior $\Pi(\cdot|\mathcal{D})$ by a distribution from a given variational family. This approximation is obtained via the minimization of the objective \mathcal{F} derived in the previous section. To achieve this a CAVI algorithm is proposed (Murphy, 2007; Blei et al., 2017) as outlined in Algorithm 1.

In this context, the objective introduced in (11) is written as $\mathcal{F}(\mu, \Sigma, \gamma, \vartheta) = \mathcal{F}(\mu_{G_k}, \mu_{G_k^c}, \Sigma_{G_k}, \Sigma_{G_k^c}, \gamma_k, \gamma_{-k}, \vartheta)$, highlighting the fact that optimization over the variational parameters occurs group-wise. Further, for each group k , while the optimization of the objective function over the inclusion probability, γ_k , can be done analytically, we use the Limited memory Broyden–Fletcher–Goldfarb–Shanno optimization algorithm (L-BFGS) to update μ_{G_k} at each iteration of the CAVI procedure. Details for the optimization with respect to Σ_{G_k} are presented in Section 3.2.1. The hyper-parameters, ϑ , are updated using L-BFGS for the Gaussian family and analytically for those under the Binomial family.

To assess convergence, the total absolute change in the parameters is tracked, terminating when this quantity falls below a specified threshold, set to 10^{-3} in our implementation. Other methods involve monitoring the absolute change in the ELBO, however we found this prohibitively expensive to compute for this purpose.

Algorithm 1 Group sparse co-ordinate ascent variational inference

Initialize $\mu, \Sigma, \gamma, \vartheta$ **while** not converged **for** $k = 1, \dots, M$

$$\mu_{G_k} \leftarrow \operatorname{argmin}_{\mu_{G_k} \in \mathbb{R}^{m_k}} \mathcal{F}(\mu_{G_k}, \mu_{G_k}^c, \Sigma_{G_k}, \Sigma_{G_k}^c, \gamma_k = 1, \gamma_{-k}, \vartheta)$$

$$\Sigma_{G_k} \leftarrow \operatorname{argmin}_{\Sigma_{G_k} \in \mathbb{R}^{m_k \times m_k}} \mathcal{F}(\mu_{G_k}, \mu_{G_k}^c, \Sigma_{G_k}, \Sigma_{G_k}^c, \gamma_k = 1, \gamma_{-k}, \vartheta)$$

$$\gamma_k \leftarrow \operatorname{argmin}_{\gamma_k \in [0,1]} \mathcal{F}(\mu_{G_k}, \mu_{G_k}^c, \Sigma_{G_k}, \Sigma_{G_k}^c, \gamma_k, \gamma_{-k}, \vartheta)$$

$$\vartheta \leftarrow \operatorname{argmin}_{\vartheta \in \Theta} \mathcal{F}(\mu_{G_k}, \mu_{G_k}^c, \Sigma_{G_k}, \Sigma_{G_k}^c, \gamma_k, \gamma_{-k}, \vartheta)$$

return $\mu, \Sigma, \gamma, \vartheta$.

3.2.1 Re-parameterization of Σ_{G_k}

Our focus now turns to the optimization of $\mathcal{F}(\mu_{G_k}, \mu_{G_k}^c, \Sigma_{G_k}, \Sigma_{G_k}^c, \gamma_k = 1, \gamma_{-k}, \vartheta)$ w.r.t. Σ_{G_k} . By using similar ideas to those of [Seeger \(1999\)](#) and [Oppor and Archambeau \(2009\)](#), it can be shown that only one free parameter is needed to describe the optimum of Σ_{G_k} under the Gaussian (see below) and Binomial family (see Section [A](#) of the Supplementary material). The objective w.r.t Σ_{G_k} under the Gaussian family, is given by,

$$\frac{a'}{2b'} \operatorname{tr}(X_{G_k}^\top X_{G_k} \Sigma_{G_k}) - \frac{1}{2} \log \det \Sigma_{G_k} + \lambda \left(\sum_{i \in G_k} \Sigma_{ii} + \mu_i^2 \right)^{1/2} + C \quad (19)$$

where C is a constant that does not depend on Σ_{G_k} . Differentiating [\(19\)](#) w.r.t. Σ_{G_k} , setting to zero and re-arranging gives,

$$\Sigma_{G_k} = \left(\frac{a'}{b'} X_{G_k}^\top X_{G_k} + 2\nu_k I_{m_k} \right)^{-1} \quad (20)$$

where $\nu_k = \frac{1}{2} \lambda (\sum_{i \in G_k} \Sigma_{ii} + \mu_i^2)^{-1/2}$. Thus, substituting $\Sigma_{G_k} = \left(\frac{a'}{b'} X_{G_k}^\top X_{G_k} + w_k I_{m_k} \right)^{-1}$ where $w_k \in \mathbb{R}$ into the objective and optimizing over w_k is equivalent to optimizing over Σ_{G_k} , and carries the added benefit of requiring one free parameter to perform rather than

$m_k(m_k - 1)/2$. Note that under this re-parametrization the inversion of an $m_k \times m_k$ matrix is required which can be a time consuming operation for large m_k .

For the Poisson family the same re-parameterization cannot be used, so Σ_{G_k} is parameterized by $U_k^\top U_k$ where $U_k \in \mathbb{R}^{m_k \times m_k}$ is an upper triangular matrix. Optimization is then performed on the upper triangular elements of U_k .

3.2.2 Initialization

As with any gradient-descent based approach, our CAVI algorithm is sensitive to the choice of initial values. We suggest to initialize μ using the group LASSO from the package `gglasso` using a small regularization parameter. This ensures many of the elements are non-zero. The covariance matrix Σ can be initialised by using the re-parametrization outlined in Section 3.2.1 with an initial value of $w_k = 1$ for $k = 1, \dots, M$ for both the Gaussian and the Binomial families. For the Poisson family we propose the use of an initial covariance matrix $\Sigma = \text{diag}(0.2, \dots, 0.2)$. Finally, the inclusion probabilities are initialised as $\gamma = (0.5, \dots, 0.5)^\top$. For the additional hyper-parameters ϑ are used: $a' = b' = 10^{-3}$ and $t_i = \left(\sum_{k=1}^M \gamma_k [\langle \mu_{G_k}, x_{i,G_k} \rangle^2 + x_{i,G_k}^\top \Sigma_{G_k} x_{i,G_k}] \right)^{1/2}$ for all $i = 1, \dots, n$.

4 Theoretical results for grouped linear regression

4.1 Notation and assumptions

This section establishes frequentist theoretical guarantees for the proposed VB approach in sparse high-dimensional linear regression with group structure. The full proofs of the following results are provided in the Supplementary Material.

To simplify technicalities, the variance parameter τ^2 is taken as known and equal to 1,

giving model $Y = X\beta + \varepsilon$ with $Y \in \mathbb{R}^n$, $X \in \mathbb{R}^{n \times p}$ and $\varepsilon \sim N_n(0, I_n)$. Under suitable conditions, contraction rates for the variational posterior are established, which quantify its spread around the ‘ground truth’ parameter $\beta_0 \in \mathbb{R}^p$ generating the data as $n, p \rightarrow \infty$.

Recall that the covariates are split into M pre-specified disjoint groups G_1, \dots, G_M of size $|G_k| = m_k$ with $\sum_{k=1}^M m_k = p$ and maximal group size $m_{\max} = \max_{k=1, \dots, M} m_k$. The above model can then be written as

$$Y = \sum_{k=1}^M X_{G_k} \beta_{G_k} + \varepsilon, \quad (21)$$

with $\beta_{G_k} \in \mathbb{R}^{m_k}$ and $X_{G_k} \in \mathbb{R}^{n \times m_k}$. Let P_β denote the law of Y under (21), $S_\beta \subseteq \{G_1, \dots, G_M\}$ be the set containing the indices of the non-zero *groups* in $\beta \in \mathbb{R}^p$. For a vector $\beta \in \mathbb{R}^p$ and set S , we also write $\beta_S = (\beta_i)_{i \in G_k: G_k \in S} \in \mathbb{R}^{\sum_{G_k \in S} m_k}$ for its vector restriction to S . We write β_0 for the ground truth generating the data, $S_0 = S_{\beta_0}$ and $s_0 = |S_0|$ for its *group-sparsity*. For a matrix $A \in \mathbb{R}^{m \times n}$, let $\|A\|_F^2 = \sum_{i=1}^m \sum_{j=1}^n A_{ij}^2 = \text{Tr}(A^T A)$ be the Frobenius norm and define the group matrix norm of $X \in \mathbb{R}^{n \times p}$ by

$$\|X\| = \max_{k=1, \dots, M} \|X_{G_k}\|_F.$$

If all the groups are singletons, $\|X\|$ reduces to the same norm as in [Castillo et al. \(2015\)](#).

We further define the $\ell_{2,1}$ -norm of a vector by $\|\beta\|_{2,1} = \sum_{k=1}^M \|\beta_{G_k}\|_2$. We assume that the prior slab scale λ satisfies

$$\underline{\lambda} \leq \lambda \leq 2\bar{\lambda}, \quad \underline{\lambda} = \frac{\|X\|}{M^{1/m_{\max}}}, \quad \bar{\lambda} = 3\|X\|\sqrt{\log M}, \quad (22)$$

mirroring the situation without grouping ([Castillo et al., 2015](#); [Ray and Szabó, 2022](#)).

The parameter β in (21) is not estimable without additional assumptions on the design matrix X , for instance that $X^T X$ is invertible for sparse subspaces of \mathbb{R}^p . These notions of invertibility can be precise via the following definitions, which are the natural adaptations

of compatibility conditions to the group sparse setting (Bühlmann and van de Geer, 2011; Castillo et al., 2015).

Definition 1. *A model $S \subseteq \{1, \dots, M\}$ has compatibility number*

$$\phi(S) = \inf \left\{ \frac{\|X\beta\|_2 |S|^{1/2}}{\|X\| \|\beta\|_{2,1}} : \|\beta_{S^c}\|_{2,1} \leq 7 \|\beta_S\|_{2,1}, \beta_S \neq 0 \right\}.$$

Compatibility considers only vectors whose coordinates are small outside S , and hence is a (weaker) notion of approximate rather than exact sparsity. For all β in the above set, it holds that $\|X\beta\|_2 |S|^{1/2} \geq \phi(S) \|X\| \|\beta\|_{2,1}$, which can be interpreted as a form of continuous invertibility of X for approximately sparse vectors in the sense that changes in β_S lead to sufficiently large changes in $X\beta$ that can be detected by the data. The number 7 is not important and is taken in Definition 2.1 of Castillo et al. (2015) to provide a specific numerical value; since we use several of their techniques, we employ the same convention. We next consider two further notions of invertibility for exact group sparsity.

Definition 2. *The compatibility number for vectors of dimension s is*

$$\bar{\phi}(s) = \inf \left\{ \frac{\|X\beta\|_2 |S_\beta|^{1/2}}{\|X\| \|\beta\|_{2,1}} : 0 \neq |S_\beta| \leq s \right\}.$$

Definition 3. *The smallest scaled sparse singular value of dimension s is*

$$\tilde{\phi}(s) = \inf \left\{ \frac{\|X\beta\|_2}{\|X\| \|\beta\|_2} : 0 \neq |S_\beta| \leq s \right\}.$$

While $X^T X$ is not generally invertible in the high-dimensional setting, these last two definitions weaken this requirement to sparse vectors. These are natural extensions of the definitions in Castillo et al. (2015) to the group setting, and similar interpretations and relations to the usual sparse setting apply also here, see Bühlmann and van de Geer (2011) or Section 2.2 in Castillo et al. (2015) for further discussion.

The interplay of the group structure and individual sparsity can lead to multiple regimes see for instance [Lounici et al. \(2011\)](#) and [Bühlmann and van de Geer \(2011\)](#). To make our results more interpretable, we restrict to the main case of practical interest where the group sizes are not too large, and hence the group sparsity drives the estimation rate.

Assumption (K). *There exists $K > 0$ such that $m_{\max} \log m_{\max} \leq K \log M$.*

While related works make similar assumptions ([Bai et al., 2020](#)), introducing an explicit constant $K > 0$ above allows us to clarify the uniformity in our results.

4.2 Asymptotic results

We now state our main result on variational posterior contraction for both *prediction loss* $\|X(\beta - \beta_0)\|_2$ and the usual ℓ_2 -loss. Our results are uniform over vectors in sets of the form

$$\mathcal{B}_{\rho_n, s_n} = \mathcal{B}_{\rho_n, s_n}(c_0) := \{\beta_0 \in \mathbb{R}^p : \phi(S_{\beta_0}) \geq c_0, |S_{\beta_0}| \leq s_n, \tilde{\phi}(\rho_n |S_{\beta_0}|) \geq c_0\}, \quad (23)$$

where $s_n \geq 1$, $c > 0$ and $\rho_n \rightarrow \infty$ (arbitrarily slowly).

Theorem 1 (Contraction). *Suppose that Assumption (K) holds, the prior (2) satisfies (22) and s_n satisfies $m_{\max} \log s_n = O(\log M)$. Then the variational posterior $\tilde{\Pi}$ based on either the variational family \mathcal{Q} in (5) or \mathcal{Q}' in (6) satisfies, with $s_0 = |S_{\beta_0}|$,*

$$\sup_{\beta_0 \in \mathcal{B}_{\rho_n, s_n}} E_{\beta_0} \tilde{\Pi} \left(\beta : \|X(\beta - \beta_0)\|_2 \geq \frac{H_0 \rho_n^{1/2} \sqrt{s_0 \log M}}{\tilde{\phi}(\rho_n s_0)} \right) \rightarrow 0$$

$$\sup_{\beta_0 \in \mathcal{B}_{\rho_n, s_n}} E_{\beta_0} \tilde{\Pi} \left(\beta : \|\beta - \beta_0\|_2 \geq \frac{H_0 \rho_n^{1/2} \sqrt{s_0 \log M}}{\|X\| \tilde{\phi}(\rho_n s_0)^2} \right) \rightarrow 0,$$

for any $\rho_n \rightarrow \infty$ (arbitrarily slowly), $\mathcal{B}_{\rho_n, s_n}$ defined in (23) and where H_0 depends only on the prior.

Contraction rates for the full posterior based on a group spike and slab prior were established in [Ning et al. \(2020\)](#), as well as for the spike and slab group LASSO in [Bai](#)

et al. (2020). Our proofs are instead based on those in Castillo et al. (2015) and Ray and Szabó (2022). We have extended their theoretical results from the coordinate sparse setting to the group sparse setting. This approach permits more explicit proofs and allows us to consider somewhat different assumptions from previous group sparse works (Ning et al., 2020; Bai et al., 2020), for instance removing the boundedness assumption for the parameter spaces.

Remark 1. *The optimization problem (4) is in general non-convex, and hence there is no guarantee that CAVI (or any other algorithm) will converge to the global minimizer $\tilde{\Pi}$. However, an inspection of the proofs shows that the conclusions of Theorems 1 and 2 apply also to any element $Q^* \in \mathcal{Q} \subset \mathcal{Q}'$ in the variational families for which*

$$0 \leq D_{KL}(Q^* \|\Pi(\cdot|Y)) - D_{KL}(\tilde{\Pi} \|\Pi(\cdot|Y)) = \mathcal{L}_{Q^*}(\mathcal{D}) - \mathcal{L}_{\tilde{\Pi}}(\mathcal{D}) = O(s_0 \log M),$$

where s_0 is the true group sparsity and $\mathcal{L}_Q(\mathcal{D})$ is the negative ELBO. Thus, as long as the ELBO is within $O(s_0 \log M)$ of its maximum, the resulting variational approximation will satisfy the above conclusions, even if it is not the global optimum.

The next result shows that the variational posterior puts most of its mass on models of size at most a multiple of the true number of groups, meaning that it concentrates on sparse sets.

Theorem 2 (Dimension). *Suppose that Assumption (K) holds, the prior (2) satisfies (22) and s_n satisfies $m_{\max} \log s_n = O(\log M)$. Then the variational posterior $\tilde{\Pi}$ based on either the variational family \mathcal{Q} in (5) or \mathcal{Q}' in (6) satisfies*

$$\sup_{\beta_0 \in \mathcal{B}_{\rho_n, s_n}} E_{\beta_0} \tilde{\Pi}(\beta : |S_\beta| \geq \rho_n |S_{\beta_0}|) \rightarrow 0$$

for any $\rho_n \rightarrow \infty$ (arbitrarily slowly) and $\mathcal{B}_{\rho_n, s_n}$ defined in (23).

5 Numerical experiments

In this section a numerical evaluation of our method, referred to as Group Spike-and-slab Variational Bayes (GSVB), is presented¹. Where necessary we distinguish between the two families, \mathcal{Q} and \mathcal{Q}' , by the suffix ‘-D’ and ‘-B’ respectively, i.e. GSVB-D denotes the method under the variational family \mathcal{Q} . Notably, throughout all our numerical experiments the prior parameters are set to $\lambda = 1$, $\alpha_0 = 1$, $b_0 = M$, and $a = b = 10^{-3}$ for the inverse-Gamma prior on τ^2 under the Gaussian family.

This section begins with a comparison of GSVB² against MCMC to assess the quality of the variational posterior in terms of variable selection and uncertainty quantification. Following this, a large-scale comparison with the Spike-and-Slab Group LASSO (SSGL) proposed by Bai et al. (2020), a state-of-the-art MAP Bayesian group selection method described in detail in Section 5.3, is performed.

5.1 Simulation set-up

Data is simulated for $i = 1, \dots, n$ observations each having a response $y_i \in \mathbb{R}$ and p continuous predictors $x_i \in \mathbb{R}^p$. The response is sampled independently from the respective family with mean given by $f(\beta_0^\top x_i)$ where f is the link function (and variance applicable to the Gaussian family of $\tau^2 = 1$). The true coefficient vector $\beta_0 = (\beta_{0,G_1}, \dots, \beta_{0,G_M})^\top \in \mathbb{R}^p$ consists of M groups each of size m i.e. $M \times m = p$. Of these groups, s are chosen at random to be non-zero and have each of their element values sampled independently and uniformly from $[-\beta_{\max}, 0.2] \cup [0.2, \beta_{\max}]$ where $\beta_{\max} = 1.5, 1.0$ and 0.45 for the Gaussian, Binomial and Poisson families respectively. Finally, the predictors are generated from one

¹Scripts to reproduce our results can be found at <https://github.com/mkomod/p3>

²An R package is available at <https://github.com/mkomod/gsvb>.

of four settings:

- **Setting 1:** $x_i \stackrel{\text{iid}}{\sim} N(0_p, I_p)$
- **Setting 2:** $x_i \stackrel{\text{iid}}{\sim} N(0_p, \Sigma)$ where $\Sigma_{ij} = 0.6^{|i-j|}$ for $i, j = 1, \dots, p$.
- **Setting 3:** $x_i \stackrel{\text{iid}}{\sim} N(0, \Sigma)$ where Σ is a block diagonal matrix where each block A is a 50×50 square matrix such that $A_{jl} = 0.6, j \neq l$ and $A_{jj} = 1$ otherwise.
- **Setting 4:** $x_i \stackrel{\text{iid}}{\sim} N(0, \Sigma)$ where $\Sigma = (1 - \alpha)W^{-1} + \alpha V^{-1}$ with $W \sim \text{Wishart}(p + \nu, I_p)$ and V is a block diagonal matrix of M blocks, where each block V_k , for $k = 1, \dots, M$, is an $m_k \times m_k$ matrix given by $V_k \sim \text{Wishart}(m_k + \nu, I_{m_k})$; we let $(\nu, \alpha) = (3, 0.9)$ so that predictors within groups are more correlated than variables between groups.

To evaluate the performance of the methods four different metrics are considered: (i) the ℓ_2 -error, $\|\widehat{\beta} - \beta_0\|_2$, between the true vector of coefficients and the estimated coefficient $\widehat{\beta}$ defined as the posterior mean where applicable, or the maximum a posteriori (MAP) estimate if this is returned, (ii) the area under the curve (AUC) of the receiver operator characteristic curve, which compares true positive and false positive rate for different thresholds of the group posterior inclusion probability, (iii) the marginal coverage of the non-zero coefficients, which reports the proportion of times the true coefficient $\beta_{0,j}$ is contained in the marginal credible set $\{j : \beta_{0,j} \neq 0\}$, (iv) and the size of the marginal credible set for the non-zero coefficients, given by the Lebesgue measure of the set. The last two metrics can only be computed when a distribution for β is available i.e. via MCMC or VB. The 95% marginal credible sets for each variable $j \in G_k$ for $k = 1, \dots, M$ are given by:

$$S_j = \begin{cases} \{0\} & \text{if } \gamma_k < \alpha \\ \left[\mu_j \pm \Sigma_{jj}^{1/2} \Phi^{-1}\left(\frac{\alpha\gamma_k}{2}\right) \right] & \text{if } \gamma_k \geq \alpha \text{ and } 0 \notin \left[\mu_j \pm \Sigma_{jj}^{1/2} \Phi^{-1}\left(\frac{\alpha\gamma_k}{2}\right) \right] \\ \left[\mu_j \pm \Sigma_{jj}^{1/2} \Phi^{-1}\left(\frac{\alpha\gamma_k}{2} + \frac{1-\gamma_k}{2}\right) \right] \cup \{0\} & \text{otherwise} \end{cases}$$

where $\alpha_{\gamma_k} = 1 - \frac{\alpha}{\gamma_k}$ and Φ^{-1} is the quantile function of the standard Normal distribution.

5.2 Comparison to MCMC

Our first numerical study compares GSVB to the posterior obtained via MCMC, often considered the gold standard in Bayesian inference. The details of the MCMC sampler used for this comparison are outlined in Section C of the Supplementary material³. The MCMC sampler is ran for 100,000 iterations taking a burn-in period of 50,000 iterations.

Within this comparison we set $p = 1,000$, $m = 5$, and vary the number of non-zero groups, s . As highlighted in Figure 1 GSVB-B performs excellently in nearly all settings, demonstrating comparable results to MCMC in terms of ℓ_2 -error and AUC. This indicates that GSVB-B exhibits similar characteristics to MCMC, both in terms of the selected groups and the posterior mean. As anticipated, all the methods exhibit better performance in simpler settings and show a decline in performance as the problem complexity increases.

Regarding coverage, whilst MCMC shows slightly better performance compared to GSVB, the proposed method still provides credible sets that capture a significant portion of the true non-zero coefficients (particularly GSVB-B). However, the credible sets of the variational posterior are sometimes not large enough to capture the true non-zero coefficients. This observation is further supported by the size of the marginal credible sets, with MCMC producing the largest sets, followed by GSVB-B and GSVB-D. These findings confirm the well-known fact that VB tends to underestimate the posterior variance (Carbonetto and Stephens, 2012; Blei et al., 2017; Zhang et al., 2019; Ray et al., 2020).

Interestingly, the set size is larger under \mathcal{Q}' , highlighting the fact that the mean field variational family (\mathcal{Q}) lacks the necessary flexibility to accurately capture the underlying

³An implementation is available at <https://github.com/mkomod/spsl>.

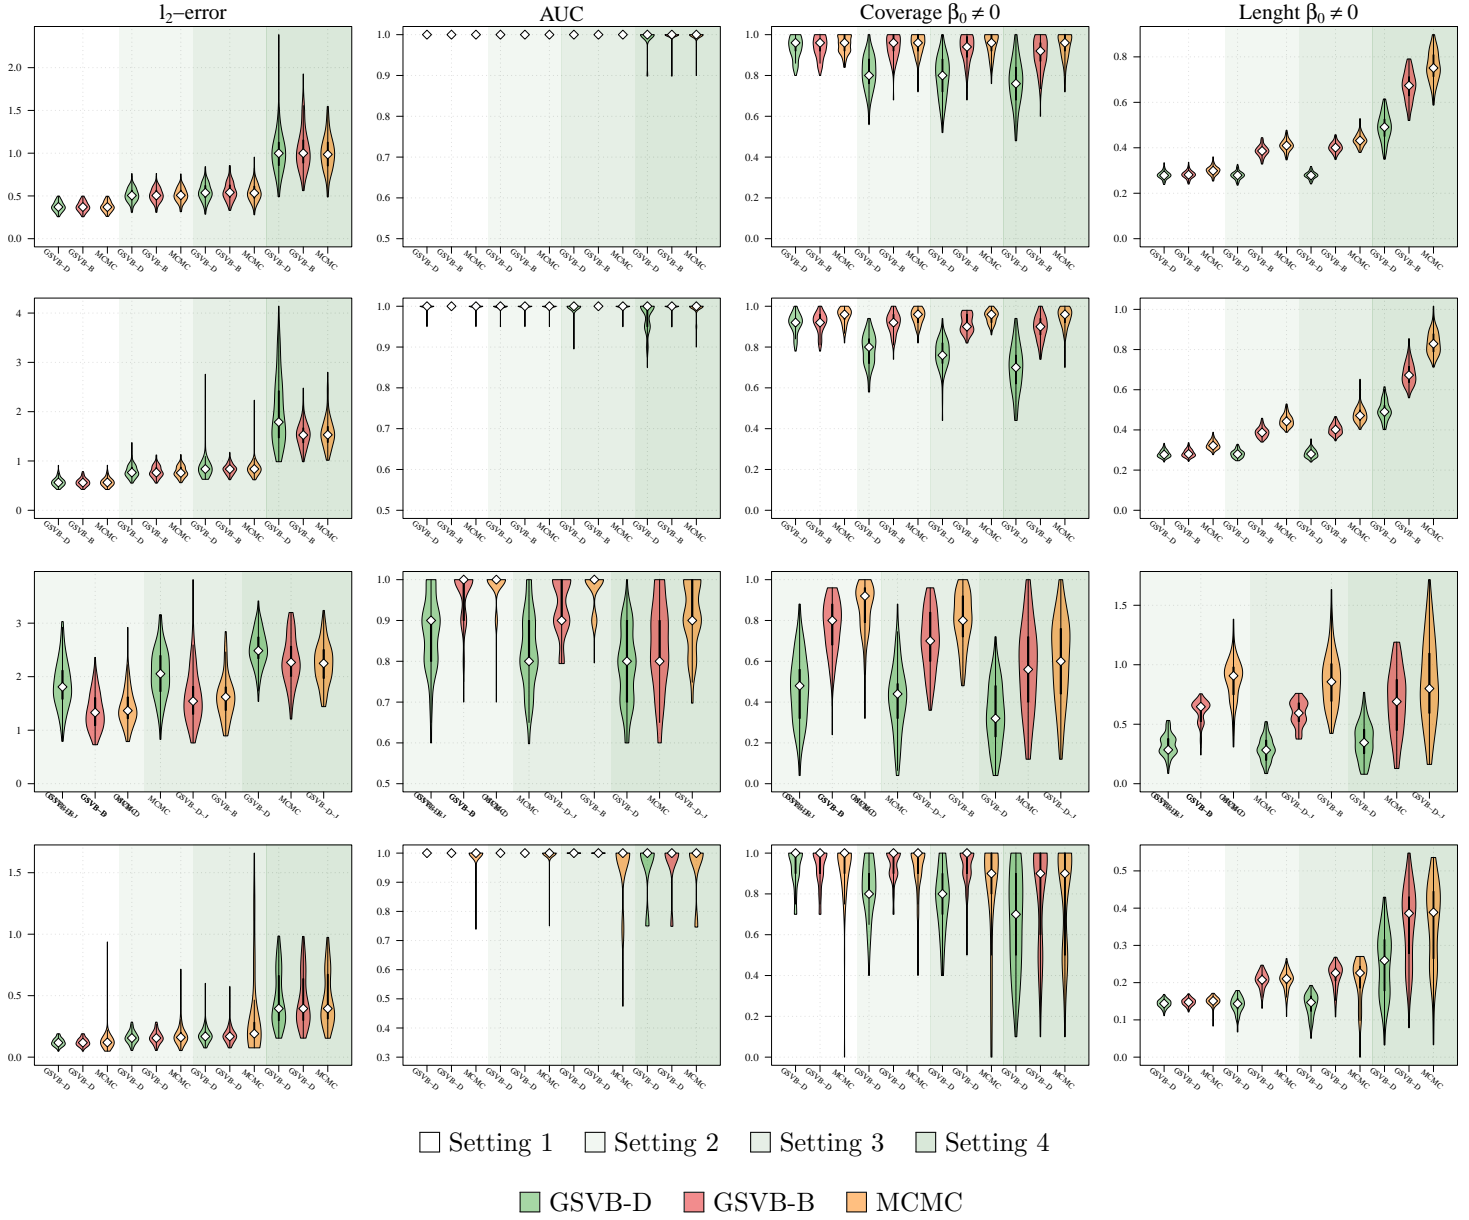


Figure 1: Performance evaluation of GSVB and MCMC for Settings 1–4 with $p = 1,000$ across 100 runs. For each method the white diamond (\diamond) indicates the median of the metric, the thick black line (—) the interquartile range, and the black line (—) 1.5 times the interquartile range. **Rows 1–2:** Gaussian family with $(n, m, s) = (200, 5, 5), (200, 5, 10)$. **Row 3:** Binomial family with $(n, m, s) = (400, 5, 5)$. **Row 4:** Poisson family with $(n, m, s) = (400, 5, 2)$. Note that for the Binomial family Setting 1 results in perfect separation of classes and is excluded from the study

structure in the data. Furthermore, this result indicates that the full marginal credible quantity improves through the consideration of the interactions within the group.

5.3 Large scale simulations

In this section, our proposed approach is compared against SSGL (Bai et al., 2020) a state-of-the-art Bayesian method. Notably, SSGL employs a similar prior to that introduced in (2), however the multivariate Dirac mass on zero is replaced with a multivariate double exponential distribution, giving a continuous mixture with one density acting as the spike and the other as the slab, parameterized by λ_0 and λ_1 respectively. Under this prior Bai et al. (2020) derive an EM algorithm, which allows for fast updates, however, only MAP estimates are returned, meaning a posterior distribution for β is not available. In addition, the performance of the proposed method is evaluated on larger datasets. As both methods are scalable with p , in this section it is increased to $p = 5,000$. Both the sample size and the number of active (non-zero) groups, s are varied as illustrated in Figure 2.

In this comparison, we set $\lambda_1 = 1$ for SSGL meaning the slab is identical between the two methods. For the spike for SSGL we took a value of $\lambda_0 = 100$ for the Gaussian and Poisson family and $\lambda_0 = 20$ under the Binomial family. These values were selected to ensure that sufficient mass is concentrated about zero without turning to cross validation to select the value. Finally, we let $a_0 = 1$ and $b_0 = M$ for both methods.

Given that SSGL only returns a point estimate for the vector of coefficients β , the methods are compared in terms of ℓ_2 -error between the estimated coefficient and the true one, as well as in terms of AUC using the group inclusion probabilities. Overall GSVB performs comparatively or better than SSGL in most settings, obtaining a lower ℓ_2 -error and higher AUC (Figure 2). As expected, across the different methods there is a decrease in

performance as the difficulty of the setting increases, meaning all methods perform better when there is less correlation in the design matrix. We note that the runtime of SSGL is marginally faster than our method in Settings 1-3, and faster in Settings 4. This is explained by the fact that SSGL only provides point estimate for β , rather than the full posterior distribution. A full breakdown of the runtimes is presented in Section D.2 of the Supplementary Material. Within this large scale simulation our method provides excellent uncertainty quantification. In particular, GSVB-B provides better coverage of the non-zero coefficients than GSVB-D, which can be justified by the set size. As in our comparison to MCMC we notice that there is an increase in the posterior set size as the difficulty of the setting increases, i.e. when there is an increase in the correlation of the design matrix.

6 Real data analysis

This section presents the analysis of two real world datasets with a third dataset analysis included in the Supplementary Material⁴. The first problem is a linear regression problem wherein $p \gg n$ for which GSVB and SSGL are applied. The second is a logistic regression problem where $p < n$, for which the logistic group LASSO is applied in addition to GSVB and SSGL. As before, $a_0 = 1$, $b_0 = M$ and GSVB was ran with a value of $\lambda = 1$. For SSGL, we let $\lambda_1 = 1$ and λ_0 was chosen via five fold cross validation on the training set.

Overall, our results highlight that GSVB achieves state-of-the-art performance, producing parsimonious models with excellent predictive accuracy. Furthermore, in our analysis of the two datasets, variational posterior predictive distributions are constructed. These highlight the practical utility of GSVB, demonstrating how uncertainty in the predictions is quantified – an added benefit otherwise not available via MAP methods such as SSGL.

⁴R scripts used to produce these results are available at <https://github.com/mkomod/p3>

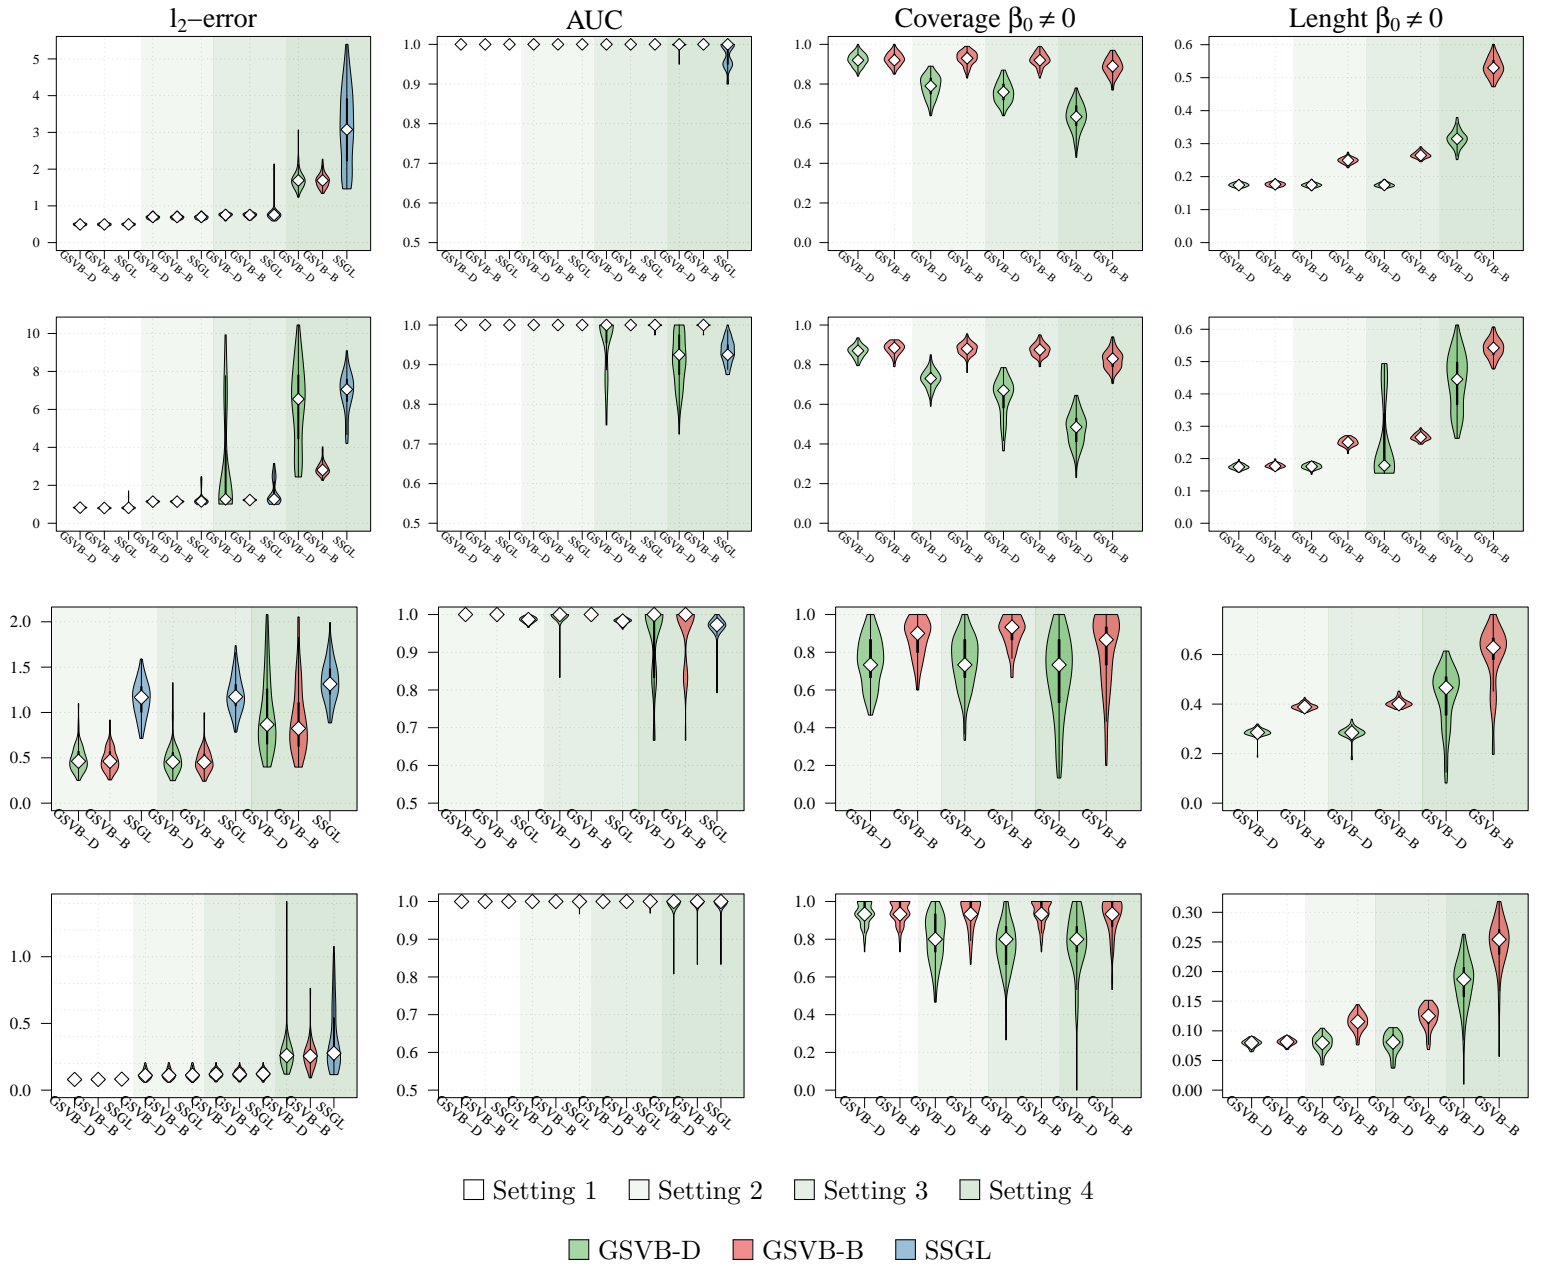


Figure 2: Performance evaluation of GSVB and SSGL for Settings 1–4 with $p = 5,000$ across 100 runs. For each method the white diamond (\diamond) indicates the median of the metric, the thick black line (—) the interquartile range, and the black line (—) 1.5 times the interquartile range. **Rows 1 – 2:** Gaussian family with $(n, m, s) = (500, 10, 10), (500, 10, 20)$. **Row 3:** Binomial family with $(n, m, s) = (1000, 5, 5)$. **Row 4:** Poisson family with $(n, m, s) = (1000, 5, 3)$. Note that for the Binomial family Setting 1 results in perfect separation of classes and is excluded from the study.

6.1 Bardet-Biedl Syndrome Gene Expression Study

The first dataset we examine is of microarray data, consisting of gene expression measurements for 120 laboratory rats⁵. Originally the dataset was studied by [Scheetz et al. \(2006\)](#) and has subsequently been used to demonstrate the performance of several group selection algorithms ([Huang and Zhang, 2010](#); [Breheny and Huang, 2015](#); [Bai et al., 2020](#)). Briefly, the dataset consists of normalized microarray data harvested from the eye tissue of 12 week old male rats. The outcome of interest is the expression of TRIM32 a gene known to cause Bardet-Biedl syndrome, a genetic disease of multiple organs including the retina.

To pre-process the original dataset, which consists of 31,099 probe sets (the predictors), we follow [Breheny and Huang \(2015\)](#) and select the 5,000 probe sets which exhibit the largest variation in expression on the log scale. Further, following [Breheny and Huang \(2015\)](#) and [Bai et al. \(2020\)](#) a non-parametric additive model is used, wherein, $y_i = \mu + \sum_{j=1}^p f_j(x_{ij}) + \epsilon_i$ with $\epsilon_i \sim N(0, \tau^2)$ and $f_j : \mathbb{R} \rightarrow \mathbb{R}$. Here, $y_i \in \mathbb{R}$ corresponds to the expression of TRIM32 for the i th observation, and x_{ij} the expression of the j th probe set for the i th observation. To approximate f_j a three term natural cubic spline basis expansion is used. The resulting processing gave a group regression problem with $n = 120$ and $M = 1,000$ groups of size $m = 3$. Denoting by $\phi_{j,k}$ the k th basis for f_j , we have

$$y_i = \mu + \sum_{j=1}^M \sum_{k=1}^m \phi_{j,k}(x_{ij}) + \epsilon_i . \quad (24)$$

The performance of the methods is evaluated a ten-fold cross validation, where the methods are fit to the validation set and the test set is used for evaluation. More details about the computation of the metrics used to assess method performance can be found in the Supplementary Material. Overall, GSVB performed excellently, in particular GSVB-D obtained a smaller 10-fold cross validated MSE and model size than SSGL, meaning the

⁵available at <ftp://ftp.ncbi.nlm.nih.gov/geo/series/GSE5nnn/GSE5680/matrix/>

models produced are parsimonious and comparably predictive to those of SSGL (Table 1). Furthermore, the PP coverage is excellent, particularly for GSVB–D which has a larger PP coverage than GSVB–B, whilst also having a smaller PP interval length.

Method	MSE	Num. selected groups	PP. coverage	PP. length
GSVB-D	0.017 (0.019)	1.10 (0.316)	0.983 (0.035)	0.5460 (0.100)
GSVB-B	0.018 (0.013)	1.50 (0.527)	0.950 (0.070)	0.5673 (0.091)
SSGL	0.019 (0.001)	2.30 (1.060)	–	–

Table 1: *Bardet-Biedl Syndrome Gene Expression Study*, evaluation metrics computed on the held out test data. Average (sd.) over the 10 folds for the metrics are provided: mean squared error, number of selected groups, posterior predictive coverage, and mean posterior predictive interval length. Notably, GSVB failed to converge on the third fold.

To identify genes associated with *TRIM32* expression, GSVB was ran on the full dataset. Both methods identify one gene each, with GSVB–D identifying *Clec3a* and GSVB–B identifying *Slc25a34*. Interestingly, the MSE was 0.012 and 0.008 for each method, suggesting *Slc25a34* is more predictive of *TRIM32* expression. Further, the PP coverage was 0.960 for both methods and the PP interval length was 0.460 and 0.367 for GSVB–D and GSVB–B, reflecting there is less uncertainty in the prediction under GSVB–B.

6.2 MEMset splice site detection

The following application is based on the prediction of short read DNA motifs, a task which plays an important role in many areas of computational biology. For example, in gene finding, wherein algorithms such as GENIE (Burge and Karlin, 1997) rely on the prediction of splice sites (regions between coding and non-coding DNA segments).

The MEMset donor data has been used to build predictive models for splice sites and

consists of a large training and test set⁶. The original training set contains 8,415 true and 179,438 false donor sites and the test set 4,208 true and 89,717 false donor sites. The predictors are given by 7 factors with 4 levels each (A, T, C, G), a more detailed description is given in [Yeo and Burge \(2004\)](#). Initially analyzed by [Yeo and Burge \(2004\)](#), the data has subsequently been used by [Meier et al. \(2008\)](#) to evaluate the logistic group LASSO.

To create a predictive model we follow [Meier et al. \(2008\)](#) and consider all 3rd order interactions of the 7 factors, which gives a total of $M = 64$ groups and $p = 1,156$ predictors. To balance the sets we randomly sub-sampled the training set without replacement, creating a training set of 1,500 true ($Y = 1$) and 1,500 false donor sites. Regarding the test set a balanced set of 4,208 true and 4,208 false donor sites was created.

In addition to GSVB, we fit SSGL and group LASSO (used in the original analysis), for which we performed 5-fold cross validation to tune the hyperparameters. To assess the different methods, we use the test data, dividing it into 10 folds, and reporting the: (i) precision ($\frac{TP}{TP + FP}$), (ii) recall ($\frac{TP}{TP + FN}$), (iii) F-score ($\frac{TP}{TP + 0.5FP + 0.5FN}$), (iv) AUC, and (v) the maximum correlation coefficient between true class membership and the predicted membership over a range of thresholds, ρ_{\max} , as used in [Yeo and Burge \(2004\)](#) and [Meier et al. \(2008\)](#). In addition we report (vi) the number selected groups, which for the group LASSO is given by the number which have a non-zero $\ell_{2,1}$ norm.

Overall the methods performed comparably. With GSVB-D obtaining the largest recall and the group LASSO the largest precision, F-score, AUC, and ρ_{\max} by a small margin (Table 2). However, GSVB returned models that were of smaller size and therefore more parsimonious than those of SSGL and the group LASSO. This is further highlighted in Figure 5 (found in the Supplementary Material), which showcases the fact that the models obtained by GSVB are far simpler, selecting groups 1–7 as well as the interactions between

⁶available at <http://hollywood.mit.edu/burgelab/maxent/ssdata/>

Method	Precision	Recall	F-score	AUC	ρ_{\max}	Num. selected groups
GSVB–D	0.915 (0.013)	0.951 (0.011)	0.933 (0.009)	0.975 (0.004)	0.870 (0.016)	10
GSVB–B	0.916 (0.011)	0.958 (0.013)	0.936 (0.008)	0.975 (0.004)	0.875 (0.016)	12
SSGL	0.921 (0.012)	0.950 (0.011)	0.935 (0.009)	0.977 (0.004)	0.879 (0.015)	48
Group LASSO	0.924 (0.013)	0.952 (0.010)	0.938 (0.010)	0.977 (0.004)	0.882 (0.016)	60

Table 2: *Memset splice site detection*, evaluation metrics computed on the held out test data. The test set was split into 10 folds and the mean (sd.) of the metrics are computed.

groups 2:4 and 6:7 for both methods, with GSVB–D selecting 5:7 and GSVB–B selecting 2:6, 4:6 and 2:4:6 in addition to these. Further, Figure 5 showcases the fact that uncertainty is available about the norms of the groups.

7 Discussion

In this manuscript we have introduced GSVB, a scalable method for group sparse general linear regression. We have shown how a fast co-ordinate ascent variational inference algorithms can be constructed and used to compute the variational posterior. We have provided theoretical guarantees for the proposed VB approach in the setting of grouped sparse linear regression by extending the theoretical work of [Castillo et al. \(2015\)](#) and [Ray and Szabó \(2022\)](#) for group sparse settings. Through extensive numerical studies we have demonstrated that GSVB provides state-of-the-art performance, offering a computationally inexpensive substitute to MCMC, whilst also performing comparably or better than MAP methods. Additionally, through our analysis of real world datasets we have highlighted the practical utility of our method. Demonstrating, that GSVB provides parsimonious models with excellent predictive performance, and as demonstrated in Appendix E.1, selects

variables with established biological significance. Finally, different to MAP and frequentist methods, GSVB provides scalable uncertainty quantification, which serves as a powerful tool in several application areas.

References

- Babacan, S. D., S. Nakajima, and M. N. Do (2014). Bayesian group-sparse modeling and variational inference. *IEEE Transactions on Signal Processing* 62(11), 2906–2921.
- Bai, R. et al. (2020). Spike-and-Slab Group Lassos for Grouped Regression and Sparse Generalized Additive Models. *Journal of the American Statistical Association*, 1–41.
- Blake, J. A. et al. (2021). Mouse Genome Database (MGD): Knowledgebase for mouse-human comparative biology. *Nucleic Acids Research* 49(D1), D981–D987.
- Blei, D. M., A. Kucukelbir, and J. D. McAuliffe (2017). Variational Inference: A Review for Statisticians. *Journal of the American Statistical Association* 112(518), 859–877.
- Breheny, P. and J. Huang (2009). Penalized methods for bi-level variable selection. *Statistics and Its Interface* 2(3), 369–380.
- Breheny, P. and J. Huang (2015). Group descent algorithms for nonconvex penalized linear and logistic regression models with grouped predictors. *Statistics and Computing* 25(2), 173–187.
- Bühlmann, P. and S. van de Geer (2011). *Statistics for high-dimensional data*. Springer Series in Statistics. Springer, Heidelberg. Methods, theory and applications.
- Burge, C. and S. Karlin (1997). Prediction of complete gene structures in human genomic DNA11Edited by F. E. Cohen. *Journal of Molecular Biology* 268(1), 78–94.

- Carbonetto, P. and M. Stephens (2012). Scalable variational inference for bayesian variable selection in regression, and its accuracy in genetic association studies. *Bayesian Analysis* 7(1), 73–108.
- Castillo, I., J. Schmidt-Hieber, and A. van der Vaart (2015). Bayesian linear regression with sparse priors. *Ann. Statist.* 43(5), 1986–2018.
- Chen, R. B., C. H. Chu, S. Yuan, and Y. N. Wu (2016). Bayesian Sparse Group Selection. *Journal of Computational and Graphical Statistics* 25(3), 665–683.
- Chipman, H. (1996). Bayesian variable selection with related predictors. *Canadian Journal of Statistics* 24(1), 17–36.
- Depraetere, N. and M. Vandebroek (2017). A comparison of variational approximations for fast inference in mixed logit models. *Computational Statistics* 32(1), 93–125.
- Huang, J., P. Breheny, and S. Ma (2012). A selective review of group selection in high-dimensional models. *Statistical Science* 27(4), 481–499.
- Huang, J., J. L. Horowitz, and F. Wei (2010). Variable selection in nonparametric additive models. *Annals of Statistics* 38(4), 2282–2313.
- Huang, J., S. Ma, H. Xie, and C. H. Zhang (2009). A group bridge approach for variable selection. *Biometrika* 96(2), 339–355.
- Huang, J. and T. Zhang (2010). The benefit of group sparsity. *Annals of Statistics* 38(4), 1978–2004.
- Jaakkola, T. S. and M. I. Jordan (1996). A variational approach to Bayesian logistic regression models and their extensions.

- Jacob, L., G. Obozinski, and J. P. Vert (2009). Group lasso with overlap and graph lasso. *ACM International Conference Proceeding Series* 382.
- Jreich, R., C. Hatte, and E. Parent (2022). Review of Bayesian selection methods for categorical predictors using JAGS. *Journal of Applied Statistics* 49(9), 2370–2388.
- Komodromos, M., E. O. Aboagye, M. Evangelou, S. Filippi, and K. Ray (2022, aug). Variational Bayes for high-dimensional proportional hazards models with applications within gene expression. *Bioinformatics* 38(16), 3918–3926.
- Kyung, M., J. Gilly, M. Ghoshz, and G. Casellax (2010). Penalized regression, standard errors, and Bayesian lassos. *Bayesian Analysis* 5(2), 369–412.
- Lai, W. T. and R. B. Chen (2021). A review of Bayesian group selection approaches for linear regression models. *Wiley Interdisciplinary Reviews: Computational Statistics* 13(4), 1–22.
- Lee, K. and X. Cao (2021). Bayesian group selection in logistic regression with application to MRI data analysis. *Biometrics* 77(2), 391–400.
- Lin, B., C. Ge, and J. S. Liu (2023). A variational spike-and-slab approach for group variable selection. arXiv:2309.16855.
- Lounici, K., M. Pontil, S. Van De Geer, and A. B. Tsybakov (2011). Oracle inequalities and optimal inference under group sparsity. *Annals of Statistics* 39(4), 2164–2204.
- Matsui, M. et al. (2010, dec). Activation of LDL Receptor Expression by Small RNAs Complementary to a Noncoding Transcript that Overlaps the LDLR Promoter. *Chemistry & Biology* 17(12), 1344–1355.

- Meier, L., S. Van De Geer, and P. Bühlmann (2008). The group lasso for logistic regression. *Journal of the Royal Statistical Society. Series B: Statistical Methodology* 70(1), 53–71.
- Mitchell, T. J. and J. J. Beauchamp (1988). Bayesian variable selection in linear regression. *Journal of the American Statistical Association* 83(404), 1023–1032.
- Murphy, K. P. (2007). Conjugate bayesian analysis of the gaussian distribution.
- Ning, B., S. Jeong, and S. Ghosal (2020). Bayesian linear regression for multivariate responses under group sparsity. *Bernoulli* 26(3), 2353–2382.
- Opper, M. and C. Archambeau (2009, 03). The variational gaussian approximation revisited. *Neural Computation* 21(3), 786–792.
- Ormerod, J. T., C. You, and S. Müller (2017). A variational bayes approach to variable selection. *Electronic Journal of Statistics* 11(2), 3549–3594.
- Pinelis, I. F. and A. I. Sakhanenko (1986). Remarks on inequalities for large deviation probabilities. *Theory of Probability & Its Applications* 30(1), 143–148.
- Raman, S. et al. (2009). The bayesian group-lasso for analyzing contingency tables. *ACM International Conference Proceeding Series* 382, 881–888.
- Ray, K. and B. Szabó (2022). Variational Bayes for high-dimensional linear regression with sparse priors. *J. Amer. Statist. Assoc.* 117(539), 1270–1281.
- Ray, K., B. Szabo, and G. Clara (2020). Spike and slab variational Bayes for high dimensional logistic regression. In *Advances in Neural Information Processing Systems*, Volume 33, pp. 14423–14434. Curran Associates, Inc.
- Ročková, V. and E. I. George (2018). The Spike-and-Slab LASSO. *Journal of the American Statistical Association* 113(521), 431–444.

- Scheetz, T. E. et al. (2006). Regulation of gene expression in the mammalian eye and its relevance to eye disease. *Proceedings of the National Academy of Sciences of the United States of America* 103(39), 14429–14434.
- Seeger, M. (1999). Bayesian methods for support vector machines and Gaussian processes. Master’s thesis, University of Karlsruhe.
- Silverman, M. G. et al. (2016). Association between lowering LDL-C and cardiovascular risk reduction among different therapeutic interventions: A systematic review and meta-analysis. *Journal of the American Medical Association* 316(12), 1289–1297.
- Simon, N., J. Friedman, T. Hastie, and R. Tibshirani (2013). A sparse-group lasso. *Journal of Computational and Graphical Statistics* 22(2), 231–245.
- Wang, L., G. Chen, and H. Li (2007). Group SCAD regression analysis for microarray time course gene expression data. *Bioinformatics* 23(12), 1486–1494.
- Xu, X. and M. Ghosh (2015). Bayesian variable selection and estimation for group lasso. *Bayesian Analysis* 10(4), 909–936.
- Yeo, G. and C. B. Burge (2004). Maximum entropy modeling of short sequence motifs with applications to RNA splicing signals. *Journal of Computational Biology* 11(2-3), 377–394.
- Yuan, M. and Y. Lin (2006). Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society. Series B: Statistical Methodology* 68(1), 49–67.
- Zhang, C. et al. (2019). Advances in Variational Inference. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41(8), 2008–2026.

SUPPLEMENTARY MATERIAL

A Co-ordinate ascent algorithms

Derivations are presented for the variational family \mathcal{Q}' , noting that $\mathcal{Q} \subset \mathcal{Q}'$, hence the update equations under \mathcal{Q} follow directly from those under \mathcal{Q}' . Recall, the family is given as,

$$\mathcal{Q}' = \left\{ \mathcal{Q}'(\mu, \Sigma, \gamma) = \bigotimes_{k=1}^M \mathcal{Q}'_k(\mu_{G_k}, \Sigma_{G_k}, \gamma_k) := \bigotimes_{k=1}^M [\gamma_k N(\mu_{G_k}, \Sigma_{G_k}) + (1 - \gamma_k)\delta_0] \right\} \quad (25)$$

where $\Sigma \in \mathbb{R}^{p \times p}$ is a covariance matrix for which $\Sigma_{ij} = 0$, for $i \in G_k, j \in G_l, k \neq l$ (i.e. there is independence between groups) and $\Sigma_{G_k} = (\Sigma_{ij})_{i,j \in G_k} \in \mathbb{R}^{m_k \times m_k}$ denotes the covariance matrix of the k th group.

A.1 Gaussian Family

Under the Gaussian family, $Y_i \stackrel{\text{iid}}{\sim} N(x_i^\top \beta, \tau^2)$, and the log-likelihood is given by,

$$\ell(\mathcal{D}; \beta, \tau^2) = -\frac{n}{2} \log(2\pi\tau^2) - \frac{1}{2\tau^2} \|y - X\beta\|^2.$$

Recall, we have chosen to model τ^2 by an inverse-Gamma prior which has density,

$$\frac{b^a}{\Gamma(a)} x^{-a-1} \exp\left(\frac{-b}{x}\right)$$

where $a, b > 0$ and in turn we extend \mathcal{Q}' to $\mathcal{Q}'_\tau = \mathcal{Q}' \times \{\Gamma^{-1}(a', b') : a', b' > 0\}$.

Under this variational family, the expectation of the negative log-likelihood is given by,

$$\begin{aligned} \mathbb{E}_{\mathcal{Q}'_\tau} [-\ell(\mathcal{D}; \beta)] &= \mathbb{E}_{\mathcal{Q}'_\tau} \left[\frac{n}{2} \log(2\pi\tau^2) + \frac{1}{2\tau^2} \|y - X\beta\|^2 \right] \\ &= \mathbb{E}_{\mathcal{Q}'_\tau} \left[\frac{n}{2} \log(2\pi\tau^2) + \frac{1}{2\tau^2} (\|y\|^2 + \|X\beta\|^2 - 2\langle y, X\beta \rangle) \right] \\ &= \frac{n}{2} (\log(2\pi) + \log(b') - \kappa(a')) \end{aligned}$$

$$+ \frac{a'}{2b'} \left(\|y\|^2 + \left(\sum_{i=1}^p \sum_{j=1}^p (X^\top X)_{ij} \mathbb{E}_{Q'_\tau} [\beta_i \beta_j] \right) - 2 \sum_{k=1}^M \gamma_k \langle y, X_{G_k} \mu_{G_k} \rangle \right) \quad (26)$$

where the expectation

$$\mathbb{E}_{Q'_\tau} [\beta_i \beta_j] = \begin{cases} \gamma_k (\Sigma_{ij} + \mu_i \mu_j) & i, j \in G_k \\ \gamma_k \gamma_h \mu_i \mu_j & i \in G_k, j \in G_h, h \neq k \end{cases}$$

A.2 Binomial Family Re-parameterization

To re-parametrization of Σ_{G_k} we follow the same process as in Section 3.2.1. Formally, we have,

$$\Sigma_{G_k} = (X_{G_k}^\top A_t X_{G_k} + wI)^{-1} \quad (27)$$

where $w \in \mathbb{R}$ is the free parameter to be optimized and $A_t = \text{diag}(a(t_1), \dots, a(t_n))$. Notably, this result follows due to the quadratic nature of the bound employed. Meaning we are able to re-parameterize Σ_{G_k} .

B Derivation of Variational Posterior Predictive

To compute the variational posterior predictive distribution for the Binomial and Poisson model, we sample $\beta \sim \tilde{\Pi}$ and then sample $y^* | x^*, \mathcal{D}$ from the respective distribution with mean given by $f(x^* \top \beta)$ where $x^* \in \mathbb{R}^p$ and f is the link function. To sample from $\tilde{\Pi}$ we sample $z_k \stackrel{\text{ind}}{\sim} \text{Bernoulli}(\gamma_k)$ and then $\beta_{G_k} \stackrel{\text{ind}}{\sim} N(\mu_{G_k}, \Sigma_{G_k})$ if $z_k = 1$ otherwise we set $\beta_{G_k} = \mathbf{0}_{m_k}$.

For the Gaussian family we are able to simplify the process by integrating out the variance τ^2 . Following [Murphy \(2007\)](#), substituting $\xi = \tau^2$ and recalling the independence

of τ^2 and β in our variational family, we have,

$$\begin{aligned} \tilde{p}(y^*|x^*, \mathcal{D}) &\propto \int_{\mathbb{R}^p} \int_{R^+} \xi^{-(a'+1/2)-1} \exp\left(-\frac{(y^* - x^{*\top}\beta)^2 + 2b'}{2\xi}\right) d\xi d\tilde{\Pi}(\beta) \\ &\propto \int_{\mathbb{R}^p} \left(\frac{(y^* - x^{*\top}\beta)^2}{2b'} + 1\right)^{-(a'+1/2)} d\tilde{\Pi}(\beta) \end{aligned}$$

Recognizing that the expression within the integral has the same functional form as a generalized t -distribution, whose density is denoted as $t(x; \mu, \sigma^2, \nu) = \Gamma((\nu+1)/2)/(\Gamma(\nu/2)\sqrt{\nu\pi}\sigma)(1+(x - \mu)^2/(\nu\sigma^2))^{-(\nu+1)/2}$, yields,

$$\tilde{p}(y^*|x^*, \mathcal{D}) = \int_{\mathbb{R}^p} t(y^*; x^{*\top}\beta, b'/a', 2a') d\tilde{\Pi}(\beta) \quad (28)$$

As (28) is intractable we instead sample from $p(y^*|x^*, \mathcal{D})$, by

1. Sampling $\beta \sim \tilde{\Pi}$.
2. Sampling y^* from $Y^* = \mu + \sigma t_\nu$ where $\mu = x^{*\top}\beta$, $\sigma = \sqrt{b'/a'}$ and $\nu = 2a'$, where t_ν denotes a t -distribution with ν degrees of freedom.

C Gibbs Sampler

We present a Gibbs sampler for the Gaussian family of models, noting that the samplers for the Binomial and Poisson family use the same principles. We begin by considering a slight alteration of the prior given in (2). Formally,

$$\begin{aligned} \beta_{G_k} &\stackrel{\text{iid}}{\sim} \Psi(\beta_{G_k}; \lambda) \\ z_k|\theta_k &\stackrel{\text{iid}}{\sim} \text{Bernoulli}(\theta_k) \\ \theta_k &\stackrel{\text{iid}}{\sim} \text{Beta}(a_0, b_0) \end{aligned} \quad (29)$$

for $k = 1, \dots, M$ and $\tau^2 = \xi \stackrel{\text{iid}}{\sim} \Gamma^{-1}(\xi; a, b)$. Under this prior writing the likelihood as,

$$p(\mathcal{D}|\beta, z, \xi) = \prod_{i=1}^n \phi\left(y_i; f\left(\sum_{k=1}^M z_k \langle x_{G_k}, \beta_{G_k} \rangle\right), \xi\right) \quad (30)$$

where $\phi(\cdot; \mu, \sigma^2)$ is the density of the Normal distribution with mean μ and variance σ^2 , and f is the link function (which in this case is the identity), yields the posterior

$$p(\beta, z, \xi | \mathcal{D}) \propto p(\mathcal{D} | \beta, z, \xi) \pi(\xi) \prod_{k=1}^M \pi(\beta_{G_k}) \pi(z_k | \theta_k) \pi(\theta_k). \quad (31)$$

Notably, the posterior is equivalent to our previous formulation.

To sample from (31), we construct a Gibbs sampler as outlined in Algorithm 2. Ignoring the superscript for clarity, the distribution $\theta_k | \mathcal{D}, \beta, z, \theta_{-k}, \xi$ is conditionally independent of $\mathcal{D}, \beta, z_{-k}, \xi$ and θ_{-k} . Therefore, θ_k is sampled from $\theta_k | z_k$, which has a Beta($a_0 + z_k, b_0 + 1 - z_k$) distribution. Regarding $z_j^{(i)}$, the conditional density

$$\begin{aligned} p(z_k | \mathcal{D}, \beta, z_{-k}, \theta, \xi) &\propto p(\mathcal{D} | \beta, z_{-k}, z_k, \theta, \xi) \pi(z_k | \beta, z_{-k}, \theta, \xi) \\ &= p(\mathcal{D}; \beta, z, \xi) \pi(z_k | \theta_k). \end{aligned} \quad (32)$$

As z_k is discrete, evaluating the RHS of (32) for $z_k = 0$ and $z_k = 1$, gives the unnormalised conditional probabilities. Summing gives the normalisation constant and thus we can sample z_k from a Bernoulli distribution with parameter

$$p_k = \frac{p(z_k = 1 | \mathcal{D}, \beta, z_{-k}, \theta, \xi)}{p(z_k = 0 | \mathcal{D}, \beta, z_{-k}, \theta, \xi) + p(z_k = 1 | \mathcal{D}, \beta, z_{-k}, \theta, \xi)}. \quad (33)$$

Finally, regarding $\beta_{G_{k,j}}^{(i)}$ we use a Metropolis-Hastings within Gibbs step, wherein a proposal $\beta_{G_{k,j}}^{(i)}$ is sampled from a random-walk proposition kernel K , where in our implementation $K(x | \vartheta, \varepsilon) = N(x; \vartheta, \sqrt{2} (10^{1-\varepsilon})^{1/2})$. The proposal is then accepted with probability A or rejected with probability $1 - A$, in which case $\beta_{G_{k,j}}^{(i)} \leftarrow \beta_{G_{k,j}}^{(i-1)}$, where A is given by,

$$A = \min \left(1, \frac{p(\mathcal{D}; \beta_{G_k^c}, \beta_{G_{k,-j}}, \beta_{G_{k,j}}^{(i)}, z, \xi) \pi(\beta_{G_{k,j}}^{(i)} | \beta_{G_{k,-j}})}{p(\mathcal{D}; \beta_{G_k^c}, \beta_{G_{k,-j}}, \beta_{G_{k,j}}^{(i-1)}, z, \xi) \pi(\beta_{G_{k,j}}^{(i-1)} | \beta_{G_{k,-j}})} \frac{K(\beta_j^{(i-1)} | \beta_j^{(i)}, z_k^{(i)})}{K(\beta_j^{(i)} | \beta_j^{(i-1)}, z_k^{(i-1)})} \right)$$

Algorithm 2 MCMC sampler for the Gaussian family GSpSL regression

Initialize $\beta^{(0)}, z^{(0)}, \theta^{(0)}, \xi^{(0)}$

for $i = 1, \dots, N$

for $k = 1, \dots, M$

$$\theta_k^{(i)} \stackrel{\text{iid.}}{\sim} \text{Beta}(a_0 + z_k^{(i-1)}, b_0 + 1 - z_k^{(i-1)})$$

for $k = 1, \dots, M$

$$z_k^{(i)} \stackrel{\text{ind.}}{\sim} \text{Bernoulli}(p_k)$$

for $k = 1, \dots, M$

for $j = 1, \dots, m_k$

$$\beta_{G_k,j}^{(i)} \sim p(\beta_{G_k,j}^{(i)} | \mathcal{D}, z^{(i)}, \beta_{G_{1:k-1}}^{(i)}, \beta_{G_{k,1:j-1}}^{(i)}, \beta_{G_{k,j+1:m_k}}^{(i-1)}, \beta_{G_{k+1:M}}^{(i-1)}, \xi^{(i-1)})$$

 Sample $\xi^{(i)} \stackrel{\text{iid.}}{\sim} \Gamma^{-1}(a + 0.5n, b + 0.5\|y - \sum_{k=1}^M z_k X_{G_k} \beta_{G_k}^{(i)}\|^2)$

return $\{\beta^{(i)}, z^{(i)}, \theta^{(i)}, \xi^{(i)}\}_{i=1}^N$.

D Numerical Study

D.1 Performance with large groups

The performance of GSVB with large groups is examined. For this we fix $(n, p, s) = (1,000, 5,000, 10)$ and vary the group size to be $m = 10, 20, 25, 50, 100$. The results, presented in Figure 3 highlight that the performance of the method is excellent for groups of size $m = 10, 20, 25$, and begins to suffer when groups are of size $m = 50$ and 100 , particularly in settings 3 and 4 wherein GSVB-B did not run.

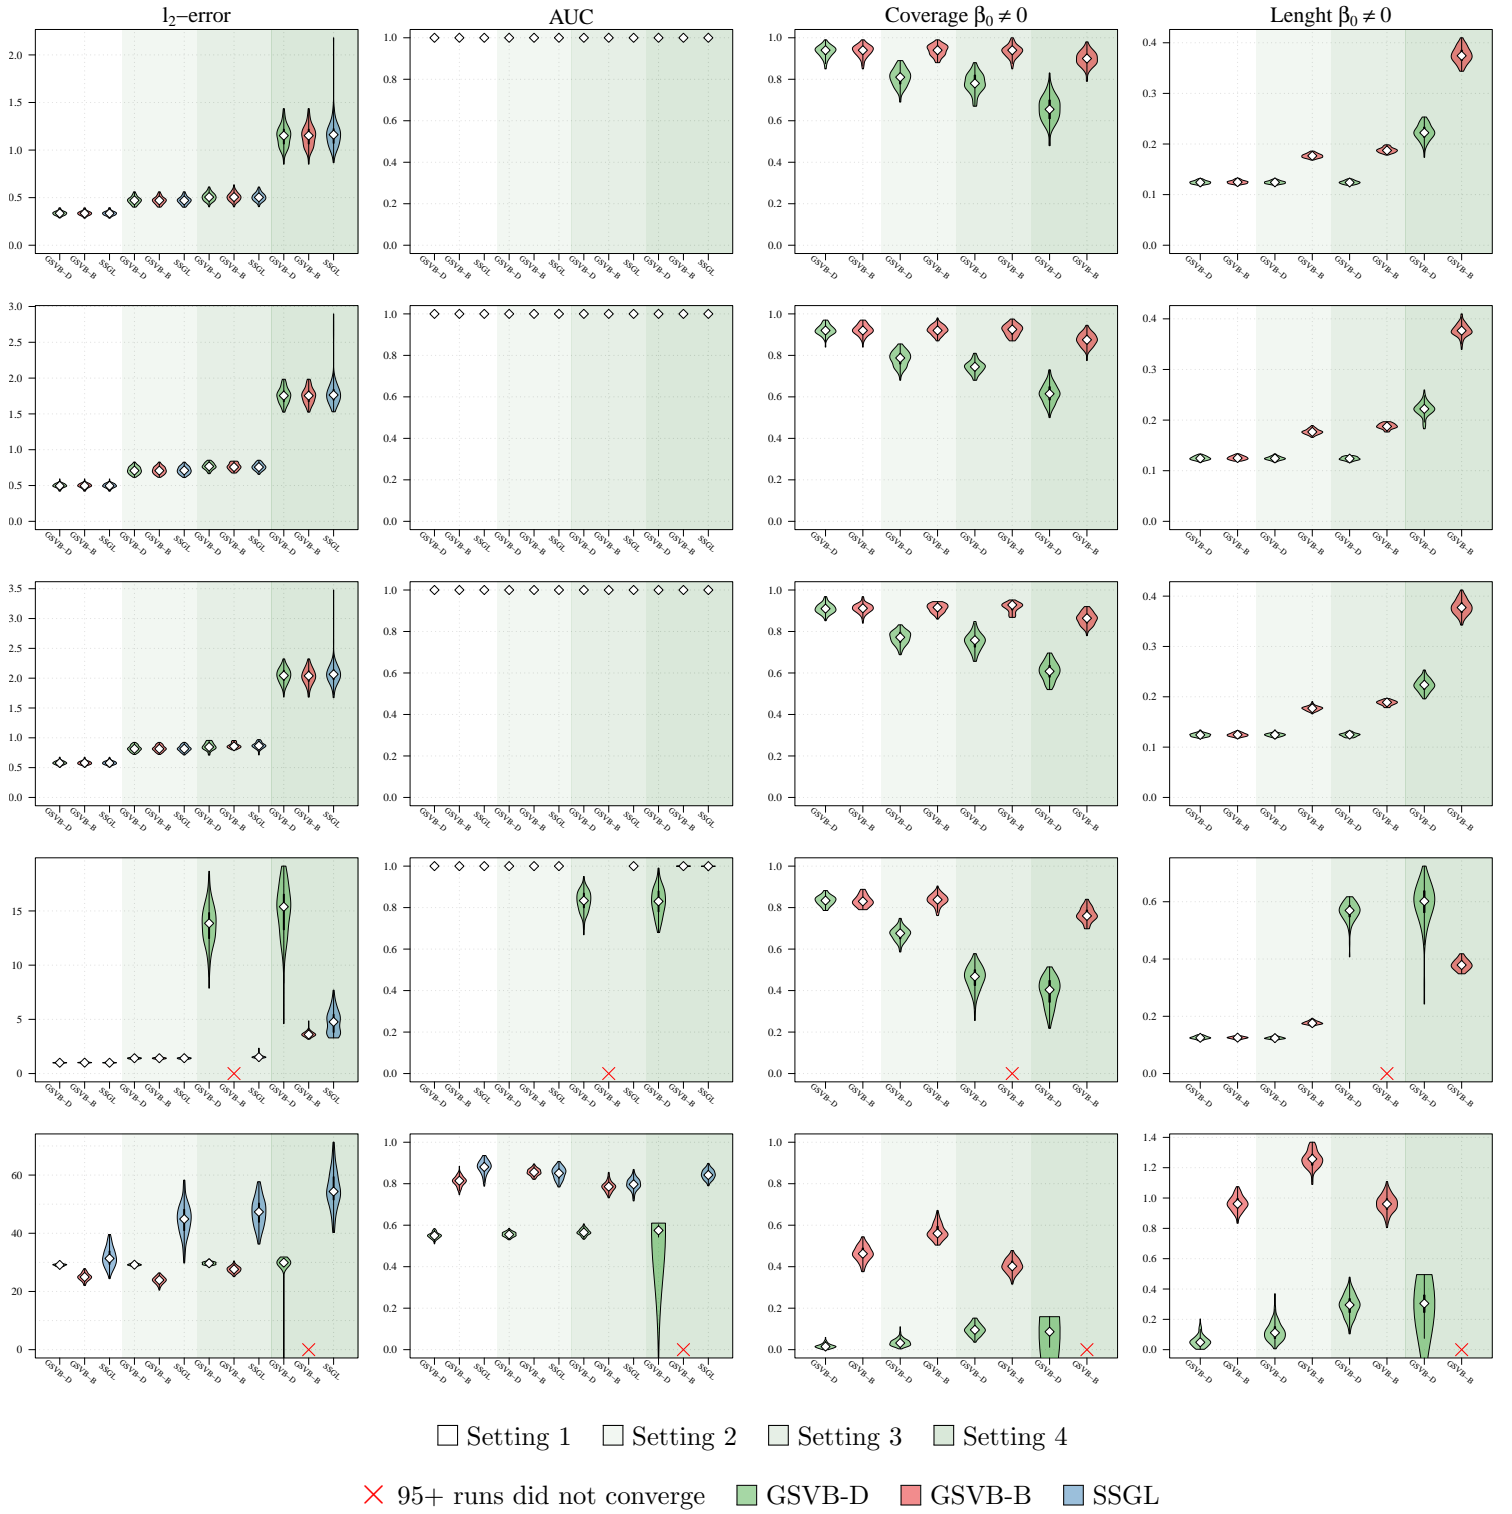


Figure 3: Performance evaluation of GSVB and SSGL for Settings 1–4 with $(n, p, s) = (1,000, 5,000, 10)$ across 100 runs. For each method the white diamond (\diamond) indicates the median of the metric, the thick black line (—) the interquartile range, and the black line (—) 1.5 times the interquartile range. **Rows 1–5:** Gaussian family with increasing group sizes of $m = 10, 20, 25, 50, 100$.

D.2 Runtime

Runtimes for the experiments presented in Section 5 are presented in Table 3 and Table 4. Notably, the runtime for GSVB over MCMC is orders of magnitudes faster. The runtime of GSVB in comparison to SSGL is slower in Settings 4, but marginally slower in the remaining settings. Note all simulations were ran on AMD EPYC 7742 CPUs (128 core, 1TB RAM).

	Method	Setting 1	Setting 2	Setting 3	Setting 4
Gaus. s=5	GSVB-D	2.1s (1.6s, 3.0s)	1.8s (1.1s, 2.9s)	3.5s (1.6s, 8.0s)	2.2s (1.6s, 5.2s)
	GSVB-B	1.8s (1.2s, 2.9s)	2.0s (1.3s, 3.1s)	3.4s (1.7s, 7.3s)	2.6s (1.8s, 6.7s)
	MCMC	7m 53s (7m 31s, 8m 12s)	7m 55s (7m 34s, 8m 2s)	7m 44s (7m 20s, 7m 48s)	8m 14s (7m 49s, 8m 17s)
Gaus. s=10	GSVB-D	2.3s (1.7s, 3.8s)	2.6s (1.9s, 3.9s)	7.4s (2.9s, 9.7s)	4.2s (2.4s, 12.1s)
	GSVB-B	2.9s (1.8s, 4.5s)	2.6s (2.2s, 3.2s)	9.5s (4.8s, 13.9s)	4.3s (2.7s, 14.0s)
	MCMC	9m 55s (9m 45s, 10m 8s)	10m 13s (9m 48s, 10m 30s)	9m 19s (9m 24s, 9m 33s)	10m 11s (9m 46s, 10m 27s)
Binom. s=3	GSVB-D	-	5.5s (3.5s, 9.3s)	7.7s (3.7s, 12.0s)	5.2s (2.7s, 8.0s)
	GSVB-B	-	7.2s (6.0s, 11.9s)	7.8s (5.6s, 14.4s)	6.0s (3.7s, 9.5s)
	MCMC	-	16m 45s (16m, 17m 15s)	16m 45s (15m 45s, 17m)	15m 45s (14m 49s, 16m 30s)
Pois. s=2	GSVB-D	3.6s (3.3s, 4.1s)	3.3s (2.9s, 3.7s)	3.1s (2.9s, 5.2s)	3.4s (2.4s, 4.1s)
	GSVB-B	12.5s (10.0s, 18.0s)	10.1s (7.9s, 14.5s)	15.5s (10.1s, 41.6s)	13.3s (10.5s, 24.0s)
	MCMC	15m 45s (14m 49s, 16m 45s)	15m 45s (15m 15s, 16m 30s)	16m (14m 59s, 18m)	18m 30s (16m 45s, 19m 45s)

Table 3: Median (5%, 95% quartile) runtimes for numerical experiments presented in Figure 1. Note $(n, p) = (200, 1000)$ for the Gaussian family and $(n, p) = (400, 1000)$ for the Binomial and Poisson family. Note that under setting 1 for the Binomial family there is perfect separation.

	Method	Setting 1	Setting 2	Setting 3	Setting 4
Gaus.	GSVB-D	59.1s (47.2s, 1m 16s)	54.3s (47.5s, 1m 8s)	2m 53s (1m 4s, 6m 35s)	8m 7s (4m 30s, 25m 36s)
m=10,	GSVB-B	57.7s (50.7s, 1m 18s)	57.0s (49.5s, 1m 14s)	2m 51s (1m 7s, 6m 38s)	9m 45s (4m 53s, 25m 40s)
s=10	SSGL	1m 3s (59.5s, 1m 13s)	50.8s (41.6s, 59.5s)	1m 1s (58.1s, 1m 16s)	1m 24s (1m 11s, 2m 3s)
Gaus.	GSVB-D	3m 7s (3m 46s, 4m 31s)	3m 13s (3m 42s, 4m 39s)	9m 42s (3m 56s, 14m 54s)	20m 48s (6m 38s, 39m 23s)
m=10,	GSVB-B	3m 29s (3m 6s, 4m 20s)	3m 23s (3m 57s, 4m 48s)	12m 18s (4m 59s, 16m 20s)	34m 59s (12m 54s, 45m 2s)
s=20	SSGL	2m 42s (1m 23s, 2m 3s)	1m 16s (59.2s, 2m 40s)	2m 6s (2m 37s, 2m 30s)	2m 16s (2m 31s, 6m 32s)
Binom.	GSVB-D	-	3m 43s (2m 54s, 4m 54s)	2m 26s (2m 34s, 4m 5s)	2m 50s (1m 16s, 3m 30s)
m=5,	GSVB-B	-	4m 32s (3m 40s, 5m 8s)	3m 26s (2m 20s, 5m 26s)	2m 25s (2m 47s, 5m 18s)
s=5	SSGL	-	2m 54s (2m 44s, 3m 42s)	3m 33s (2m 57s, 3m 26s)	2m 40s (2m 31s, 2m 3s)
Pois.	GSVB-D	16.3s (9.8s, 21.5s)	11.2s (8.2s, 20.8s)	16.0s (10.4s, 43.0s)	12.9s (8.6s, 48.7s)
m=5,	GSVB-B	2m 26s (2m 59s, 3m 20s)	2m 55s (1m 29s, 3m 20s)	3m 53s (2m 49s, 9m 25s)	3m 33s (2m 4s, 9m 11s)
s=3	SSGL	4m 11s (3m 49s, 8m 38s)	4m 34s (2m 11s, 13m 4s)	2m 29s (2m 56s, 16m 33s)	2m 32s (1m 25s, 3m 2s)

Table 4: Median (5%, 95% quartile) runtimes for numerical experiments presented in Figure 2. Note $(n, p) = (500, 5000)$ for the Gaussian family and $(n, p) = (1000, 5000)$ for the Binomial and Poisson family. Note that under setting 1 for the Binomial family there is perfect separation

E Real data analysis

E.1 Genetic determinants of Low-density Lipoprotein in mice

Here we analyze a dataset from the Mouse Genome Database (Blake et al., 2021) and consists of $p = 10,346$ single nucleotide polymorphisms (SNPs) collected from $n = 1,637$ laboratory mice. Notably, each SNP, x_{ij} for $j = 1, \dots, p$ and $i = 1, \dots, n$, takes a value of 0, 1 or 2. This value indicates how many copies of the risk allele are present. Alongside the genotype data, a phenotype (response) is also collected. The phenotype we consider is low-density lipoprotein cholesterol (LDL-C), which has been shown to be a major risk factor for conditions like coronary artery disease, heart attacks, and strokes (Silverman et al., 2016).

To pre-process the original dataset SNPs with a rare allele frequency, given by $\text{RAF}_j = \frac{1}{2n} \sum_{i=1}^n (\mathbb{I}(x_{ij} = 1) + 2\mathbb{I}(x_{ij} = 2))$ for $j = 1, \dots, p$, below the first quartile were discarded. Further, due to the high collinearity, covariates with $|\text{corr}(x_j, x_k)| \geq 0.97$ for $k > j$, $j = 1, \dots, p - 1$, were removed. After pre-processing, 3,341 SNPs remained. These were used to construct groups by coding each $x_{ij} : \{0, 1, 2\} \mapsto \{(0, 0), (0, 1), (1, 1)\}$, in turn giving groups of size $m = 2$.

To evaluate the methods ten fold cross validation is used. Specifically, methods are fit to the validation set and the test set is used for evaluation. This is done by computing the: (i) MSE between the true and predicted value of the response, given by $\frac{1}{\tilde{n}} \|y - X\hat{\beta}\|^2$ where \tilde{n} is the size of the test set, (ii) posterior predictive (PP) coverage, which measures the proportion of times the response is included in the 95% PP set, and (iii) the average size of the 95% PP set. In addition, (iv) the number of groups selected, given by $\sum_{k=1}^M \mathbb{I}(\gamma_k > 0.5)$ is also reported.

Importantly the PP distribution is available for GSVB, and given by, $\tilde{p}(y^*|x^*, \mathcal{D}) = \int_{\mathbb{R}^p \times \mathbb{R}^+} p(y^*|\beta, \tau^2, x^*, \mathcal{D}) d\tilde{\Pi}(\beta, \tau^2)$ where $x^* \in \mathbb{R}^p$ is a feature vector (see Section B of the Supplementary material for details). In genetics studies, $x^{*\top}\beta$, is referred to as the polygenic risk score and is commonly used to evaluate genetic risk. While frequentist or MAP methods offer only point estimates, Figure 4 demonstrates that GSVB can offer a distribution for this score, highlighting the uncertainty associated with the estimate.

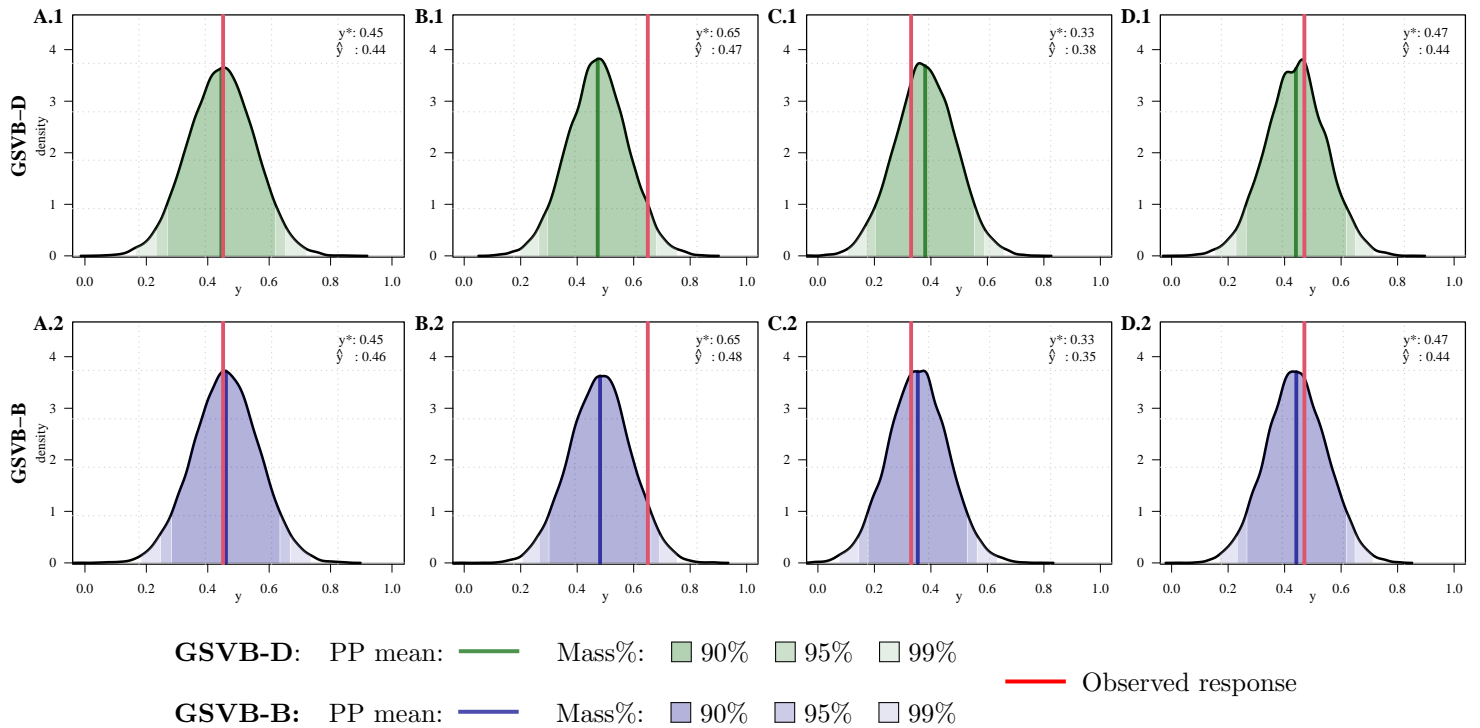


Figure 4: *Genetic determinants of LDL-C in mice*, variational posterior predictive distribution for GSVB-D (panels A.1 - D.1) and GSVB-B (panels A.2 - D.2) constructed for four observations in the held out test set (fold 1). Notably the red line (—) represents the measured LDL-C level, with the adjacent text giving this value (y^*) and the posterior mean (\hat{y}). The shading of the variational posterior indicates where 90%, 95% and 99% of the mass is contained.

Regarding the performance of each method results are presented in Table 5, these high-

light that GSVB performs excellently obtaining parsimonious models with a comparable MSE to SSGL. In addition, GSVB provides impressive uncertainty quantification, yielding a coverage of 0.945 and 0.941 for GSVB-D and GSVB-B respectively. The SNPs selected by GSVB-D and GSVB-B are similar and reported in Table 6 of the Supplementary material. When the methods are fit to the full dataset, the SNPs selected by both methods are: rs13476241, rs13477968, and rs13483814, with GSVB-B selecting rs13477939 in addition. Notably, rs13477968 corresponds to the *Ago3* gene, which has been shown to play a role in activating LDL receptors, which in turn regulate LDL-C (Matsui et al., 2010).

Method	MSE	Num. selected groups	PP. coverage	PP. length
GSVB-D	0.012 (0.002)	3.50 (0.527)	0.945 (0.027)	0.425 (0.005)
GSVB-B	0.012 (0.002)	4.00 (0.667)	0.941 (0.029)	0.421 (0.005)
SSGL	0.012 (0.002)	57.40 (5.481)	-	-

Table 5: *Genetic determinants of LDL-C in mice*, evaluation of methods on the held out genotype data. Reported is the 10 fold cross validated MSE, the number of selected groups, the posterior predictive coverage, and mean posterior predictive interval length.

RSID (or probe ID)	Freq. GSVB-D	Freq. GSVB-B
rs13476279	6	6
rs13483823	6	5
CEL.4_130248229	5	3
rs13477968	3	5
rs13476241	4	3
CEL.X_65891570	2	4
rs13477903	2	2
rs3688710	2	1
rs13477939	1	-
rs3157124	1	-
rs13478204	-	3
rs6186902	-	2
rs13483814	-	1
rs4222922	-	1
rs6187266	-	1
UT_4_128.521481	-	1

Table 6: *Genetic determinants of LDL-C in mice: SNPs selected by GSVB.*

E.2 Bardet-Biedl Syndrome Gene Expression Study

Probe ID	Gene Name	Freq. GSVB-D	Freq. GSVB-B	Freq. SSGL
1383183_at		1	1	8
1386237_at		1	1	-
1397359_at		4	2	-
1386552_at		-	1	-
1378682_at		-	1	-
1385926_at		-	-	2
1396814_at		-	-	3
1386811_at	<i>Hacd4</i>	1	2	-
1386069_at	<i>Sp2</i>	4	2	-
1371209_at	<i>RT1-CE5</i>	-	1	-
1391624_at	<i>Sec14l3</i>	-	1	-
1383829_at	<i>Bbx</i>	-	2	-
1389274_at	<i>Dcakd</i>	-	-	1
1373995_at	<i>Abcg1</i>	-	-	1
1375361_at	<i>Snx18</i>	-	-	1
1368915_at	<i>Kmo</i>	-	-	7

Table 7: Bardet-Biedl syndrome gene expression study: variables selected by GSVB and SSGL.

E.3 MEMset splice site detection

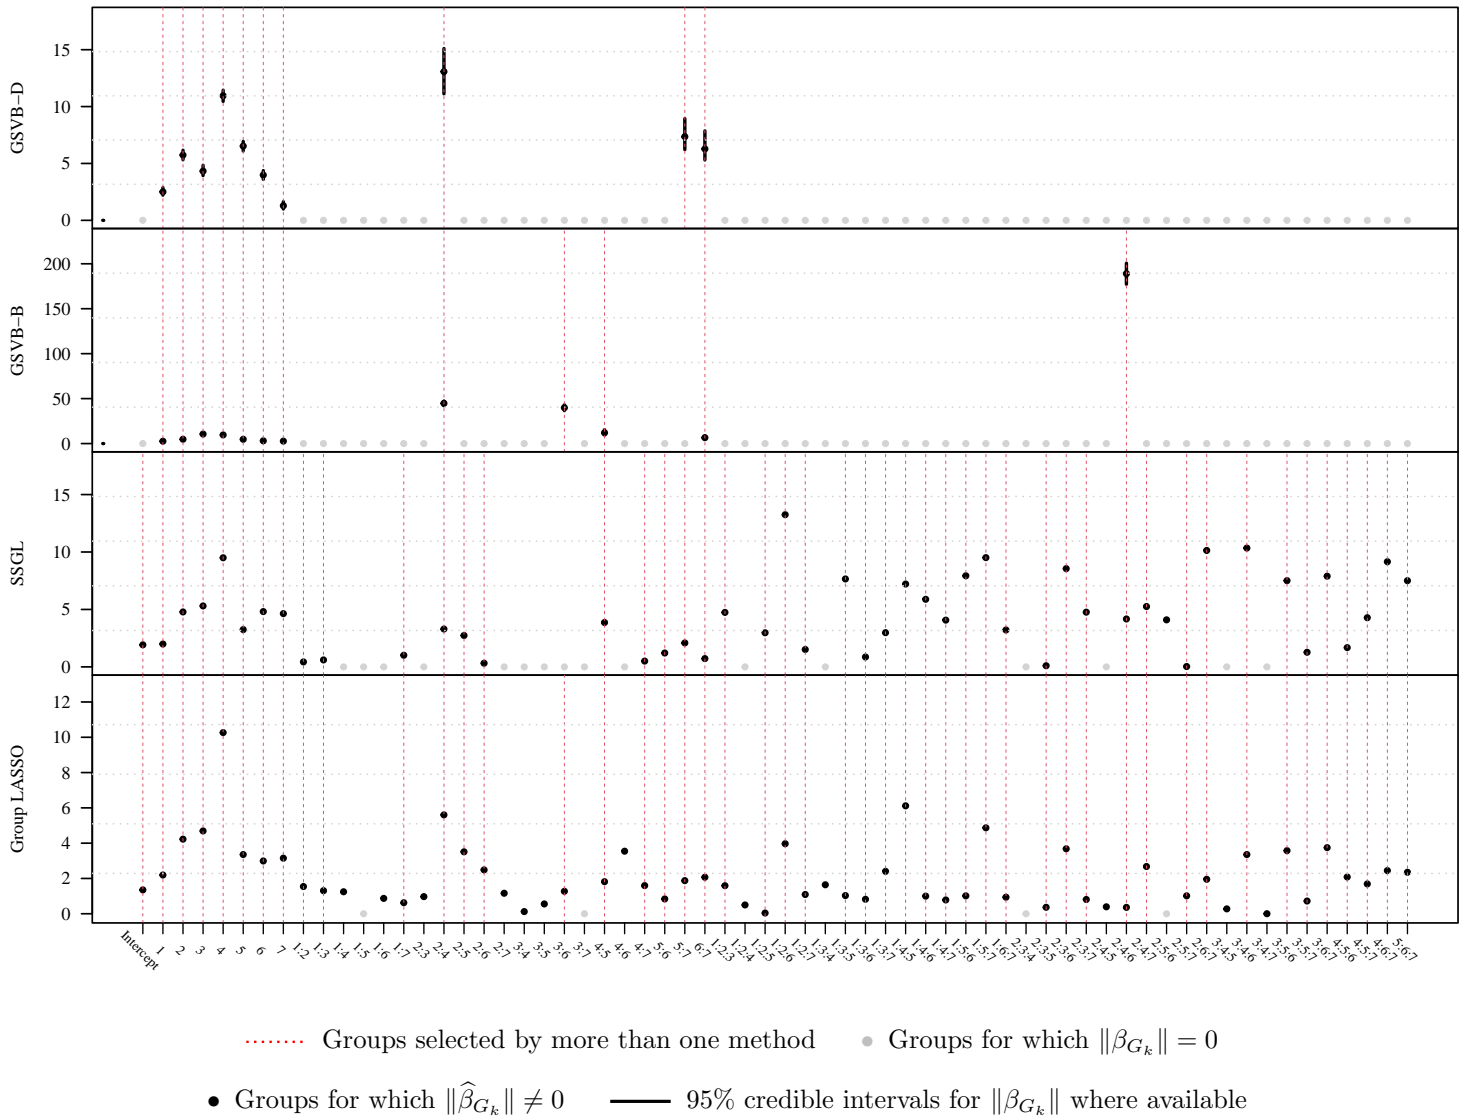


Figure 5: *MEMset splice site detection*, comparison of ℓ_2 -norms for each group of coefficients $\hat{\beta}_{G_k}$ for $k = 1, \dots, 64$. The red dotted line (.....) shows groups which have been selected by more than one method. The points indicate $\|\hat{\beta}_{G_k}\|$ where $\hat{\beta}_{G_k}$ is the posterior mean for GSVB, the MAP for SSGl and the MLE estimate under the group LASSO. The points are coloured in black (●) when they are non-zero and grey (●) otherwise. Finally when available the 95% credible set for $\|\beta_{G_k}\|$ is given by the solid black line (—).

F Proofs of asymptotic results

F.1 A general class of model selection priors and an overview of the proof

Our theoretical results apply to a wider class of model selection priors than the group spike and slab prior (2) underlying our variational approximation. In this section, we thus consider this more general class of model selection priors [Castillo et al. \(2015\)](#); [Ning et al. \(2020\)](#); [Ray and Szabó \(2022\)](#) defined hierarchically via:

$$\begin{aligned}
 s &\sim \pi_M(s) \\
 S||S| = s &\sim \text{Unif}_{M,s} \\
 \beta_{G_k} &\sim_{\text{ind}} \begin{cases} \Psi_{\lambda, m_k}(\beta_{G_k}), & G_k \in S, \\ \delta_0, & G_k \notin S, \end{cases}
 \end{aligned} \tag{34}$$

where π_M is a prior on $\{0, 1, \dots, M\}$, $\text{Unif}_{M,s}$ is the uniform distribution on subsets $S \subseteq \{1, \dots, M\}$ of size s and

$$\Psi_{\lambda, m_k}(\beta_{G_k}) = \Delta_{m_k} \lambda^{m_k} \exp(-\lambda \|\beta_{G_k}\|_2) \tag{35}$$

is a density on \mathbb{R}^{m_k} with $\Delta_{m_k} = 2^{-m_k} \pi^{(1-m_k)/2} \Gamma((m_k + 1)/2)^{-1}$. This can be concisely written as

$$(S, \beta) \mapsto \pi_M(|S|) \frac{1}{\binom{M}{|S|}} \delta_0(S^c) \prod_{G_k \in S} \Psi_{\lambda, m_k}(\beta_{G_k}). \tag{36}$$

Following [Castillo et al. \(2015\)](#); [Ning et al. \(2020\)](#); [Ray and Szabó \(2022\)](#), we assume as usual that

$$A_1 M^{-A_3} \pi_M(s-1) \leq \pi_M(s) \leq A_2 M^{-A_4} \pi_M(s-1), \quad s = 1, \dots, M. \tag{37}$$

We further recall the assumption (38) made on the scale parameter:

$$\underline{\lambda} \leq \lambda \leq 2\bar{\lambda}, \quad \underline{\lambda} = \frac{\|X\|}{M^{1/m_{\max}}}, \quad \bar{\lambda} = 3\|X\|\sqrt{\log M}. \quad (38)$$

The group spike and slab prior fits within this framework by taking $\pi_M = \text{Bin}(M, \frac{a_0}{a_0+b_0})$, and hence the following results immediately imply Theorems 1 and 2. Recall the parameter set

$$\mathcal{B}_{\rho_n, s_n} = \{\beta_0 \in \mathbb{R}^p : \phi(S_{\beta_0}) \geq c_0, |S_{\beta_0}| \leq s_n, \tilde{\phi}(\rho_n | S_{\beta_0}|) \geq c_0\}$$

defined in (23).

Theorem 3 (Contraction). *Suppose that Assumption (K) holds, the prior satisfies (37)-(38) and s_n satisfies $m_{\max} \log s_n \leq K' \log M$ for some $K' > 0$. Then the variational posterior $\tilde{\Pi}$ based on either the variational family \mathcal{Q} in (5) or \mathcal{Q}' in (6) satisfies, with $s_0 = |S_{\beta_0}|$,*

$$\begin{aligned} \sup_{\beta_0 \in \mathcal{B}_{\rho_n, s_n}} E_{\beta_0} \tilde{\Pi} \left(\beta : \|X(\beta - \beta_0)\|_2 \geq \frac{H_0 \rho_n^{1/2} \sqrt{s_0 \log M}}{\bar{\phi}(\rho_n s_0)} \right) &\rightarrow 0 \\ \sup_{\beta_0 \in \mathcal{B}_{\rho_n, s_n}} E_{\beta_0} \tilde{\Pi} \left(\beta : \|\beta - \beta_0\|_{2,1} \geq \frac{H_0 \rho_n s_0 \sqrt{\log M}}{\|X\| \bar{\phi}(\rho_n s_0)^2} \right) &\rightarrow 0, \\ \sup_{\beta_0 \in \mathcal{B}_{\rho_n, s_n}} E_{\beta_0} \tilde{\Pi} \left(\beta : \|\beta - \beta_0\|_2 \geq \frac{H_0 \rho_n^{1/2} \sqrt{s_0 \log M}}{\|X\| \tilde{\phi}(\rho_n s_0)^2} \right) &\rightarrow 0, \end{aligned}$$

for any $\rho_n \rightarrow \infty$ (arbitrarily slowly), $\mathcal{B}_{\rho_n, s_n}$ defined in (23) and where H_0 depends only on the prior.

Theorem 4 (Dimension). *Suppose that Assumption (K) holds, the prior satisfies (37)-(38) and s_n satisfies $m_{\max} \log s_n \leq K' \log M$ for some $K' > 0$. Then the variational posterior $\tilde{\Pi}$ based on either the variational family \mathcal{Q} in (5) or \mathcal{Q}' in (6) satisfies*

$$\sup_{\beta_0 \in \mathcal{B}_{\rho_n, s_n}} E_{\beta_0} \tilde{\Pi} (\beta : |S_\beta| \geq \rho_n |S_{\beta_0}|) \rightarrow 0$$

for any $\rho_n \rightarrow \infty$ (arbitrarily slowly) and $\mathcal{B}_{\rho_n, s_n}$ defined in (23).

To prove Theorems 3 and 4, we use the following result which relates the VB probability of sets having exponentially small probability under the true posterior.

Lemma 1 (Theorem 5 of Ray and Szabó (2022)). *Let B_n be a subset of the parameter space, A_n be events and Q_n be distributions for β . If there exists $C > 0$ and $\delta_n > 0$ such that*

$$E_{\beta_0} \Pi(\beta \in B_n^c | Y) 1_{A_n} \leq C e^{-\delta_n},$$

then

$$E_{\beta_0} Q_n(\beta \in B_n^c) 1_{A_n} \leq \frac{2}{\delta_n} [E_{\beta_0} D_{KL}(Q_n \| \Pi(\cdot | Y)) 1_{A_n} + C e^{-\delta/2}].$$

The proof thus reduces to finding events A_n with $P_{\beta_0}(A_n) \rightarrow 1$ on which:

1. The true posterior places only exponentially small probability outside B_n , that is $\Pi(B_n^c | Y) \leq C e^{-\delta_n}$ for some rate $\delta_n \rightarrow \infty$,
2. The D_{KL} -divergence between the VB posterior $\tilde{\Pi}$ and the full posterior is $o(\delta_n)$.

In our setting, we shall take $\delta_n = C s_0 \log M$. Full posterior results are dealt with in Section F.2, the D_{KL} -divergence in Section F.3 and the proofs of Theorems 3 and 4 are completed in Section F.4.

F.2 Asymptotic theory for the full posterior

We now establish contraction rates for the full computationally expensive posterior distribution, keeping track of the exponential tail bounds needed to apply Lemma 1. While the proofs in this section largely follow those in Castillo et al. Castillo et al. (2015), the precise arguments adapting these results to the group sparse setting are rather technical and hence we provide them for convenience. We first establish a Gaussian tail bound in terms of the group structure.

Lemma 2. *The event*

$$\mathcal{T}_0 = \left\{ \max_{k=1, \dots, M} \|X_{G_k}^T (Y - X\beta_0)\|_2 \leq 3\|X\|\sqrt{\log M} \right\} \quad (39)$$

satisfies $\sup_{\beta_0 \in \mathbb{R}^p} P_{\beta_0}(\mathcal{T}_0^c) \leq 2/M$.

Proof. Since $Y - X\beta_0 = \varepsilon \sim N_n(0, I_n)$ under P_{β_0} , applying a union bound over the group structure yields $P_{\beta_0}(\max_k \|X_{G_k}^T (Y - X\beta_0)\|_2 > t) \leq \sum_{k=1}^M P_{\beta_0}(\|X_{G_k}^T \varepsilon\|_2 > t)$. Recall that for a multivariate normal $W \sim N_m(0, \Sigma)$, we have $P(\|W\|_2 - E\|W\|_2 \geq x) \leq 2 \exp(-x^2/(2E\|W\|_2^2))$ by Corollary 3 of [Pinelis and Sakhnenko \(1986\)](#). Since $X_{G_k}^T \varepsilon \sim N_{m_k}(0, X_{G_k}^T X_{G_k})$, this implies

$$(E\|X_{G_k}^T \varepsilon\|_2)^2 \leq E\|X_{G_k}^T \varepsilon\|_2^2 = \text{Tr}(X_{G_k}^T X_{G_k}) \leq \|X\|^2,$$

and hence $P(\|X_{G_k}^T \varepsilon\|_2 \geq \|X\| + x) \leq 2 \exp(-x^2/(2\|X\|^2))$ for any $x > 0$. Substituting this bound with $x = 2\|X\|\sqrt{\log M}$ into the above union bound thus yields

$$P_{\beta_0}(\max_k \|X_{G_k}^T (Y - X\beta_0)\|_\infty > 3\|X\|\sqrt{\log M}) \leq 2Me^{-2\log M} = 2M^{-1}.$$

□

Let $\ell_n(\beta)$ denote the log-likelihood of the $N_n(X\beta, I_n)$ -distribution. For any $\beta \in \mathbb{R}^p$, we have log-likelihood ratio

$$\Lambda_\beta(Y) = e^{\ell_n(\beta) - \ell_n(\beta_0)} = e^{-\frac{1}{2}\|X(\beta - \beta_0)\|_2^2 + (Y - X\beta_0)^T X(\beta - \beta_0)}. \quad (40)$$

The next result establishes an almost sure lower bound on the denominator of the Bayes formula. It follows Lemma 2 of [Castillo et al. \(2015\)](#), but must be adapted to account for the uneven prior normalizing factors coming from the group structure.

Lemma 3. *Suppose Assumption (K) holds and that $\beta_0 \in \mathbb{R}^p$ satisfies $m_{\max} \log s_0 \leq K' \log M$ for $K' > 0$. Then it holds that with P_{β_0} -probability one,*

$$\int \Lambda_{\beta}(Y) d\Pi(\beta) \geq C \pi_m(s_0) s_0! e^{-\lambda \|\beta_0\|_{2,1}} e^{-cs_0 \log M},$$

where $C, c > 0$ depend only on K, K' .

The mild condition $m_{\max} \log s_0 \leq K' \log M$, which will be assumed throughout, relates the true sparsity with the maximal group size. As in Assumption (K), the constant K' is used to provide uniformity, but this can be ignored at first reading.

Proof. The bound trivially holds true for $s_0 = 0$, hence we assume $s_0 \geq 1$. Write $p_0 := \sum_{G_k \in S_0} m_k$ to be maximal number of non-zero coefficients in β_0 . In an abuse of notation, we shall sometimes interchangeably use β_{S_0} for both the vector in \mathbb{R}^{p_0} and the vector in \mathbb{R}^p with the entries in S_0^c set to zero. Using the form (36) of the prior,

$$\int \Lambda_{\beta}(Y) d\Pi(\beta) \geq \frac{\pi_M(s_0)}{\binom{M}{s_0}} \int_{\mathbb{R}^{p_0}} \Lambda_{\beta}(Y) \prod_{G_k \in S_0} \Psi_{\lambda, m_k}(\beta_{G_k}) d\beta_{G_k}.$$

By the change of variable $b_{S_0} = \beta_{S_0} - \beta_{0, S_0}$ and the form of the log-likelihood (40), the last display is lower bounded by

$$\frac{\pi_M(s_0)}{\binom{M}{s_0}} e^{-\lambda \|\beta_0\|_{2,1}} \int_{\mathbb{R}^{p_0}} e^{-\frac{1}{2} \|X_{S_0} b_{S_0}\|_2^2 + (Y - X\beta_0)^T X_{S_0} b_{S_0}} \prod_{G_k \in S_0} \Psi_{\lambda, m_k}(b_{G_k}) db_{G_k}.$$

Define the measure μ on \mathbb{R}^{p_0} by $d\mu(b_{S_0}) = e^{-\frac{1}{2} \|X_{S_0} b_{S_0}\|_2^2} \prod_{G_k \in S_0} \Psi_{\lambda, m_k}(b_{G_k}) db_{G_k}$. Let $\bar{\mu} = \mu / \mu(\mathbb{R}^{p_0})$ denote the normalized probability measure with corresponding expectation $E_{\bar{\mu}}$.

Defining $Z(b_{S_0}) = (Y - X\beta_0)^T X_{S_0} b_{S_0}$, Jensen's inequality implies $E_{\bar{\mu}} e^Z \geq e^{E_{\bar{\mu}} Z} = 1$, since $E_{\bar{\mu}} Z = 0$ as $\bar{\mu}$ is a symmetric probability distribution about zero. Thus the last display is lower bounded by $\pi_m(s_0) e^{-\lambda \|\beta_0\|_{2,1}} \mu(\mathbb{R}^{p_0}) / \binom{M}{s_0}$, which equals

$$\frac{\pi_M(s_0)}{\binom{M}{s_0}} e^{-\lambda \|\beta_0\|_{2,1}} \int_{\mathbb{R}^{p_0}} e^{-\frac{1}{2} \|X_{S_0} b_{S_0}\|_2^2} \prod_{G_k \in S_0} \Psi_{\lambda, m_k}(b_{G_k}) db_{G_k}. \quad (41)$$

Using the group structure, $\|Xb\|_2 = \|\sum_{k=1}^M X_{\cdot G_k} b_{G_k}\|_2 \leq \sum_{k=1}^M \|X_{\cdot G_k}\|_2 \|b_{G_k}\|_2 \leq \|X\| \|b\|_1$ since $\|b\|_2 \leq \|b\|_1$. Using the form (35) of the density Ψ_{λ, m_k} and recalling that

$$\int_{\|\beta_S\|_1 \leq r} (\lambda/2)^{|S|} e^{-\lambda \|\beta_S\|_1} d\beta_S \geq e^{-\lambda r} (\lambda r)^{|S|} / |S|!$$

by (6.2) in Castillo et al. (2015), the integral in the last display is bounded below by

$$\begin{aligned} e^{-1/2} \int_{\|X\| \|b_{S_0}\|_1 \leq 1} \prod_{G_k \in S_0} \Delta_{m_k} \lambda^{m_k} e^{-\lambda \|\beta_{G_k}\|_1} d\beta_{G_k} \\ \geq e^{-1/2} e^{-\lambda \|X\|} \frac{\lambda^{p_0}}{\|X\|^{p_0} p_0!} \prod_{G_k \in S_0} \Delta_{m_k} 2^{m_k}. \end{aligned} \quad (42)$$

Deviating from Castillo et al. (2015), we must now take careful account of the normalizing constants Δ_{m_k} .

Recall the form of the normalized constants $\Delta_{m_k} = 2^{-m_k} \pi^{(1-m_k)/2} \Gamma((m_k+1)/2)^{-1}$. The non-asymptotic upper bound in Stirling's approximation for the Gamma function gives for $z \geq 2$: $\Gamma(z) \leq \sqrt{2\pi(z-1)} \left(\frac{z-1}{e}\right)^{z-1} e^{\frac{1}{12(z-1)}}$. Taking $z = (m_k+1)/2 \geq 2$ for $m_k \geq 3$,

$$\Gamma((m_k+1)/2) \leq 2\sqrt{\pi(m_k-1)} \left(\frac{m_k-1}{2e}\right)^{(m_k-1)/2} \leq 2e^{-1/2} \sqrt{\pi} m_k^{m_k/2} \frac{1}{(2e)^{(m_k-1)/2}}, \quad (43)$$

where we have used that $\left(\frac{m_k-1}{m_k}\right)^{m_k/2} \leq \lim_{x \rightarrow \infty} (1 - \frac{1}{x})^{x/2} = e^{-1/2}$ since the function in the limit is strictly increasing on $(1, \infty)$. One can directly verify that the upper bound (43) also holds for $m_k = 1, 2$. Using (43),

$$\prod_{G_k \in S_0} \Delta_{m_k} 2^{m_k} = \prod_{G_k \in S_0} \frac{\pi^{(1-m_k)/2}}{\Gamma((m_k+1)/2)} \geq \frac{1}{(\sqrt{2}e^{-1})^{s_0} (2\pi e)^{p_0/2}} \prod_{G_k \in S_0} m_k^{-m_k/2} \geq c^{p_0} m_{\max}^{-p_0/2}$$

for some universal constant $c > 0$. Using this last display, we lower bound (42) by a constant multiple of

$$e^{-\lambda \|X\|} \left(\frac{\lambda}{\|X\|}\right)^{p_0} \frac{1}{p_0!} c^{p_0} m_{\max}^{-p_0/2}.$$

Using (38), if $\lambda/\|X\| \leq 1/2$, then $e^{-\lambda/\|X\|} (\lambda/\|X\|)^{p_0} \geq e^{-1/2} M^{-p_0/m_{\max}} \geq e^{-1/2} e^{-s_0 \log M}$, while if $\lambda/\|X\| \geq 1/2$, then $e^{-\lambda/\|X\|} (\lambda/\|X\|)^{p_0} \geq e^{-6\sqrt{\log M}} 2^{-p_0} \geq e^{-C(K)s_0 \log M}$ since

$m_{\max} \leq K \log M$ by assumption. Since also $m_{\max}^{-p_0/2}/p_0! \geq e^{-2p_0 \log p_0} \geq e^{-2m_{\max}s_0 \log(m_{\max}s_0)}$, the last display is lower bounded by $e^{-C_{s_0} \log M}$ under the lemma's hypotheses. The result then follows by substituting this lower bound for the integral in (41) and using that $\binom{M}{s_0} \leq M^{s_0}/s_0! = e^{s_0 \log M}/s_0!$. \square

The next result follows Theorem 10 of [Castillo et al. \(2015\)](#).

Lemma 4 (Dimension). *Suppose that Assumption (K) holds and the prior satisfies (37) and (38). Further assume $M > 0$ is large enough that $\log M \geq \max\{\frac{4m_{\max} \log 4}{A_2}, \frac{4 \log A_2}{A_2}\}$ and $M \geq (4^{m_{\max}+1/2} A_2)^{1/A_4}$. Then for any $\beta_0 \in \mathbb{R}^p$ such that $m_{\max} \log s_0 \leq K' \log M$ and any $L > 0$,*

$$\begin{aligned} E_{\beta_0} \Pi(\beta : |S_\beta| \geq (L+1)s_0 | Y) 1_{\mathcal{T}_0} \\ \leq C(K, K') \exp \left\{ \left(c(K, K', A_2) + \frac{144}{\phi(S_0)^2} - \frac{LA_4}{2} \right) s_0 \log M \right\}, \end{aligned}$$

where $s_0 = |S_{\beta_0}|$ and \mathcal{T}_0 is the event (39).

Proof. Using Bayes formula with likelihood ratio (40) and Lemma 3, for any measurable set $B \subset \mathbb{R}^p$,

$$\begin{aligned} \Pi(B|Y)1_{\mathcal{T}_0} &= 1_{\mathcal{T}_0} \frac{\int_B \Lambda_\beta(Y) d\Pi(\beta)}{\int \Lambda_\beta(Y) d\Pi(\beta)} \\ &\leq C \frac{e^{\lambda \|\beta_0\|_{2,1} + c_{s_0} \log M}}{s_0! \pi_M(s_0)} \int_B 1_{\mathcal{T}_0} e^{-\frac{1}{2} \|X(\beta - \beta_0)\|_2^2 + (Y - X\beta_0)^T X(\beta - \beta_0)} d\Pi(\beta) \end{aligned} \quad (44)$$

for \mathcal{T}_0 the event defined in (39). Applying Cauchy-Schwarz on the event \mathcal{T}_0 gives

$$|(Y - X\beta_0)^T X(\beta - \beta_0)| \leq \max_{k=1, \dots, M} \|X_{G_k}^T (Y - X\beta_0)\|_2 \sum_{k=1}^M \|\beta_{G_k} - \beta_{0, G_k}\|_2 \leq \bar{\lambda} \|\beta - \beta_0\|_{2,1}. \quad (45)$$

Therefore, since $(Y - X\beta)^T X(\beta - \beta_0) \sim N(0, \|X(\beta - \beta_0)\|_2^2)$ under P_{β_0} , on \mathcal{T}_0 , the integrand in the second last display is bounded by

$$e^{-\frac{1}{2} \|X(\beta - \beta_0)\|_2^2} E_{\beta_0} [1_{\mathcal{T}_0} e^{(1 - \frac{\lambda}{2\bar{\lambda}})(Y - X\beta_0)^T X(\beta - \beta_0)}] e^{\frac{\lambda}{2} \|\beta - \beta_0\|_{2,1}} \leq e^{-\frac{1}{2} [1 - (1 - \lambda/(2\bar{\lambda}))^2] \|X(\beta - \beta_0)\|_2^2} e^{\frac{\lambda}{2} \|\beta - \beta_0\|_{2,1}}$$

and hence

$$E_{\beta_0} \Pi(B|Y) 1_{\mathcal{T}_0} \leq C \frac{e^{\lambda \|\beta_0\|_{2,1} + cs_0 \log M}}{s_0! \pi_M(s_0)} \int_B e^{-\frac{1}{2}[1 - (1 - \lambda/(2\bar{\lambda}))^2] \|X(\beta - \beta_0)\|_2^2} e^{\frac{\lambda}{2} \|\beta - \beta_0\|_{2,1}} d\Pi(\beta).$$

Arguing exactly as on p. 2007-8 in [Castillo et al. \(2015\)](#),

$$\|\beta_0\|_{2,1} + \frac{1}{2} \|\beta - \beta_0\|_{2,1} \leq \frac{1}{8\bar{\lambda}} \|X(\beta - \beta_0)\|_2^2 + \frac{8s_0\bar{\lambda}}{\|X\|^2 \phi(S_0)^2} - \frac{1}{4} \|\beta - \beta_0\|_{2,1} + \|\beta\|_{2,1}$$

and hence

$$E_{\beta_0} \Pi(B|Y) 1_{\mathcal{T}_0} \leq C \frac{e^{cs_0 \log M}}{s_0! \pi_M(s_0)} e^{\frac{8s_0\bar{\lambda}}{\|X\|^2 \phi(S_0)^2}} \int_B e^{-\frac{\lambda}{4} \|\beta - \beta_0\|_{2,1} + \lambda \|\beta\|_{2,1}} d\Pi(\beta). \quad (46)$$

Setting now $B = \{\beta : |S_\beta| > R\}$ with $R \geq s_0$, the integral in the last display is bounded

by

$$\sum_{S: |S| > R} \frac{\pi_M(|S|)}{\binom{M}{|S|}} \int e^{-\frac{\lambda}{4} \|\beta - \beta_0\|_{2,1}} \prod_{G_k \in S} \Delta_{m_k} \lambda^{m_k} d\beta_{G_k} = \sum_{S: |S| > R} \frac{\pi_M(|S|)}{\binom{M}{|S|}} \prod_{G_k \in S} 4^{m_k}.$$

Using the prior condition (37), this is then bounded by

$$\sum_{s=R+1}^M \pi_M(s) 4^{m_{\max} s} \leq \pi_M(s_0) 4^{m_{\max} s_0} \left(\frac{4^{m_{\max} A_2}}{M^{A_4}} \right)^{R+1-s_0} \sum_{j=0}^{\infty} \left(\frac{4^{m_{\max} A_2}}{M^{A_4}} \right)^j.$$

Since $4^{m_{\max} A_2}/M^{A_4} \leq 1/2$ by the lemma hypothesis, the last sum is bounded by 2 and

hence the first term in (46) is bounded by

$$2C \exp \left\{ cs_0 \log M + \frac{8s_0\bar{\lambda}}{\|X\|^2 \phi(S_0)^2} + m_{\max} s_0 \log 4 + (R+1-s_0) \log(4^{m_{\max} A_2}/M^{A_4}) \right\},$$

where $C, c > 0$ depend only on $K, K' > 0$. Taking $R = (L+1)s_0 - 1$, using the lemma

hypotheses and that $\lambda \leq 2\bar{\lambda} = 6\|X\|\sqrt{\log M}$, the last display is bounded by

$$2C \exp \left\{ \left(c + \frac{144}{\phi(S_0)^2} + \frac{(L+1)m_{\max} \log 4}{\log M} + \frac{L \log A_2}{\log M} - LA_4 \right) s_0 \log M \right\}.$$

For $M > 0$ large enough that $\log M \geq \max\left\{ \frac{4m_{\max} \log 4}{A_2}, \frac{2 \log A_2}{A_2} \right\}$, this is then bounded by

$$2C \exp \left\{ \left(c + \frac{144}{\phi(S_0)^2} + \frac{A_2}{4} - LA_4/2 \right) s_0 \log M \right\}.$$

□

We next obtain a contraction rate for the full posterior underlying the variational approximation. The proof follows that of Theorem 3 of [Castillo et al. \(2015\)](#), modified for the group setting.

Lemma 5 (Contraction). *Suppose that Assumption (K) holds and the prior satisfies (37) and (38). Further assume $M > 0$ is large enough that $\log M \geq \max\{\frac{4m_{\max}\log 4}{A_2}, \frac{4\log A_2}{A_2}\}$ and $M \geq (4^{m_{\max}+1/2}A_2)^{1/A_4}$. Then there exists a constant $H_0 = H_0(A_1, A_3, A_4) > 0$ such that for any $\beta_0 \in \mathbb{R}^p$ such that $m_{\max}\log s_0 \leq K' \log M$ and any $L > 0$,*

$$\begin{aligned} E_{\beta_0} \Pi \left(\beta : \|X(\beta - \beta_0)\|_2 \geq \frac{H_0 \sqrt{(L+2)s_0 \log M}}{\bar{\phi}((L+2)s_0)} \middle| Y \right) 1_{\mathcal{T}_0} \\ \leq C(K, K') \exp \left\{ \left(c(K, K', A_2) + \frac{144}{\phi(S_0)^2} - \frac{LA_4}{2} \right) s_0 \log M \right\}, \end{aligned}$$

where $s_0 = |S_{\beta_0}|$ and \mathcal{T}_0 is the event (39). Moreover, both

$$\begin{aligned} E_{\beta_0} \Pi \left(\beta : \|\beta - \beta_0\|_{2,1} \geq \frac{H_0(L+2)s_0 \sqrt{\log M}}{\|X\| \bar{\phi}((L+2)s_0)^2} \middle| Y \right) 1_{\mathcal{T}_0}, \\ E_{\beta_0} \Pi \left(\beta : \|\beta - \beta_0\|_2 \geq \frac{H_0 \sqrt{(L+2)s_0 \log M}}{\|X\| \tilde{\phi}((L+2)s_0)^2} \middle| Y \right) 1_{\mathcal{T}_0}, \end{aligned}$$

satisfy the same inequality.

Proof. Consider the set $E = E_L = \{\beta \in \mathbb{R}^p : |S_\beta| \leq (L+1)s_0\}$, which satisfies

$$E_{\beta_0} \Pi(E^c | Y) 1_{\mathcal{T}_0} \leq C(K, K') \exp \left\{ \left(c(K, K', A_2) + \frac{144}{\phi(S_0)^2} - \frac{LA_4}{2} \right) s_0 \log M \right\} \quad (47)$$

by Lemma 4. Recall that $\lambda \|\beta_0\|_{2,1} \leq 2\bar{\lambda} \|\beta - \beta_0\|_{2,1} + \lambda \|\beta\|_{2,1}$ by (38). Using this, for any set $B \subseteq E$ and $\mathcal{T}_0 = \{\max_k \|X_{G_k}^T(Y - X\beta_0)\|_2 \leq \bar{\lambda}\}$, (44) and (45) give

$$\Pi(B|Y) 1_{\mathcal{T}_0} \leq C \frac{e^{c s_0 \log M}}{s_0! \pi_M(s_0)} \int_B 1_{\mathcal{T}_0} e^{-\frac{1}{2} \|X(\beta - \beta_0)\|_2^2 + 3\bar{\lambda} \|\beta - \beta_0\|_{2,1} + \lambda \|\beta\|_{2,1}} d\Pi(\beta).$$

Note that for any $\beta \in E$, $|S_{\beta - \beta_0}| \leq |S_\beta| + s_0 \leq (L+2)s_0 =: D_L s_0$. Using Definition 2 of

the uniform compatibility $\bar{\phi}(s)$,

$$\begin{aligned} (4-1)\bar{\lambda}\|\beta - \beta_0\|_{2,1} &\leq \frac{4\bar{\lambda}\|X(\beta - \beta_0)\|_2|S_{\beta-\beta_0}|^{1/2}}{\|X\|\bar{\phi}(|S_{\beta-\beta_0}|)} - \bar{\lambda}\|\beta - \beta_0\|_{2,1} \\ &\leq \frac{1}{8}\|X(\beta - \beta_0)\|_2^2 + \frac{32\bar{\lambda}^2 D_L s_0}{\|X\|^2 \bar{\phi}(D_L s_0)^2} - \bar{\lambda}\|\beta - \beta_0\|_{2,1}. \end{aligned}$$

Thus for any $B \subseteq E$,

$$\Pi(B|Y)1_{\mathcal{T}_0} \leq C \frac{e^{cs_0 \log M}}{s_0! \pi_M(s_0)} e^{\frac{32\bar{\lambda}^2 D_L s_0}{\|X\|^2 \bar{\phi}(D_L s_0)^2}} \int_B 1_{\mathcal{T}_0} e^{-\frac{3}{8}\|X(\beta-\beta_0)\|_2^2 - \bar{\lambda}\|\beta-\beta_0\|_{2,1} + \lambda\|\beta\|_{2,1}} d\Pi(\beta).$$

Note that by (37), $\pi_M(s_0) \geq (A_1 M^{-A_3})^{s_0} \pi_M(0)$. For $B = \{\beta \in E : \|X(\beta - \beta_0)\|_2 \geq R\}$,

the last display implies

$$\Pi(B|Y)1_{\mathcal{T}_0} \leq C \frac{e^{cs_0 \log M}}{s_0! A_1^{s_0} \pi_M(0)} M^{A_3 s_0} e^{\frac{32\bar{\lambda}^2 D_L s_0}{\|X\|^2 \bar{\phi}(D_L s_0)^2}} e^{-3R^2/8} \int_B e^{-\bar{\lambda}\|\beta-\beta_0\|_{2,1} + \lambda\|\beta\|_{2,1}} d\Pi(\beta).$$

The last integral is upper bounded by

$$\sum_{S:|S|\leq(L+1)s_0} \frac{\pi_M(|S|)}{\binom{M}{|S|}} \int e^{-\bar{\lambda}\|\beta-\beta_0\|_{2,1}} \prod_{G_k \in S} \Delta_{m_k} \lambda^{m_k} d\beta_{G_k} = \sum_{s=0}^{(L+1)s_0} \pi_M(s) \prod_{G_k \in S} (\lambda/\bar{\lambda})^{m_k}.$$

Using (37) and (38), the last display is bounded by $\pi_M(0) \sum_{s=0}^{\infty} (A_2 M^{-A_4})^s 2^{m_{\max} s} \leq \pi_M(0) \sum_{s=0}^{\infty} 4^{-s} \leq$

$2\pi_M(0)$, since $2^{m_{\max}} A_2 M^{-A_4} \leq 2^{-m_{\max}+1} \leq 1/4$ by the lemma hypothesis. Setting $R^2 =$

$H_0^2 D_L s_0 \log M / \bar{\phi}(D_L s_0)^2$ for a constant $H_0 > 0$, the second last display is bounded by

$$\begin{aligned} C(K, K') \exp \left\{ \left(c(K, K') + A_3 - \frac{\log A_1}{\log M} + \frac{288 D_L}{\bar{\phi}(D_L s_0)^2} \right) s_0 \log M - \frac{3}{8} R^2 \right\} \\ \leq C \exp \left\{ - \left[\frac{3H_0^2}{8} - 288 - c - A_3 - \frac{|\log A_1|}{\log M} \right] \frac{D_L s_0 \log M}{\bar{\phi}(D_L s_0)^2} \right\}, \end{aligned}$$

where we have used that $\bar{\phi}(D_L s_0) \leq \bar{\phi}(1) \leq 1$. Taking H_0 large enough depending on

A_1, A_3, A_4 and using again that $\bar{\phi}(D_L s_0) \leq 1$, the last display can be made smaller than

(47), which is thus the dominant term in the tail probability bound.

For the $\|\cdot\|_{2,1}$ -loss, using Definition 2 of the uniform compatibility, for any $\beta \in E$,

$$\bar{\lambda}\|\beta - \beta_0\|_{2,1} \leq \frac{\bar{\lambda}\|X(\beta - \beta_0)\|_2|S_{\beta-\beta_0}|^{1/2}}{\|X\|\bar{\phi}(|S_{\beta-\beta_0}|)} \leq \frac{1}{2} \frac{\bar{\lambda}^2 D_L s_0}{\|X\|^2 \bar{\phi}(D_L s_0)^2} + \frac{1}{2} \|X(\beta - \beta_0)\|_2^2.$$

The result then follows from the first assertion for the prediction loss $\|X(\beta - \beta_0)\|_2$.

For the $\|\cdot\|_2$ -loss, note that $\|X(\beta - \beta_0)\|_2 \geq \tilde{\phi}(D_L s_0) \|X\| \|\beta - \beta_0\|_2$ for any $\beta \in E$. Since $\tilde{\phi}(s) \leq \bar{\phi}(s)$ for all s by Lemma 1 of [Castillo et al. \(2015\)](#) (which extends to the group setting), the result follows. \square

F.3 D_{KL} -divergence between the variational and true posterior

We next bound the KL-divergence between the variational family and the true posterior on the following event:

$$\mathcal{T}_1 = \mathcal{T}_1(\Gamma, \varepsilon, \kappa) = \mathcal{T}_0 \cap \{\Pi(\beta : |S_\beta| > \Gamma | Y) \leq 1/4\} \cap \{\Pi(\beta : \|\beta - \beta_0\|_2 > \varepsilon | Y) \leq e^{-\kappa}\}, \quad (48)$$

where $\Gamma, \kappa, \varepsilon > 0$. The proof largely follows Section B.2 of [Ray and Szabó \(2022\)](#), again modified to the group sparse setting, which we produce in full for completeness.

Lemma 6. *If $4e^{1+\Gamma \log M - \kappa} \leq 1$, then there exists an element $Q \in \mathcal{Q} \subset \mathcal{Q}'$ of the variational families such that*

$$\begin{aligned} D_{\text{KL}}(Q || \Pi(\cdot | Y)) 1_{\mathcal{T}_1} &\leq \Gamma \left(\log M + m_{\max} \log \frac{1}{\tilde{\phi}(\Gamma)} \right) + \frac{\lambda \Gamma}{\tilde{\phi}(\Gamma)^2} \left(3s_0^{1/2} \varepsilon + \frac{3\sqrt{\log M}}{\|X\|} + \frac{m_{\max}^{1/2}}{\|X\|} \right) \\ &\quad + \log(4e). \end{aligned}$$

If Assumption (K) also holds, then

$$\begin{aligned} D_{\text{KL}}(Q || \Pi(\cdot | Y)) 1_{\mathcal{T}_1} &\leq \Gamma(K + 1) \log M \log \frac{1}{\tilde{\phi}(\Gamma)} + \frac{\lambda \Gamma}{\tilde{\phi}(\Gamma)^2} \left(3s_0^{1/2} \varepsilon + \frac{(3 + \sqrt{K})\sqrt{\log M}}{\|X\|} \right) \\ &\quad + \log(4e). \end{aligned}$$

Proof. The full posterior takes the form

$$\Pi(\cdot | Y) = \sum_{S: S \subseteq \{1, \dots, M\}} \hat{q}_S \Pi_S(\cdot | Y) \otimes \delta_{S^c}, \quad (49)$$

with weights $0 \leq \hat{q}_S \leq 1$ satisfying $\sum_S \hat{q}_S = 1$ and where $\Pi_S(\cdot|Y)$ denotes the posterior for $\beta_S \in \mathbb{R}^{\sum_{G_k \in S} m_k}$ in the restricted model $Y = X_S \beta_S + Z$. Arguing exactly as in the proof of Lemma B.2 of [Ray and Szabó \(2022\)](#), one has that on \mathcal{T}_1 and for $4e^{1+\Gamma} \log M^{-\kappa} \leq 1$, there exists a set $\tilde{S} \subseteq \{G_1, \dots, G_M\}$ such that

$$|\tilde{S}| \leq \Gamma, \quad \|\beta_{0,S^c}\|_2 \leq \varepsilon, \quad \hat{q}_{\tilde{S}} \geq (2e)^{-1} M^{-\Gamma}. \quad (50)$$

Writing $\tilde{p} = \sum_{G_k \in \tilde{S}} m_k$, consider the element of the variational family

$$\tilde{Q} = \left(\bigotimes_{G_k \in \tilde{S}} N_{m_k}(\mu_{G_k}, D_{G_k}) \right) \otimes \left(\bigotimes_{G_k \in \tilde{S}^c} \delta_{G_k} \right) =: N_{\tilde{p}}(\mu_{\tilde{S}}, D_{\tilde{S}}) \otimes \delta_{\tilde{S}^c} \quad (51)$$

where D_{G_k} are diagonal matrices to be defined below. This is the distribution that assigns mass one to the model \tilde{S} and then fits an independent normal distribution with diagonal covariance on each of the groups in \tilde{S} . Since \tilde{Q} is only absolutely continuous with respect to the $\hat{q}_{\tilde{S}} \Pi_{\tilde{S}}(\cdot|Y) \otimes \delta_{\tilde{S}^c}$ component of the posterior (49),

$$\begin{aligned} \inf_{Q \in \mathcal{Q}} D_{\text{KL}}(Q \| \Pi(\cdot|Y)) &\leq D_{\text{KL}}(\tilde{Q} \| \Pi(\cdot|Y)) = E_{\beta \sim \tilde{Q}} \left[\log \frac{dN_{\tilde{p}}(\mu_{\tilde{S}}, D_{\tilde{S}}) \otimes \delta_{\tilde{S}^c}(\beta)}{\hat{q}_{\tilde{S}} d\Pi_{\tilde{S}}(\cdot|Y) \otimes \delta_{\tilde{S}^c}(\beta)} \right] \\ &= \log \frac{1}{\hat{q}_{\tilde{S}}} + D_{\text{KL}}(N_{\tilde{p}}(\mu_{\tilde{S}}, D_{\tilde{S}}) \| \Pi_{\tilde{S}}(\cdot|Y)). \end{aligned} \quad (52)$$

On \mathcal{T}_1 , the first term is bounded by $\log(2eM^\Gamma)$ by (50), so that it remains to bound the second term in (52).

Define

$$\mu_{\tilde{S}} = (X_{\tilde{S}}^T X_{\tilde{S}})^{-1} X_{\tilde{S}}^T Y, \quad D_{\tilde{S}} = \text{diag}((D_{G_k})_{G_k \in \tilde{S}}), \quad \Sigma_{\tilde{S}} = (X_{\tilde{S}}^T X_{\tilde{S}})^{-1}, \quad (53)$$

where $D_{G_k} \in \mathbb{R}^{m_k \times m_k}$, $G_k \in \tilde{S}$, is the diagonal matrix with entries

$$(D_{G_k})_{ii} = \frac{1}{(X_{G_k}^T X_{G_k})_{ii}}, \quad i = 1, \dots, m_k.$$

Writing $E_{\mu_{\tilde{S}}, D_{\tilde{S}}}$ for the expectation under the distribution $\beta_{\tilde{S}} \sim N_{\tilde{p}}(\mu_{\tilde{S}}, D_{\tilde{S}})$,

$$\begin{aligned} D_{\text{KL}}(N_{\tilde{p}}(\mu_{\tilde{S}}, D_{\tilde{S}}) \| \Pi_{\tilde{S}}(\cdot|Y)) &= E_{\mu_{\tilde{S}}, D_{\tilde{S}}} \left[\log \frac{dN_{\tilde{p}}(\mu_{\tilde{S}}, D_{\tilde{S}})}{dN_{\tilde{p}}(\mu_{\tilde{S}}, \Sigma_{\tilde{S}})} + \log \frac{dN_{\tilde{p}}(\mu_{\tilde{S}}, \Sigma_{\tilde{S}})}{d\Pi_{\tilde{S}}(\cdot|Y)} \right] \\ &=: (I) + (II). \end{aligned} \quad (54)$$

We next deal with each term separately.

Term (I) in (54). Using the formula for the Kullback-Leibler divergence between two multivariate Gaussian distributions,

$$(I) = D_{\text{KL}}(N_{\tilde{p}}(\mu_{\tilde{S}}, D_{\tilde{S}}) \| N_{\tilde{p}}(\mu_{\tilde{S}}, \Sigma_{\tilde{S}})) = \frac{1}{2} \left(\text{Tr}(\Sigma_{\tilde{S}}^{-1} D_{\tilde{S}}) - \tilde{p} + \log(|\Sigma_{\tilde{S}}|/|D_{\tilde{S}}|) \right),$$

where $|A|$ denotes the determinant of a square matrix A . Further define the matrix

$$V_{\tilde{S}} = \text{diag}(((X_{G_k}^T X_{G_k})^{-1})_{G_k \in \tilde{S}}),$$

which is a block-diagonalization of $\Sigma_{\tilde{S}} = (X_{\tilde{S}}^T X_{\tilde{S}})^{-1}$. Let $\tilde{S} = \{G_{k_1}, \dots, G_{k_s}\}$ for $s = |\tilde{S}|$.

By considering multiplication along the block structure, $\Sigma_{\tilde{S}}^{-1} V_{\tilde{S}}$ has $(i, j)^{\text{th}}$ -block equal to $(X_{G_{k_i}}^T X_{G_{k_j}})(X_{G_{k_j}}^T X_{G_{k_j}})^{-1}$, $i, j = 1, \dots, s$. Furthermore, $V_{\tilde{S}}^{-1} D_{\tilde{S}}$ is a block-diagonal matrix with $(i, i)^{\text{th}}$ -block $(X_{G_{k_i}}^T X_{G_{k_i}}) D_{G_{k_i}}$, $i = 1, \dots, s$. Thus, by considering the block-diagonal terms,

$$\begin{aligned} \text{Tr}(\Sigma_{\tilde{S}}^{-1} D_{\tilde{S}}) &= \text{Tr}(\Sigma_{\tilde{S}}^{-1} V_{\tilde{S}} V_{\tilde{S}}^{-1} D_{\tilde{S}}) = \text{Tr} \left(\text{diag}((X_{G_{k_i}}^T X_{G_{k_i}} D_{G_{k_i}})_{G_{k_i} \in \tilde{S}}) \right) \\ &= \sum_{G_k \in \tilde{S}} \text{Tr} \left(X_{G_{k_i}}^T X_{G_{k_i}} D_{G_{k_i}} \right) = \sum_{G_k \in \tilde{S}} m_k = \tilde{p}, \end{aligned}$$

and hence $(I) = \frac{1}{2} \log(|\Sigma_{\tilde{S}}| |D_{\tilde{S}}^{-1}|)$. Using (53),

$$|D_{\tilde{S}}^{-1}| = \prod_{G_k \in \tilde{S}} \prod_{i \in G_k} (X_{G_k}^T X_{G_k})_{ii} \leq \prod_{G_k \in \tilde{S}} \|X\|^{2m_k} = \|X\|^{2\tilde{p}}.$$

Let $\lambda_{\max}(A)$ and $\lambda_{\min}(A)$ denote the largest and smallest eigenvalues, respectively, of a matrix A . Arguing as in equation (B.12) in Ray and Szabó (2022), for any $S \subseteq \{G_1, \dots, G_M\}$,

$$\lambda_{\min}(X_S^T X_S) \geq \|X\|^2 \tilde{\phi}(|S|)^2. \quad (55)$$

Therefore,

$$|\Sigma_{\tilde{S}}| = 1/|X_{\tilde{S}}^T X_{\tilde{S}}| \leq 1/\lambda_{\min}(X_{\tilde{S}}^T X_{\tilde{S}})^{\tilde{p}} \leq 1/(\|X\| \tilde{\phi}(|\tilde{S}|))^{2\tilde{p}},$$

and hence $(I) = \frac{1}{2} \log(|\Sigma_{\tilde{S}}| |D_{\tilde{S}}^{-1}|) \leq \tilde{p} \log(1/\tilde{\phi}(|\tilde{S}|)) \leq m_{\max} \Gamma \log(1/\tilde{\phi}(\Gamma))$ using (50).

Term (II) in (54). One can check that the $N_{\tilde{p}}(\mu_{\tilde{S}}, \Sigma_{\tilde{S}})$ distribution has density function proportional to $e^{-\frac{1}{2}\|Y - X_{\tilde{S}}\beta_{\tilde{S}}\|_2^2}$ for $\beta_{\tilde{S}} \in \mathbb{R}^{\tilde{p}}$. Therefore,

$$\begin{aligned} (II) &= E_{\mu_{\tilde{S}}, D_{\tilde{S}}} \left[\log \frac{D_{\Pi} e^{-\frac{1}{2}\|Y - X_{\tilde{S}}\beta_{\tilde{S}}\|_2^2 - \lambda\|\beta_{0,\tilde{S}}\|_{2,1}}}{D_N e^{-\frac{1}{2}\|Y - X_{\tilde{S}}\beta_{\tilde{S}}\|_2^2 - \lambda\|\beta_{\tilde{S}}\|_{2,1}}} \right] \\ &= \log(D_{\Pi}/D_N) + \lambda E_{\mu_{\tilde{S}}, D_{\tilde{S}}} (\|\beta_{\tilde{S}}\|_{2,1} - \|\beta_{0,\tilde{S}}\|_{2,1}), \end{aligned}$$

where $D_{\Pi} = \int_{\mathbb{R}^{\tilde{p}}} e^{-\frac{1}{2}\|Y - X_{\tilde{S}}\beta_{\tilde{S}}\|_2^2 - \lambda\|\beta_{\tilde{S}}\|_{2,1}} d\beta_{\tilde{S}}$ and $D_N = \int_{\mathbb{R}^{\tilde{p}}} e^{-\frac{1}{2}\|Y - X_{\tilde{S}}\beta_{\tilde{S}}\|_2^2 - \lambda\|\beta_{0,\tilde{S}}\|_{2,1}} d\beta_{\tilde{S}}$ are the normalizing constants for the densities. Arguing exactly as in the proof of Lemma B.2 of Ray and Szabó (2022), one can show that on the event \mathcal{T}_1 it holds that $\log(D_{\Pi}/D_N) \leq 2\lambda\Gamma^{1/2}\varepsilon + \log 2$ if $4e^{1+\Gamma \log M - \kappa} \leq 1$. On \mathcal{T}_1 , the second term in the last display is bounded by

$$\begin{aligned} \lambda E_{\mu_{\tilde{S}}, D_{\tilde{S}}} \|\beta_{\tilde{S}} - \beta_{0,\tilde{S}}\|_{2,1} &\leq \lambda |\tilde{S}|^{1/2} E_{\mu_{\tilde{S}}, D_{\tilde{S}}} \|\beta_{\tilde{S}} - \beta_{0,\tilde{S}}\|_2 \\ &\leq \lambda \Gamma^{1/2} (\|\mu_{\tilde{S}} - \beta_{0,\tilde{S}}\|_2 + E_{0, D_{\tilde{S}}} \|\beta_{\tilde{S}}\|_2) \end{aligned} \tag{56}$$

using Cauchy-Schwarz.

Under P_{β_0} , so that $Y = X\beta_0 + \varepsilon$, and using (53),

$$\|\mu_{\tilde{S}} - \beta_{0,\tilde{S}}\|_2 \leq \|(X_{\tilde{S}}^T X_{\tilde{S}})^{-1} X_{\tilde{S}}^T X_{\tilde{S}^c} \beta_{0,\tilde{S}^c}\|_2 + \|(X_{\tilde{S}}^T X_{\tilde{S}})^{-1} X_{\tilde{S}}^T \varepsilon\|_2.$$

Arguing for this term as in Lemma B.2 of Ray and Szabó (2022), one can show that the first term is bounded by $\Gamma^{1/2} s_0^{1/2} \varepsilon / \tilde{\phi}(|\tilde{S}|)^2$ on \mathcal{T}_1 . Using that the ℓ_2 -operator norm of $(X_{\tilde{S}}^T X_{\tilde{S}})^{-1}$ is bounded by $1/(\|X\| \tilde{\phi}(|\tilde{S}|))^2$ by (55), the second term in the last display is bounded by $\|X_{\tilde{S}}^T \varepsilon\|_2 / (\|X\|^2 \tilde{\phi}(|\tilde{S}|)^2)$. But on $\mathcal{T}_1 \subset \mathcal{T}_0$,

$$\|X_{\tilde{S}}^T \varepsilon\|_2^2 = \sum_{G_k \in \tilde{S}} \|X_{G_k}^T (Y - X\beta_0)\|_2^2 \leq 9|\tilde{S}| \|X\|^2 \log M.$$

Combining the above bounds thus yields

$$\|\mu_{\tilde{S}} - \beta_{0,\tilde{S}}\|_2 \leq \frac{\Gamma^{1/2} s_0^{1/2} \varepsilon}{\tilde{\phi}(\Gamma)^2} + \frac{3\Gamma^{1/2} \sqrt{\log M}}{\|X\| \tilde{\phi}(|\tilde{S}|)^2}$$

on \mathcal{T}_1 , thereby controlling the first term in (56).

It remains only to bound the second term in (56). Let e_i denote the i^{th} unit vector in \mathbb{R}^{m_k} , $i = 1, \dots, m_k$, and let \bar{e}_i denote its extension to $\mathbb{R}^{\tilde{p}}$ with unit entry in the i^{th} coordinate of group G_k . Then $(X_{G_k}^T X_{G_k})_{ii} = \|X_{G_k} e_i\|_2^2 = \|X \bar{e}_i\|_2^2 \geq \|X\|^2 \tilde{\phi}(1)^2$ for $i = 1, \dots, m_k$. Therefore,

$$E_{0, D_{\tilde{S}}} \|\beta_{\tilde{S}}\|_2^2 = \text{Tr}(D_{\tilde{S}}) = \sum_{G_k \in \tilde{S}} \sum_{i \in G_k} \frac{1}{(X_{\tilde{S}}^T X_{\tilde{S}})_{ii}} \leq \frac{\tilde{p}}{\|X\|^2 \tilde{\phi}(1)^2}.$$

Putting together of all the above bounds yields

$$(II) \leq 2\lambda\Gamma^{1/2}\varepsilon + \log 2 + \frac{\lambda\Gamma}{\tilde{\phi}(\Gamma)^2} \left(s_0^{1/2}\varepsilon + \frac{3\sqrt{\log M}}{\|X\|} + \frac{m_{\max}^{1/2}}{\|X\|} \right),$$

using that $\tilde{p} \leq m_{\max}|\tilde{S}|$, $|\tilde{S}| \leq \Gamma$ and $\tilde{\phi}(|\tilde{S}|) \leq \tilde{\phi}(1) \leq 1$ for $\tilde{S} \neq \emptyset$.

Substituting the bounds for (I) and (II) just derived into (52) and (54) then gives the first result. The second result follows from using that $m_{\max} \leq K \log M$ under Assumption (K). \square

F.4 Proofs of Theorems 3 and 4

To complete the proofs, we apply Lemma 1 with the event \mathcal{T}_1 defined in (48) for suitably chosen constants $\Gamma, \varepsilon, \kappa > 0$.

Lemma 7. *Suppose that Assumption (K) holds, the prior satisfies (37)-(38) and s_n satisfies $m_{\max} \log s_n \leq K' \log M$ for some $K' > 0$. Then*

$$\inf_{\substack{\beta_0 \in \mathbb{R}^p: \phi(S_{\beta_0}) \geq c_0 \\ |S_{\beta_0}| \leq s_n}} P_{\beta_0}(\mathcal{T}_1(\Gamma_{s_0}, \varepsilon_{s_0}, \kappa_{s_0})) \rightarrow 1$$

as $n \rightarrow \infty$, where

$$\Gamma_{s_0} = C_{\Gamma} s_0, \quad \varepsilon_{s_0} = C_{\varepsilon} \frac{\sqrt{s_0 \log M}}{\|X\| \tilde{\phi}(C_{\phi} s_0)}, \quad \kappa_{s_0} = (C_{\Gamma} s_0 + 1) \log M, \quad (57)$$

and the constants have dependence $C_\Gamma = C_\Gamma(K, K', A_2, A_4, c_0)$, $C_\varepsilon = C_\varepsilon(K, K', A_1 - A_4, c_0)$ and $C_\phi = C_\phi(K, K', A_2, A_4, c_0)$. Moreover, $4e^{1+\Gamma_{s_0} \log M - \kappa_{s_0}} \leq 1$ for $M > 0$ large enough.

Proof. We consider each of the sets in \mathcal{T}_1 individually. In what follows, $C = C(K, K')$ and $c = c(K, K', A_2)$ will be constants that may change line-by-line but which will not depend on other parameters.

Applying Lemma 4 with $L = \frac{2}{A_4}(1 + \log C + c + 144/c_0^2)$, where C, c are the constants in that lemma, we have $E_{\beta_0} \Pi(\beta : |S_\beta| > (L + 1)s_0 | Y) 1_{\mathcal{T}_0} \leq e^{-s_0 \log M} \leq 1/4$ for M large enough. Setting $C_\Gamma = L + 1$, the second event in \mathcal{T}_1 is thus a subset of \mathcal{T}_0 for $\Gamma = \Gamma_{s_0}$ and all β_0 in the infimum in the lemma.

For the third event in \mathcal{T}_1 , we apply Lemma 5 with $L = \frac{2}{A_4}(1 + \log C + c + 144/c_0^2 + C_\Gamma)$, where now C, c are the constants in Lemma 5, to obtain

$$E_{\beta_0} \Pi \left(\beta : \|\beta - \beta_0\|_2 \geq \frac{H_0 \sqrt{(L + 2)s_0 \log M}}{\|X\| \tilde{\phi}((L + 2)s_0)^2} \middle| Y \right) 1_{\mathcal{T}_0} \leq e^{-\kappa_{s_0}}.$$

Setting $C_\varepsilon = H_0 \sqrt{L + 2}$ and $C_\phi = L + 2$ shows that the third event in \mathcal{T}_1 is also contained in \mathcal{T}_0 for these choices and all β_0 in the infimum in the lemma. It thus suffices to control the probability of \mathcal{T}_0 , which tends to one uniformly in $\beta_0 \in \mathbb{R}^p$ by Lemma 2. Lastly, note that $4e^{1+\Gamma_{s_0} \log M - \kappa_{s_0}} = 4e^{1 - \log M} \leq 1$ for M large enough. \square

Proof of Theorem 3. Write $\mathcal{B}_n = \mathcal{B}_{\rho_n, s_n} = \{\beta_0 \in \mathbb{R}^p : \phi(S_{\beta_0}) \geq c_0, |S_{\beta_0}| \leq s_n, \tilde{\phi}(\rho_n | S_{\beta_0}) \geq c_0\}$ for the parameter set and let

$$\Omega_n = \left\{ \beta : \|X(\beta - \beta_0)\|_2 \leq \frac{H_0 \rho_n^{1/2} \sqrt{s_0 \log M}}{\tilde{\phi}(\rho_n s_0)} \right\}$$

for $H_0 = H_0(A_1, A_3, A_4) > 0$ the constant in Lemma 5. Let $\mathcal{T}_1 = \mathcal{T}_1 = \mathcal{T}_1(\Gamma_{s_0}, \varepsilon_{s_0}, \kappa_{s_0})$ be the event (48) with parameters (57), so that Lemma 7 gives $E_{\beta_0} \tilde{\Pi}(\Omega_n^c) = E_{\beta_0} \tilde{\Pi}(\Omega_n^c) 1_{\mathcal{T}_1} + o(1)$, uniformly over $\beta_0 \in \mathcal{B}_n$.

Applying Lemma 5 with $L+2 = \rho_n \rightarrow \infty$ gives that for n large enough, $E_{\beta_0} \Pi(\Omega_n^c | Y) 1_{\mathcal{T}_0} \leq C e^{-c\rho_n s_0 \log M}$, where C, c depend only on K, K', A_2, A_4, c_0 . We can then use Lemma 1 with event $A_n = \mathcal{T}_1 \subset \mathcal{T}_0$ and $\delta_n = c\rho_n s_0 \log M$ to obtain

$$E_{\beta_0} \tilde{\Pi}(\Omega_n^c) 1_{\mathcal{T}_1} \leq \frac{2}{c\rho_n s_0 \log M} \left[D_{\text{KL}}(\tilde{\Pi} \| \Pi(\cdot | Y)) 1_{\mathcal{T}_1} + C^{-c\rho_n s_0 \log M} \right].$$

Since $\rho_n s_0 \log M \rightarrow \infty$ and $\tilde{\Pi}$ is by definition the KL-minimizer of the variational family to the posterior, we obtain that for any $Q \in \mathcal{Q}$,

$$E_{\beta_0} \tilde{\Pi}(\Omega_n^c) \leq \frac{2}{c\rho_n s_0 \log M} D_{\text{KL}}(Q \| \Pi(\cdot | Y)) 1_{\mathcal{T}_1} + o(1),$$

where the $o(1)$ term is uniform over $\beta_0 \in \mathcal{B}_n$. But by Lemma 6, there exists an element $Q \in \mathcal{Q}$ of the variational family such that

$$D_{\text{KL}}(Q \| \Pi(\cdot | Y)) 1_{\mathcal{T}_1} \leq C \left(\log \frac{1}{\tilde{\phi}(C_\Gamma s_0)} + \frac{1}{\tilde{\phi}(C_\Gamma s_0)^2 \tilde{\phi}(C_\phi s_0)} \right) s_0 \log M,$$

where C is uniform over $\beta_0 \in \mathcal{B}_n$ and C_Γ, C_ϕ have dependences specified in Lemma 7. But since C_Γ, C_ϕ are bounded under the hypothesis of the present theorem and $\rho_n \rightarrow \infty$, we have $\rho_n \geq C_\Gamma, C_\phi$ for n large enough and hence $\tilde{\phi}(C_\Gamma s_0), \tilde{\phi}(C_\phi s_0) \geq \tilde{\phi}(\rho_n s_0) \geq c_0$. Using this and the last two displays, we thus have

$$\sup_{\beta_0 \in \mathcal{B}_n} E_{\beta_0} \tilde{\Pi}(\Omega_n^c) \leq C/\rho_n + o(1) = o(1)$$

as $n \rightarrow \infty$, thereby establishing the first assertion. The second two assertions follow similarly for using the corresponding results in Lemma 5. \square

Proof of Theorem 4. The proof follows similarly to the proof of Theorem 3, using Lemma 4 to control the probability of the posterior set instead of Lemma 5. \square