

Laplace Approximation in High-Dimensional Bayesian Regression

Rina Foygel Barber, Mathias Drton, and Kean Ming Tan

Abstract We consider Bayesian variable selection in sparse high-dimensional regression, where the number of covariates p may be large relative to the sample size n , but at most a moderate number q of covariates are active. Specifically, we treat generalized linear models. For a single fixed sparse model with well-behaved prior distribution, classical theory proves that the Laplace approximation to the marginal likelihood of the model is accurate for sufficiently large sample size n . We extend this theory by giving results on uniform accuracy of the Laplace approximation across all models in a high-dimensional scenario in which p and q , and thus also the number of considered models, may increase with n . Moreover, we show how this connection between marginal likelihood and Laplace approximation can be used to obtain consistency results for Bayesian approaches to variable selection in high-dimensional regression.

1 Introduction

A key issue in Bayesian approaches to model selection is the evaluation of the marginal likelihood, also referred to as the evidence, of the different models that are being considered. While the marginal likelihood may sometimes be available in closed form when adopting suitable priors, most problems require approximation techniques. In particular, this is the case for variable selection in generalized linear models such as logistic regression, which are the models treated in this paper. Different strategies to approximate the marginal likelihood are reviewed in [7]. Our

R.F. Barber

Department of Statistics, The University of Chicago, Chicago, IL 60637, USA

e-mail: rina@uchicago.edu

M. Drton (✉)

Department of Statistics, University of Washington, Seattle, WA 98195, USA

e-mail: md5@uw.edu

K.M. Tan

Department of Biostatistics, University of Washington, Seattle, WA 98195, USA

e-mail: keanming@uw.edu

focus will be on the accuracy of the Laplace approximation that is derived from large-sample theory; see also Sect. 4.4 in [3].

Suppose we have n independent observations of a response variable, and along with each observation we record a collection of p covariates. Write $L(\beta)$ for the likelihood function of a generalized linear model relating the response to the covariates, where $\beta \in \mathbb{R}^p$ is a vector of coefficients in the linear predictor [13]. Let $f(\beta)$ be a prior distribution, and let $\hat{\beta}$ be the maximum likelihood estimator (MLE) of the parameter vector $\beta \in \mathbb{R}^p$. Then the evidence for the (saturated) regression model is the integral

$$\int_{\mathbb{R}^p} L(\beta)f(\beta) d\beta,$$

and the Laplace approximation is the estimate

$$\text{Laplace} := L(\hat{\beta})f(\hat{\beta}) \left(\frac{(2\pi)^p}{\det H(\hat{\beta})} \right)^{1/2},$$

where H denotes the negative Hessian of the log-likelihood function $\log L$.

Classical asymptotic theory for large sample size n but fixed number of covariates p shows that the Laplace approximation is accurate with high probability [9]. With p fixed, this then clearly also holds for variable selection problems in which we would consider every one of the finitely many models given by the 2^p subsets of covariates. This accuracy result justifies the use of the Laplace approximation as a proxy for an actual model evidence. The Laplace approximation is also useful for proving frequentist consistency results about Bayesian methods for variable selection for a general class of priors. This is again discussed in [9]. The ideas go back to the work of Schwarz [15] on the Bayesian information criterion (BIC).

In this paper, we set out to give analogous results on the interplay between Laplace approximation, model evidence, and frequentist consistency in variable selection for regression problems that are high-dimensional, possibly with $p > n$, and sparse in that we consider only models that involve small subsets of covariates. We denote by q an upper bound on the number of active covariates. In variable selection for sparse high-dimensional regression, the number of considered models is very large, on the order of p^q . Our interest is then in bounds on the approximation error of Laplace approximations that, with high probability, hold uniformly across all sparse models. Theorem 1, our main result, gives such uniform bounds (see Sect. 3). A numerical experiment supporting the theorem is described in Sect. 4.

In Sect. 5, we show that when adopting suitable priors on the space of all sparse models, model selection by maximizing the product of model prior and Laplace approximation is consistent in an asymptotic scenario in which p and q may grow with n . As a corollary, we obtain a consistency result for fully Bayesian variable selection methods. We note that the class of priors on models we consider is the same as the one that has been used to define extensions of BIC that have consistency

properties for high-dimensional variable selection problems; see, for example, [2, 4–6, 8, 11, 12, 17]. The prior has also been discussed in [16].

2 Setup and Assumptions

In this section, we provide the setup for the studied problem and the assumptions needed for our results.

2.1 Problem Setup

We treat generalized linear models for n independent observations of a response, which we denote as $Y_1, \dots, Y_n$. Each observation Y_i follows a distribution from a univariate exponential family with density

$$p_\theta(y) \propto \exp \{y \cdot \theta - b(\theta)\} , \quad \theta \in \mathbb{R},$$

where the density is defined with respect to some measure on $\mathbb{R}$. Let θ_i be the (natural) parameter indexing the distribution of Y_i , so $Y_i \sim p_{\theta_i}$. The vector $\boldsymbol{\theta} = (\theta_1, \dots, \theta_n)^T$ is then assumed to lie in the linear space spanned by the columns of a design matrix $X = (X_{ij}) \in \mathbb{R}^{n \times p}$, that is, $\boldsymbol{\theta} = X\boldsymbol{\beta}$ for a parameter vector $\boldsymbol{\beta} \in \mathbb{R}^p$. Our work is framed in a setting with a fixed/deterministic design X . In the language of McCullagh and Nelder [13], our basic setup uses a canonical link, no dispersion parameter and an exponential family whose natural parameter space is the entire real line. This covers, for instance, logistic and Poisson regression. However, extensions beyond this setting are possible; see, for instance, the related work of Luo and Chen [11] whose discussion of Bayesian information criteria encompasses other link functions.

We write X_i for the i th row of X , that is, the p -vector of covariate values for observation Y_i . The regression model for the responses then has log-likelihood, score, and negative Hessian functions

$$\begin{aligned} \log L(\boldsymbol{\beta}) &= \sum_{i=1}^n Y_i \cdot X_i^T \boldsymbol{\beta} - b(X_i^T \boldsymbol{\beta}) \in \mathbb{R} , \\ s(\boldsymbol{\beta}) &= \sum_{i=1}^n X_i (Y_i - b'(X_i^T \boldsymbol{\beta})) \in \mathbb{R}^p , \\ H(\boldsymbol{\beta}) &= \sum_{i=1}^n X_i X_i^T \cdot b''(X_i^T \boldsymbol{\beta}) \in \mathbb{R}^{p \times p} . \end{aligned}$$

The results in this paper rely on conditions on the Hessian H , and we note that, implicitly, these are actually conditions on the design X .

We are concerned with a sparsity scenario in which the joint distribution of $Y_1, \dots, Y_n$ is determined by a true parameter vector $\beta_0 \in \mathbb{R}^p$ supported on a (small) set $J_0 \subset [p] := \{1, \dots, p\}$, that is, $\beta_{0j} \neq 0$ if and only if $j \in J_0$. Our interest is in the recovery of the set J_0 when knowing an upper bound q on the cardinality of J_0 , so $|J_0| \leq q$. To this end, we consider the different submodels given by the linear spaces spanned by subsets $J \subset [p]$ of the columns of the design matrix X , where $|J| \leq q$.

For notational convenience, we take $J \subset [p]$ to mean either an index set for the covariates or the resulting regression model. The regression coefficients in model J form a vector of length $|J|$. We index such vectors β by the elements of J , that is, $\beta = (\beta_j : j \in J)$, and we write $\mathbb{R}^J$ for the Euclidean space containing all these coefficient vectors. This way the coefficient and the covariate it belongs to always share a common index. In other words, the coefficient for the j -th coordinate of covariate vector X_i is denoted by β_j in any model J with $j \in J$. Furthermore, it is at times convenient to identify a vector $\beta \in \mathbb{R}^J$ with the vector in $\mathbb{R}^p$ that is obtained from β by filling in zeros outside of the set J . As this is clear from the context, we simply write β again when referring to this sparse vector in $\mathbb{R}^p$. Finally, $s_J(\beta)$ and $H_J(\beta)$ denote the subvector and submatrix of $s(\beta)$ and $H(\beta)$, respectively, obtained by extracting entries indexed by J . These depend only on the subvectors $X_{iJ} = (X_{ij})_{j \in J}$ of the covariate vectors X_i .

2.2 Assumptions

Recall that n is the sample size, p is the number of covariates, q is an upper bound on the model size, and β_0 is the true parameter vector. We assume the following conditions to hold for all considered regression problems:

- (A1) The Euclidean norm of the true signal is bounded, that is, $\|\beta_0\|_2 \leq a_0$ for a fixed constant $a_0 \in (0, \infty)$.
- (A2) There is a decreasing function $c_{\text{lower}} : [0, \infty) \rightarrow (0, \infty)$ and an increasing function $c_{\text{upper}} : [0, \infty) \rightarrow (0, \infty)$ such that for all $J \subset [p]$ with $|J| \leq 2q$ and all $\beta \in \mathbb{R}^J$, the Hessian of the negative log-likelihood function is bounded as

$$c_{\text{lower}}(\|\beta\|_2) \mathbf{I}_J \leq \frac{1}{n} H_J(\beta) \leq c_{\text{upper}}(\|\beta\|_2) \mathbf{I}_J.$$

- (A3) There is a constant $c_{\text{change}} \in (0, \infty)$ such that for all $J \subset [p]$ with $|J| \leq 2q$ and all $\beta, \beta' \in \mathbb{R}^J$,

$$\frac{1}{n} \|H_J(\beta) - H_J(\beta')\|_{\text{sp}} \leq c_{\text{change}} \cdot \|\beta - \beta'\|_2,$$

where $\|\cdot\|_{\text{sp}}$ is the spectral norm of a matrix.

Assumption (A2) provides control of the spectrum of the Hessian of the negative log-likelihood function, and (A3) yields control of the change of the Hessian. Together, (A2) and (A3) imply that for all $\epsilon > 0$, there is a $\delta > 0$ such that

$$(1 - \epsilon)H_J(\beta_0) \preceq H_J(\beta_J) \preceq (1 + \epsilon)H_J(\beta_0), \quad (1)$$

for all $J \supseteq J_0$ with $|J| \leq 2q$ and $\beta_J \in \mathbb{R}^J$ with $\|\beta_J - \beta_0\|_2 \leq \delta$; see Proposition 2.1 in [2]. Note also that we consider sets J with cardinality $2q$ in (A2) and (A3) because it allows us to make arguments concerning false models, with $J \not\supseteq J_0$, using properties of the true model given by the union $J \cup J_0$.

Remark 1 When treating generalized linear models, some control of the size of the true coefficient vector β_0 is indeed needed. For instance, in logistic regression, if the norm of β_0 is too large, then the binary response will take on one of its values with overwhelming probability. Keeping with the setting of logistic regression, Barber and Drton [2] show how assumptions (A2) and (A3) hold with high probability in certain settings in which the covariates are generated as i.i.d. sample. Assumptions (A2) and (A3), or the implication from (1), also appear in earlier work on Bayesian information criteria for high-dimensional problems such as [6] or [11].

Let $\{f_J : J \subset [p], |J| \leq q\}$ be a family of probability density functions $f_J : \mathbb{R}^J \rightarrow [0, \infty)$ that we use to define prior distributions in all q -sparse models. We say that the family is log-Lipschitz with respect to radius $R > 0$ and has bounded log-density ratios if there exist two constants $F_1, F_2 \in [0, \infty)$ such that the following conditions hold for all $J \subset [p]$ with $|J| \leq q$:

(B1) The function $\log f_J$ is F_1 -Lipschitz on the ball $B_R(0) = \{\beta \in \mathbb{R}^J : \|\beta\|_2 \leq R\}$, i.e., for all $\beta', \beta \in B_R(0)$, we have

$$|\log f_J(\beta') - \log f_J(\beta)| \leq F_1 \|\beta' - \beta\|_2.$$

(B2) For all $\beta \in \mathbb{R}^J$,

$$\log f_J(\beta) - \log f_J(0) \leq F_2.$$

Example 1 If we take f_J to be the density of a $|J|$ -fold product of a centered normal distribution with variance σ^2 , then (B1) holds with $F_1 = R/\sigma^2$ and $F_2 = 0$.

3 Laplace Approximation

This section provides our main result. For a high-dimensional regression problem, we show that a Laplace approximation to the marginal likelihood of each sparse model,

$$\text{Evidence}(J) := \int_{\mathbb{R}^J} L(\beta) f_J(\beta) d\beta,$$

leads to an approximation error that, with high probability, is bounded uniformly across all models. To state our result, we adopt the notation

$$a = b(1 \pm c) \quad : \Longleftrightarrow \quad a \in [b(1 - c), b(1 + c)].$$

Theorem 1 *Suppose the conditions from (A1)–(A3) hold. Then, there are constants $\nu, c_{\text{sample}}, a_{\text{MLE}} \in (0, \infty)$ depending only on $(a_0, c_{\text{lower}}, c_{\text{upper}}, c_{\text{change}})$ such that if*

$$n \geq c_{\text{sample}} \cdot q^3 \max\{\log(p), \log^3(n)\},$$

then with probability at least $1 - p^{-\nu}$ the following two statements are true for all sparse models $J \subset [p]$, $|J| \leq q$:

- (i) *The MLE $\hat{\beta}_J$ satisfies $\|\hat{\beta}_J\|_2 \leq a_{\text{MLE}}$.*
- (ii) *If additionally the family of prior densities $\{f_J : J \subset [p], |J| \leq q\}$ satisfies the Lipschitz condition from (B1) for radius $R \geq a_{\text{MLE}} + 1$, and has log-density ratios bounded as in (B2), then there is a constant $c_{\text{Laplace}} \in (0, \infty)$ depending only on $(a_0, c_{\text{lower}}, c_{\text{upper}}, c_{\text{change}}, F_1, F_2)$ such that*

$$\text{Evidence}(J) = L(\hat{\beta}_J) f_J(\hat{\beta}_J) \cdot \left(\frac{(2\pi)^{|J|}}{\det H_J(\hat{\beta}_J)} \right)^{1/2} \cdot \left(1 \pm c_{\text{Laplace}} \sqrt{\frac{|J|^3 \log^3(n)}{n}} \right).$$

Proof

- (i) *Bounded MLEs.* It follows from Sect. B.2 in [2]¹ that, with the claimed probability, the norms $\|\hat{\beta}_J\|_2$ for true models J (i.e., $J \supseteq J_0$ and $|J| \leq 2q$) are bounded by a constant. The result makes reference to an event for which all the claims we make subsequently are true. The bound on the norm of an MLE of a true model was obtained by comparing the maximal likelihood to the likelihood at the true parameter β_0 . As we show now, for false sparse models, we may argue similarly but comparing to the likelihood at 0.

Recall that a_0 is the bound on the norm of β_0 assumed in (A1) and that the functions c_{lower} and c_{upper} in (A2) are decreasing and increasing in the norm of β_0 , respectively. Throughout this part, we use the abbreviations

$$c_{\text{lower}} := c_{\text{lower}}(a_0), \quad c_{\text{upper}} := c_{\text{upper}}(a_0).$$

¹In the proof of this theorem, we cite several results from Sect. B.2 and Lemma B.1 in [2]. Although that paper treats the specific case of logistic regression, by examining the proofs of their results that we cite here, we can see that they hold more broadly for the general GLM case as long as we assume that the Hessian conditions hold, i.e., Conditions (A1)–(A3), and therefore we may use these results for the setting considered here.

First, we lower-bound the likelihood at 0 via a Taylor-expansion using the true model J_0 . For some $t \in [0, 1]$, we have that

$$\begin{aligned} \log L(0) - \log L(\beta_0) &= -\beta_0^T s_{J_0}(\beta_0) - \frac{1}{2} \beta_0^T H_{J_0}(t\beta_0) \beta_0 \geq -\beta_0^T s_{J_0}(\beta_0) \\ &\quad - \frac{1}{2} n a_0^2 c_{\text{upper}}, \end{aligned}$$

where we have applied (A2). Lemma B.1 in [2] yields that

$$|\beta_0^T s_{J_0}(\beta_0)| \leq \|H_{J_0}(\beta_0)^{-\frac{1}{2}} s_{J_0}(\beta_0)\|_2 \|H_{J_0}(\beta_0)^{\frac{1}{2}} \beta_0\|_2 \leq \tau_0 a_0 \sqrt{n c_{\text{upper}}},$$

where τ_0^2 can be bounded by a constant multiple of $q \log(p)$. By our sample size assumption (i.e., the existence of the constant c_{sample}), we thus have that

$$\log L(0) - \log L(\beta_0) \geq -n \cdot c_1 \quad (2)$$

for some constant $c_1 \in (0, \infty)$.

Second, we may consider the true model $J \cup J_0$ instead of J and apply (B.17) in [2] to obtain the bound

$$\begin{aligned} \log L(\hat{\beta}_J) - \log L(\beta_0) &\leq \\ &\|\hat{\beta}_J - \beta_0\|_2 \cdot \left(\sqrt{n c_{\text{upper}}} \cdot \tau_{J \setminus J_0} - \frac{n c_{\text{lower}}}{4} \min \left\{ \|\hat{\beta}_J - \beta_0\|_2, \frac{c_{\text{lower}}}{2 c_{\text{change}}} \right\} \right), \end{aligned} \quad (3)$$

where $\tau_{J \setminus J_0}^2$ can be bounded by a constant multiple of $q \log(p)$. Choosing our sample size constant c_{sample} large enough, we may deduce from (3) that there is a constant $c_2 \in (0, \infty)$ such that

$$\log L(\hat{\beta}_J) - \log L(\beta_0) \leq -n \|\hat{\beta}_J - \beta_0\|_2 c_2$$

whenever $\|\hat{\beta}_J - \beta_0\|_2 > c_{\text{lower}}/(2c_{\text{change}})$. Using the fact that $\log L(0) \leq \log(\hat{\beta}_J)$ for any model J , we may deduce from (3) that there is a constant $c_2 \in (0, \infty)$ such that

$$\log L(0) - \log L(\beta_0) \leq -n \|\hat{\beta}_J - \beta_0\|_2 c_2$$

whenever $\|\hat{\beta}_J - \beta_0\|_2 > c_{\text{lower}}/(2c_{\text{change}})$. Together with (2), this implies that $\|\hat{\beta}_J - \beta_0\|_2$ is bounded by a constant c_3 . Having assumed (A1), we may conclude that the norm of $\hat{\beta}_J$ is bounded by $a_0 + c_3$.

(ii) *Laplace approximation.* Fix $J \subset [p]$ with $|J| \leq q$. In order to analyze the evidence of model J , we split the integration domain $\mathbb{R}^J$ into two regions,

namely, a neighborhood $\mathcal{N}$ of the MLE $\hat{\beta}_J$ and the complement $\mathbb{R}^J \setminus \mathcal{N}$. More precisely, we choose the neighborhood of the MLE as

$$\mathcal{N} := \left\{ \beta \in \mathbb{R}^J : \|H_J(\hat{\beta}_J)^{1/2}(\beta - \hat{\beta}_J)\|_2 \leq \sqrt{5|J| \log(n)} \right\}.$$

Then the marginal likelihood, $\text{Evidence}(J)$, is the sum of the two integrals

$$\begin{aligned} \mathcal{J}_1 &= \int_{\mathcal{N}} L(\beta) f_J(\beta) d\beta, \\ \mathcal{J}_2 &= \int_{\mathbb{R}^J \setminus \mathcal{N}} L(\beta) f_J(\beta) d\beta. \end{aligned}$$

We will estimate $\mathcal{J}_1$ via a quadratic approximation to the log-likelihood function. Outside of the region $\mathcal{N}$, the quadratic approximation may no longer be accurate but due to concavity of the log-likelihood function, the integrand can be bounded by $e^{-c\|\beta_J - \hat{\beta}_J\|_2}$ for an appropriately chosen constant c , which allows us to show that $\mathcal{J}_2$ is negligible when n is sufficiently large.

We now approximate $\mathcal{J}_1$ and $\mathcal{J}_2$ separately. Throughout this part we assume that we have a bound a_{MLE} on the norms of the MLEs $\hat{\beta}_J$ in sparse models J with $|J| \leq q$. For notational convenience, we now let

$$c_{\text{lower}} := c_{\text{lower}}(a_{\text{MLE}}), \quad c_{\text{upper}} := c_{\text{upper}}(a_{\text{MLE}}).$$

(ii-a) *Approximation of integral $\mathcal{J}_1$.* By a Taylor-expansion, for any $\beta \in \mathbb{R}^J$ there is a $t \in [0, 1]$ such that

$$\log L(\beta) = \log L(\hat{\beta}_J) - \frac{1}{2}(\beta - \hat{\beta}_J)^T H_J \left(\hat{\beta}_J + t(\beta - \hat{\beta}_J) \right) (\beta - \hat{\beta}_J).$$

By (A3) and using that $|t| \leq 1$,

$$\left\| H_J \left(\hat{\beta}_J + t(\beta - \hat{\beta}_J) \right) - H_J(\hat{\beta}_J) \right\|_{\text{sp}} \leq n \cdot c_{\text{change}} \|\beta - \hat{\beta}_J\|_2.$$

Hence,

$$\log L(\beta) = \log L(\hat{\beta}_J) - \frac{1}{2}(\beta - \hat{\beta}_J)^T H_J(\hat{\beta}_J)(\beta - \hat{\beta}_J) \pm \frac{1}{2} \|\beta - \hat{\beta}_J\|_2^3 \cdot n c_{\text{change}}. \quad (4)$$

Next, observe that (A2) implies that

$$H_J(\hat{\beta}_J)^{-1/2} \preceq \sqrt{\frac{1}{nc_{\text{lower}}}} \cdot \mathbf{I}_J.$$

We deduce that for any vector $\beta \in \mathcal{N}$,

$$\|\beta - \hat{\beta}_J\|_2 \leq \sqrt{5|J| \log(n)} \|H_J(\hat{\beta}_J)^{-1/2}\|_{\text{sp}} \leq \sqrt{\frac{5|J| \log(n)}{nc_{\text{lower}}}}. \quad (5)$$

This gives

$$\begin{aligned} \log L(\beta) &= \log L(\hat{\beta}_J) - \frac{1}{2}(\beta - \hat{\beta}_J)^T H_J(\hat{\beta}_J)(\beta - \hat{\beta}_J) \\ &\quad \pm \sqrt{\frac{|J|^3 \log^3(n)}{n}} \cdot \sqrt{\frac{125c_{\text{change}}^2}{4c_{\text{lower}}^3}}. \end{aligned} \quad (6)$$

Choosing the constant c_{sample} large enough, we can ensure that the upper bound in (5) is no larger than 1. In other words, $\|\beta - \hat{\beta}_J\|_2 \leq 1$ for all points $\beta \in \mathcal{N}$. By our assumption that $\|\hat{\beta}_J\|_2 \leq a_{\text{MLE}}$, the set $\mathcal{N}$ is thus contained in the ball

$$\mathcal{B} = \{\beta \in \mathbb{R}^J : \|\beta\|_2 \leq a_{\text{MLE}} + 1\}.$$

Since, by (B1), the logarithm of the prior density is F_1 -Lipschitz on $\mathcal{B}$, it follows from (5) that

$$\log f_J(\beta) = \log f_J(\hat{\beta}_J) \pm F_1 \|\beta - \hat{\beta}_J\|_2 = \log f_J(\hat{\beta}_J) \pm F_1 \sqrt{\frac{5|J| \log(n)}{nc_{\text{lower}}}}. \quad (7)$$

Plugging (6) and (7) into $\mathcal{J}_1$, and writing $a = b \cdot \exp\{\pm c\}$ as a shorthand for $a \in [b \cdot e^{-c}, b \cdot e^c]$, we find that

$$\begin{aligned} \mathcal{J}_1 &= L(\hat{\beta}_J) f_J(\hat{\beta}_J) \exp \left\{ \pm \sqrt{\frac{|J|^3 \log^3(n)}{n}} \cdot \left(\sqrt{\frac{5F_1^2}{nc_{\text{lower}}}} + \sqrt{\frac{125c_{\text{change}}^2}{4c_{\text{lower}}^3}} \right) \right\} \\ &\quad \times \int_{\mathcal{N}} \exp \left\{ -\frac{1}{2}(\beta - \hat{\beta}_J)^T H_J(\hat{\beta}_J)(\beta - \hat{\beta}_J) \right\} d\beta. \end{aligned} \quad (8)$$

In the last integral, change variables to $\xi = H_J(\hat{\beta}_J)^{1/2}(\beta - \hat{\beta}_J)$ to see that

$$\begin{aligned} &\int_{\mathcal{N}} \exp \left\{ -\frac{1}{2}(\beta - \hat{\beta}_J)^T H_J(\hat{\beta}_J)(\beta - \hat{\beta}_J) \right\} d\beta \\ &= \left(\det H_J(\hat{\beta}_J) \right)^{-1/2} \cdot \int_{\|\xi\|_2 \leq \sqrt{5|J| \log(n)}} \exp \left\{ -\frac{1}{2}\|\xi\|_2^2 \right\} d\xi \end{aligned}$$

$$\begin{aligned}
&= \left(\frac{(2\pi)^{|J|}}{\det H_J(\hat{\beta}_J)} \right)^{1/2} \cdot \Pr \left\{ \chi_{|J|}^2 \leq 5|J| \log(n) \right\} \\
&= \left(\frac{(2\pi)^{|J|}}{\det H_J(\hat{\beta}_J)} \right)^{1/2} \cdot \exp\{\pm 1/\sqrt{n}\}, \tag{9}
\end{aligned}$$

where we use a tail bound for the χ^2 -distribution stated in Lemma 1. We now substitute (9) into (8), and simplify the result using that $e^{-x} \geq 1 - 2x$ and $e^x \leq 1 + 2x$ for all $0 \leq x \leq 1$. We find that

$$\begin{aligned}
\mathcal{J}_1 &= L(\hat{\beta}_J) f_J(\hat{\beta}_J) \left(\frac{(2\pi)^{|J|}}{\det H_J(\hat{\beta}_J)} \right)^{1/2} \\
&\quad \times \left(1 \pm 2 \left(1 + \sqrt{\frac{125c_{\text{change}}^2}{4c_{\text{lower}}^3}} + \sqrt{\frac{5F_1^2}{c_{\text{lower}}}} \right) \sqrt{\frac{|J|^3 \log^3(n)}{n}} \right) \tag{10}
\end{aligned}$$

when the constant c_{sample} is chosen large enough.

(ii-b) *Approximation of integral $\mathcal{J}_2$.* Let β be a point on the boundary of $\mathcal{N}$. It then holds that

$$(\beta - \hat{\beta}_J)^T H_J(\hat{\beta}_J) (\beta - \hat{\beta}_J) = \sqrt{5|J| \log(n)} \cdot \|H_J(\hat{\beta}_J)^{1/2} (\beta - \hat{\beta}_J)\|_2.$$

We may deduce from (6) that

$$\begin{aligned}
\log L(\beta) &\leq \log L(\hat{\beta}_J) - \frac{\sqrt{5|J| \log(n)}}{2} \|H_J(\hat{\beta}_J)^{1/2} (\beta - \hat{\beta}_J)\|_2 \\
&\quad + \sqrt{\frac{|J|^3 \log^3(n)}{n}} \cdot \sqrt{\frac{125c_{\text{change}}^2}{4c_{\text{lower}}^3}} \\
&\leq \log L(\hat{\beta}_J) - \|H_J(\hat{\beta}_J)^{1/2} (\beta - \hat{\beta}_J)\|_2 \cdot \sqrt{|J| \log(n)},
\end{aligned}$$

for $|J|^3 \log^3(n)/n$ sufficiently small, which can be ensured by choosing c_{sample} large enough. The concavity of the log-likelihood function now implies that for all $\beta \notin \mathcal{N}$ we have

$$\log L(\beta) \leq \log L(\hat{\beta}_J) - \|H_J(\hat{\beta}_J)^{1/2} (\beta - \hat{\beta}_J)\|_2 \cdot \sqrt{|J| \log(n)}. \tag{11}$$

Moreover, using first assumption (B2) and then assumption (B1), we have that

$$\log f_J(\beta) \leq \log f_J(0) + F_2 \leq \log f_J(\hat{\beta}_J) + F_1 \|\hat{\beta}_J\|_2 + F_2.$$

Since $\|\hat{\beta}_J\|_2 \leq a_{\text{MLE}}$, it thus holds that

$$\log f_J(\beta) \leq \log f_J(\hat{\beta}_J) + F_1 a_{\text{MLE}} + F_2. \quad (12)$$

Combining the bounds from (11) and (12), the integral can be bounded as

$$\begin{aligned} \mathcal{J}_2 &\leq L(\hat{\beta}_J) f_J(\hat{\beta}_J) e^{F_1 a_{\text{MLE}} + F_2} \\ &\quad \times \int_{\mathbb{R}^J \setminus \mathcal{N}} \exp \left\{ -\|H_J(\hat{\beta}_J)^{1/2}(\beta - \hat{\beta}_J)\|_2 \cdot \sqrt{|J| \log(n)} \right\} d\beta. \end{aligned} \quad (13)$$

Changing variables to $\xi = H_J(\hat{\beta}_J)^{1/2}(\beta - \hat{\beta}_J)$ and applying Lemma 2, we may bound the integral in (13) as

$$\begin{aligned} &\int_{\mathbb{R}^J \setminus \mathcal{N}} \exp \left\{ -\|H_J(\hat{\beta}_J)^{1/2}(\beta - \hat{\beta}_J)\|_2 \cdot \sqrt{|J| \log(n)} \right\} d\beta \\ &\leq \left(\det H_J(\hat{\beta}_J) \right)^{-1/2} \cdot \int_{\|\xi\|_2 > \sqrt{5|J| \log(n)}} \exp \left\{ -\sqrt{|J| \log(n)} \cdot \|\xi\|_2 \right\} d\xi \\ &\leq \left(\det H_J(\hat{\beta}_J) \right)^{-1/2} \cdot \frac{4(\pi)^{|J|/2}}{\Gamma(\frac{1}{2}|J|)} \frac{\sqrt{5|J| \log(n)}^{|J|-1}}{\sqrt{|J| \log(n)}} e^{-\sqrt{5}|J| \log(n)} \\ &= \left(\frac{(2\pi)^{|J|}}{\det H_J(\hat{\beta}_J)} \right)^{1/2} \cdot \frac{2\sqrt{5}}{\Gamma(\frac{1}{2}|J|)} \left(\frac{5}{2}|J| \log(n) \right)^{|J|/2-1} \cdot \frac{1}{n^{\sqrt{5}|J|}}. \end{aligned}$$

Stirling's lower bound on the Gamma function gives

$$\frac{(|J|/2)^{|J|/2-1}}{\Gamma(\frac{1}{2}|J|)} = \frac{(|J|/2)^{|J|/2}}{\Gamma(\frac{1}{2}|J| + 1)} \leq \frac{1}{\sqrt{|J|\pi}} e^{|J|/2}.$$

Using this inequality, and returning to (13), we see that

$$\begin{aligned} \mathcal{J}_2 &\leq L(\hat{\beta}_J) f_J(\hat{\beta}_J) \left(\frac{(2\pi)^{|J|}}{\det H_J(\hat{\beta}_J)} \right)^{1/2} \\ &\quad \times e^{F_1 a_{\text{MLE}} + F_2} \cdot \frac{2e\sqrt{5}}{\sqrt{|J|\pi}} \cdot \left(\frac{5e \log(n)}{n} \right)^{|J|/2-1} \frac{1}{n^{(\sqrt{5}-1/2) \cdot |J|+1}}. \end{aligned} \quad (14)$$

Based on this fact, we certainly have the very loose bound that

$$\mathcal{J}_2 \leq L(\hat{\beta}_J)f_J(\hat{\beta}_J) \left(\frac{(2\pi)^{|J|}}{\det H_J(\hat{\beta}_J)} \right)^{1/2} e^{F_1 a_{\text{MLE}} + F_2} \cdot \frac{1}{\sqrt{n}}, \quad (15)$$

for sufficiently large n .

(ii-c) *Combining the bounds.* From (10) and (15), we obtain that

$$\begin{aligned} \text{Evidence}(J) = \mathcal{J}_1 + \mathcal{J}_2 &= L(\hat{\beta}_J)f_J(\hat{\beta}_J) \left(\frac{(2\pi)^{|J|}}{\det H_J(\hat{\beta}_J)} \right)^{1/2} \times \\ &\left(1 \pm \left(e^{F_1 a_{\text{MLE}} + F_2} + 2 + \sqrt{\frac{125c_{\text{change}}^2}{c_{\text{lower}}^3}} + \sqrt{\frac{20F_1^2}{c_{\text{lower}}}} \right) \sqrt{\frac{|J|^3 \log^3(n)}{n}} \right) \end{aligned} \quad (16)$$

for sufficiently large n , as desired.

Remark 2 The proof of Theorem 1 could be modified to handle other situations of interest. For instance, instead of a fixed Lipschitz constant F_1 for all log prior densities, one could consider the case where $\log f_J$ is Lipschitz with respect to a constant $F_1(J)$ that grows with the cardinality of $|J|$, e.g., at a rate of $\sqrt{|J|}$ in which case the rate of square root of $|J|^3 \log^3(n)/n$ could be modified to square root of $|J|^4 \log^3(n)/n$. The term $e^{F_1(J)a_{\text{MLE}}}$ that would appear in (16) could be compensated using (14) in less crude of a way than when moving to (15).

4 Numerical Experiment for Sparse Bayesian Logistic Regression

In this section, we perform a simulation study to assess the approximation error in Laplace approximations to the marginal likelihood of logistic regression models. To this end, we generate independent covariate vectors $X_1, \dots, X_n$ with i.i.d. $N(0, 1)$ entries. For each choice of a (small) value of q , we take the true parameter vector $\beta_0 \in \mathbb{R}^p$ to have the first q entries equal to two and the rest of the entries equal zero. So, $J_0 = \{1, \dots, q\}$. We then generate independent binary responses $Y_1, \dots, Y_n$, with values in $\{0, 1\}$ and distributed as $(Y_i|X_i) \sim \text{Bernoulli}(p_i(X_i))$, where

$$p_i(x) = \left(\frac{\exp(x^T \beta_0)}{1 + \exp(x^T \beta_0)} \right) \iff \log \left(\frac{p_i(x)}{1 - p_i(x)} \right) = x \cdot \beta_0,$$

based on the usual (and canonical) logit link function.

We record that the logistic regression model with covariates indexed by $J \subset [p]$ has the likelihood function

$$L(\beta) = \exp \left\{ \sum_{i=1}^n Y_i \cdot X_{iJ}^T \beta - \log(1 + \exp(X_{iJ}^T \beta)) \right\}, \quad \beta \in \mathbb{R}^J, \quad (17)$$

where, as previously defined, $X_{iJ} = (X_{ij})_{j \in J}$ denotes the subset of covariates for model J . The negative Hessian of the log-likelihood function is

$$H_J(\beta) = \sum_{i=1}^n X_{iJ} X_{iJ}^T \cdot \frac{\exp(X_{iJ}^T \beta)}{(1 + \exp(X_{iJ}^T \beta))^2}.$$

For Bayesian inference in the logistic regression model given by J , we consider as a prior distribution a standard normal distribution on $\mathbb{R}^J$, that is, the distribution of a random vector with $|J|$ independent $N(0, 1)$ coordinates. As in previous section, we denote the resulting prior density by f_J . We then wish to approximate the evidence or marginal likelihood

$$\text{Evidence}(J) := \int_{\mathbb{R}^J} L(\beta) f_J(\beta) d\beta.$$

As a first approximation, we use a Monte Carlo approach in which we simply draw independent samples $\beta^1, \dots, \beta^B$ from the prior f_J and estimate the evidence as

$$\text{MonteCarlo}(J) = \frac{1}{B} \sum_{b=1}^B L(\beta^b),$$

where we use $B = 50,000$ in all of our simulations. As a second method, we compute the Laplace approximation

$$\text{Laplace}(J) := L(\hat{\beta}_J) f_J(\hat{\beta}_J) \left(\frac{(2\pi)^{|J|}}{\det H_J(\hat{\beta}_J)} \right)^{1/2},$$

where $\hat{\beta}_J$ is the maximum likelihood estimator in model J . For each choice of the number of covariates p , the model size q , and the sample size n , we calculate the Laplace approximation error as

$$\max_{J \subset [p], |J| \leq q} |\log \text{MonteCarlo}(J) - \log \text{Laplace}(J)|.$$

We consider $n \in \{50, 60, 70, 80, 90, 100\}$ in our experiment. Since we wish to compute the Laplace approximation error of every q -sparse model, and the number of possible models is on the order of p^q , we consider $p = n/2$ and $q \in \{1, 2, 3\}$.




Fig. 1 Maximum Laplace approximation error, averaged over 100 data sets, as a function of the sample size n . The number of covariates is $n/2$, and the number of active covariates is bounded by $q \in \{1, 2, 3\}$

The Laplace approximation error, averaged across 100 independent simulations, is shown in Fig. 1. We remark that the error in the Monte Carlo approximation to the marginal likelihood is negligible compared to the quantity plotted in Fig. 1. With two independent runs of our Monte Carlo integration routine, we found the Monte Carlo error to be on the order of 0.05.

For each $q = 1, 2, 3$, Fig. 1 shows a decrease in Laplace approximation error as n increases. We emphasize that p and thus also the number of considered q -sparse models increase with n . As we increase the number of active covariates q , the Laplace approximation error increases. These facts are in agreement with Theorem 1. This said, the scope of this experiment is clearly limited by the fact that only small values of q and moderate values of p and n are computationally feasible.

5 Consistency of Bayesian Variable Selection

In this section, we apply the result on uniform accuracy of the Laplace approximation (Theorem 1) to prove a high-dimensional consistency result for Bayesian variable selection. Here, consistency refers to the property that the probability of choosing the most parsimonious true model tends to one in a large-sample limit. As

discussed in Sect. 1, we consider priors of the form

$$P_\gamma(J) \propto \left(\frac{p}{|J|} \right)^{-\gamma} \cdot \mathbb{1}\{|J| \leq q\}, \quad J \subset [p], \quad (18)$$

where $\gamma \geq 0$ is a parameter that allows one to interpolate between the case of a uniform distribution on models ($\gamma = 0$) and a prior for which the model cardinality $|J|$ is uniformly distributed ($\gamma = 1$).

Bayesian variable selection is based on maximizing the (unnormalized) posterior probability

$$\text{Bayes}_\gamma(J) := \left(\frac{p}{|J|} \right)^{-\gamma} \text{Evidence}(J) \quad (19)$$

over $J \subset [p]$, $|J| \leq q$. Approximate Bayesian variable section can be based on maximizing instead the quantity

$$\text{Laplace}_\gamma(J) := \left(\frac{p}{|J|} \right)^{-\gamma} \text{Laplace}(J). \quad (20)$$

We will identify asymptotic scenarios under which maximization of Laplace_γ yields consistent variable selection. Using Theorem 1, we obtain as a corollary that fully Bayesian variable selection, i.e., maximization of Bayes_γ , is consistent as well.

To study consistency, we consider a sequence of variable selection problems indexed by the sample size n , where the n -th problem has p_n covariates, true parameter $\beta_0(n)$ with support $J_0(n)$, and signal strength

$$\beta_{\min}(n) = \min_{j \in J_0(n)} |(\beta_0(n))_j|.$$

In addition, let q_n be the upper bound on the size of the considered models. The following consistency result is similar to the related results for extensions of the Bayesian information criterion; see, for instance, [2, 6].

Theorem 2 *Suppose that $p_n = n^\kappa$ for $\kappa > 0$, that $q_n = n^\psi$ for $0 \leq \psi < 1/3$, that $\beta_{\min}(n) = n^{-\phi/2}$ for $0 \leq \phi < 1 - \psi$, and that $\kappa > \psi$. Assume that (A1) holds for a fixed constant a_0 and that there a fixed functions c_{lower} and c_{upper} with respect to which the covariates satisfy the Hessian conditions (A2) and (A3) for all $J \supseteq J_0(n)$ with $|J| \leq 2q_n$. Moreover, assume that for the considered family of prior densities $\{f_J(\cdot) : J \subset [p_n], |J| \leq q_n\}$ there are constants $F_3, F_4 \in (0, \infty)$ such that, uniformly for all $|J| \leq q_n$, we have*

$$\sup_{\beta} f_J(\beta) \leq F_3 < \infty, \quad \inf_{\|\beta\|_2 \leq a_{\text{MLE}}} f_J(\beta) \geq F_4 > 0,$$

where a_{MLE} is the constant from Theorem 1(i). Then, for any $\gamma > 1 - \frac{1-2\psi}{2(\kappa-\psi)}$, model selection with Laplace_γ is consistent in the sense that the event

$$J_0(n) = \arg \max \{ \text{Laplace}_\gamma(J) : J \subset [p_n], |J| \leq q_n \}$$

has probability tending to one as $n \rightarrow \infty$.

Together with Theorem 1, the proof of Theorem 2, which we give below, also shows consistency of the fully Bayesian procedure.

Corollary 1 *Under the assumptions of Theorem 2, fully Bayesian model selection is consistent, that is, the event*

$$J_0(n) = \arg \max \{ \text{Bayes}_\gamma(J) : J \subset [p_n], |J| \leq q_n \}$$

has probability tending to one as $n \rightarrow \infty$.

Proof (Proof of Theorem 2) Our scaling assumptions for p_n , q_n and $\beta_{\min}(n)$ are such that the conditions imposed in Theorem 2.2 of [2] are met for n large enough. This theorem and Theorem 1(i) in this paper then yield that there are constants $\nu, \epsilon, C_{\text{false}}, a_{\text{MLE}} > 0$ such that with probability at least $1 - p_n^\nu$ the following three statements hold simultaneously:

(a) For all $|J| \leq q_n$ with $J \supseteq J_0(n)$,

$$\log L(\hat{\beta}_J) - \log L(\hat{\beta}_{J_0(n)}) \leq (1 + \epsilon)(|J \setminus J_0(n)| + \nu) \log(p_n). \quad (21)$$

(b) For all $|J| \leq q_n$ with $J \not\supseteq J_0(n)$,

$$\log L(\hat{\beta}_{J_0(n)}) - \log L(\hat{\beta}_J) \geq C_{\text{false}} n \beta_{\min}(n)^2. \quad (22)$$

(c) For all $|J| \leq q_n$ and some constant $a_{\text{MLE}} > 0$,

$$\|\hat{\beta}_J\|_2 \leq a_{\text{MLE}}. \quad (23)$$

In the remainder of this proof we show that these three facts, in combination with further technical results from [2], imply that

$$J_0(n) = \arg \max \{ \text{Laplace}_\gamma(J) : J \subset [p_n], |J| \leq q_n \}. \quad (24)$$

For simpler notation, we no longer indicate explicitly that p_n , q_n , β_0 and derived quantities vary with n . We will then show that

$$\begin{aligned} \log \frac{\text{Laplace}_\gamma(J_0)}{\text{Laplace}_\gamma(J)} = & (\log P(J_0) - \log P(J)) + \left(\log L(\hat{\beta}_{J_0}) - \log L(\hat{\beta}_J) \right) - |J \setminus J_0| \log \sqrt{2\pi} \\ & + \left(\log f_{J_0}(\hat{\beta}_{J_0}) - \log f_J(\hat{\beta}_J) \right) + \frac{1}{2} \left(\log \det H_J(\hat{\beta}_J) - \log \det H_{J_0}(\hat{\beta}_{J_0}) \right) \end{aligned} \quad (25)$$

is positive for any model given by a set $J \neq J_0$ of cardinality $|J| \leq q$. We let

$$c_{\text{lower}} := c_{\text{lower}}(a_{\text{MLE}}), \quad c_{\text{upper}} := c_{\text{upper}}(a_{\text{MLE}}).$$

We note that this definition of c_{lower} and c_{upper} differs from the one used in the proof of Theorem 1.

False Models If $J \not\supseteq J_0$, that is, if the model is false, we observe that

$$\log P(J_0) - \log P(J) = -\gamma \log \left(\frac{p}{|J_0|} \right) + \gamma \log \left(\frac{p}{|J|} \right) \geq -\gamma \log \left(\frac{p}{|J_0|} \right) \geq -\gamma q \log p.$$

Moreover, by (A2) and (23),

$$\begin{aligned} \log \det H_J(\hat{\beta}_J) - \log \det H_{J_0}(\hat{\beta}_{J_0}) & \geq |J| \cdot \log(nc_{\text{lower}}) - |J_0| \cdot \log(nc_{\text{upper}}) \\ & \geq -q \log \left(n \frac{c_{\text{upper}}}{\min\{c_{\text{lower}}, 1\}} \right). \end{aligned}$$

Combining the lower bounds with (22), we obtain that

$$\begin{aligned} \log \frac{\text{Laplace}_\gamma(J_0)}{\text{Laplace}_\gamma(J)} & \geq C_{\text{false}} n \beta_{\min}^2 - |J \setminus J_0| \log(\sqrt{2\pi}) - q \log \left(p^\gamma n \frac{c_{\text{upper}}}{\min\{c_{\text{lower}}, 1\}} \right) + \log \left(\frac{F_4}{F_3} \right) \\ & \geq C_{\text{false}} n \beta_{\min}^2 - q \log \left(\frac{c_{\text{upper}}}{\min\{c_{\text{lower}}, 1\}} \cdot \sqrt{2\pi} n p^\gamma \right) + \log \left(\frac{F_4}{F_3} \right). \end{aligned}$$

By our scaling assumptions, the lower bound is positive for sufficiently large n .

True Models It remains to resolve the case of $J \supsetneq J_0$, that is, when model J is true. We record that from the proof of Theorem 2.2 in [2], it holds on the considered event

of probability at least $1 - p^{-\nu}$ that for any $J \supseteq J_0$,

$$\|\hat{\beta}_J - \beta_0\|_2 \leq \frac{4\sqrt{c_{\text{upper}}}}{\sqrt{nc_{\text{lower}}}} \tau_{|J \setminus J_0|}, \quad (26)$$

where

$$\tau_r^2 = \frac{2}{(1 - \epsilon')^3} \cdot \left[(J_0 + r) \log\left(\frac{3}{\epsilon'}\right) + \log(4p^\nu) + r \log(2p) \right].$$

Under our scaling assumptions on p and q , it follows that $\|\hat{\beta}_J - \beta_0\|_2$ tends to zero as $n \rightarrow \infty$.

We begin again by considering the prior on models, for which we have that

$$\begin{aligned} \log P(J_0) - \log P(J) &= \gamma \log \frac{|J_0|!(p - |J_0|)!}{|J|!(p - |J|)!} \\ &\geq -\gamma|J \setminus J_0| \log q + \gamma|J \setminus J_0| \log(p - q) \\ &\geq -\gamma|J \setminus J_0| \log q + \gamma|J \setminus J_0|(1 - \tilde{\epsilon}) \log p \end{aligned}$$

for all n sufficiently large. Indeed, we assume that $p = n^\kappa$ and $q = n^\psi$ with $\kappa > \psi$ such that $p - q \geq p^{1-\tilde{\epsilon}}$ for any small constant $\tilde{\epsilon} > 0$ as long as p is sufficiently large relative to q .

Next, if $J \supsetneq J_0$, then (A2) and (A3) allow us to relate $H_J(\hat{\beta}_J)$ and $H_{J_0}(\hat{\beta}_{J_0})$ to the respective Hessian at the true parameter, i.e., $H_J(\beta_0)$ and $H_{J_0}(\beta_0)$. We find that

$$\begin{aligned} \log \left(\frac{\det H_J(\hat{\beta}_J)}{\det H_{J_0}(\hat{\beta}_{J_0})} \right) &\geq \log \left(\frac{\det H_J(\beta_0)}{\det H_{J_0}(\beta_0)} \right) + |J| \log \left(1 - \frac{c_{\text{change}}}{c_{\text{lower}}} \|\hat{\beta}_J - \beta_0\|_2 \right) \\ &\quad - |J_0| \log \left(1 + \frac{c_{\text{change}}}{c_{\text{lower}}} \|\hat{\beta}_{J_0} - \beta_0\|_2 \right). \end{aligned}$$

Note that by assuming n large enough, we may assume that $\|\hat{\beta}_J - \beta_0\|_2$ is small enough for the logarithm to be well defined; recall (26). Using that $x \geq \log(1 + x)$ for all $x > -1$ and $\log(1 - \frac{x}{2}) \geq -x$ for all $0 \leq x \leq 1$, we see that

$$\begin{aligned} \log \left(\frac{\det H_J(\hat{\beta}_J)}{\det H_{J_0}(\hat{\beta}_{J_0})} \right) &\geq \log \left(\frac{\det H_J(\beta_0)}{\det H_{J_0}(\beta_0)} \right) - 2|J| \frac{c_{\text{change}}}{c_{\text{lower}}} \|\hat{\beta}_J - \beta_0\|_2 \\ &\quad - |J_0| \frac{c_{\text{change}}}{c_{\text{lower}}} \|\hat{\beta}_{J_0} - \beta_0\|_2. \end{aligned}$$

Under our scaling assumptions, $q^3 \log(p) = o(n)$, and thus applying (26) twice shows that

$$-2|J| \frac{c_{\text{change}}}{c_{\text{lower}}} \|\hat{\beta}_J - \beta_0\|_2 - |J_0| \frac{c_{\text{change}}}{c_{\text{lower}}} \|\hat{\beta}_{J_0} - \beta_0\|_2$$

is larger than any small negative constant for n large enough. For simplicity, we take the lower bound as -1 . By (A2), it holds that

$$\begin{aligned} \log \left(\frac{\det H_J(\beta_0)}{\det H_{J_0}(\beta_0)} \right) &= \log \det (H_{J \setminus J_0}(\beta_0) - H_{J_0, J \setminus J_0}(\beta_0)^T H_{J_0}(\beta_0)^{-1} H_{J_0, J \setminus J_0}(\beta_0)) \\ &\geq |J \setminus J_0| \log(n) + |J \setminus J_0| \log(c_{\text{lower}}), \end{aligned}$$

because the eigenvalues of the Schur complement of $H_J(\beta_0)$ are bounded the same way as the eigenvalues of $H_J(\beta_0)$; see, e.g., Chap. 2 of [18]. Hence, for sufficiently large n , the following is true for all $J \supsetneq J_0$:

$$\log \det H_J(\hat{\beta}_J) - \log \det H_{J_0}(\hat{\beta}_{J_0}) \geq |J \setminus J_0| \log(n) + |J \setminus J_0| \log(c_{\text{lower}}) - 1. \quad (27)$$

Combining the bound for the model prior probabilities with (21) and (27), we have for any true model $J \supsetneq J_0$ that

$$\begin{aligned} \log \frac{\text{Laplace}_\gamma(J_0)}{\text{Laplace}_\gamma(J)} &\geq -(1 + \epsilon)(|J \setminus J_0| + \nu) \log(p) + \gamma |J \setminus J_0| (1 - \tilde{\epsilon}) \log(p) \\ &\quad + \frac{1}{2} |J \setminus J_0| \log(n) - \gamma |J \setminus J_0| \log(q) + \frac{1}{2} |J \setminus J_0| \left(\log \frac{c_{\text{lower}}}{2\pi} \right) + \log \left(\frac{F_4}{F_3} \right) - 1. \end{aligned}$$

Collecting terms and using that $|J \setminus J_0| \geq 1$, we obtain the lower bound

$$\begin{aligned} \log \frac{\text{Laplace}_\gamma(J_0)}{\text{Laplace}_\gamma(J)} &\geq \frac{1}{2} |J \setminus J_0| \left\{ \log(n) - \log q^{2\gamma} \right. \\ &\quad \left. + 2[(1 - \tilde{\epsilon})\gamma - (1 + \epsilon)(1 + \nu)] \log(p) + \log \left(\frac{c_{\text{lower}}}{2\pi} \right) \right\} + \log \left(\frac{F_4}{F_3} \right) - 1. \end{aligned}$$

This lower bound is positive for all n large because our assumption that $p = n^\kappa$, $q = n^\psi$ for $0 \leq \psi < 1/3$, and

$$\gamma > 1 - \frac{1 - 2\psi}{2(\kappa - \psi)}$$

implies that

$$\lim_{n \rightarrow \infty} \frac{\sqrt{n}}{p^{(1+\epsilon)(1+\nu)-\gamma(1-\tilde{\epsilon})} q^\gamma} = \infty \quad (28)$$

provided the constants ϵ , ν , and $\tilde{\epsilon}$ are chosen sufficiently small.

6 Discussion

In this paper, we showed that in the context of high-dimensional variable selection problems, the Laplace approximation can be accurate uniformly across a potentially very large number of sparse models. We also demonstrated how this approximation result allows one to give results on the consistency of fully Bayesian techniques for variable selection.

In practice, it is of course infeasible to evaluate the evidence or Laplace approximation for every single sparse regression model, and some search strategy must be adopted instead. Some related numerical experiments can be found in [5, 6, 17], and [2], although that work considers BIC scores that drop some of the terms appearing in the Laplace approximation.

Finally, we emphasize that the setup we considered concerns generalized linear models without dispersion parameter and with canonical link. The conditions from [11] could likely be used to extend our results to other situations.

7 Technical Lemmas

This section provides two lemmas that were used in the proof of Theorem 1.

Lemma 1 (Chi-Square Tail Bound) *Let χ_k^2 denote a chi-square random variable with k degrees of freedom. Then, for any $n \geq 3$,*

$$\mathbb{P}\{\chi_k^2 \leq 5k \log(n)\} \geq 1 - \frac{1}{n^k} \geq \exp\{-1/\sqrt{n}\}.$$

Proof Since $\log(n) \geq 1$ when $n \geq 3$, we have that

$$k + 2\sqrt{k \cdot k \log(n)} + 2k \log(n) \leq 5k \log(n).$$

Using the chi-square tail bound in [10], it thus holds that

$$\begin{aligned} \mathbb{P}\{\chi_k^2 \leq 5k \log(n)\} &\geq \mathbb{P}\left\{\chi_k^2 \leq k + 2\sqrt{k \cdot k \log(n)} + 2k \log(n)\right\} \\ &\geq 1 - e^{-k \log(n)}. \end{aligned}$$

Finally, for the last step, by the Taylor series for $x \mapsto e^x$, for all $n \geq 3$ we have

$$\exp\{-1/\sqrt{n}\} \leq 1 - \frac{1}{\sqrt{n}} + \frac{1}{2} \cdot \frac{1}{n} \leq 1 - \frac{1}{n}.$$

Lemma 2 *Let $k \geq 1$ be any integer, and let $a, b > 0$ be such that $ab \geq 2(k-1)$. Then*

$$\int_{\|\xi\|_2 > a} \exp\{-b\|\xi\|_2\} d\xi \leq \frac{4(\pi)^{k/2}}{\Gamma(\frac{1}{2}k)} \frac{a^{k-1}}{b} e^{-ab},$$

where the integral is taken over $\xi \in \mathbb{R}^k$.

Proof We claim that the integral of interest is

$$\int_{\|\xi\|_2 > a} \exp\{-b\|\xi\|_2\} d\xi = \frac{2(\pi)^{k/2}}{b^k \Gamma(\frac{1}{2}k)} \int_{r=a}^{\infty} r^{k-1} e^{-r} dr. \quad (29)$$

Indeed, in $k = 1$ dimension,

$$\int_{\|\xi\|_2 > a} \exp\{-b\|\xi\|_2\} d\xi = 2 \int_{r=a}^{\infty} e^{-br} dr = \frac{2}{b} e^{-ab},$$

which is what (29) evaluates to. If $k \geq 2$, then using polar coordinates (see Exercises 7.1–1.3 in [1]), we find that

$$\begin{aligned} \int_{\|\xi\|_2 > a} \exp\{-b\|\xi\|_2\} d\xi &= 2\pi \int_{r=a}^{\infty} r^{k-1} e^{-br} dr \cdot \prod_{i=1}^{k-2} \int_{-\pi/2}^{\pi/2} \cos^i(\theta_i) d\theta_i \\ &= 2\pi \int_{r=a}^{\infty} r^{k-1} e^{-br} dr \cdot \prod_{i=1}^{k-2} \frac{\sqrt{\pi} \Gamma(\frac{1}{2}(i+1))}{\Gamma(\frac{1}{2}(i+2))}, \end{aligned}$$

which again agrees with the formula from (29).

Now, the integral on the right-hand side of (29) defines the upper incomplete Gamma function and can be bounded as

$$\Gamma(k, ab) = \int_{r=ab}^{\infty} r^{k-1} e^{-r} dr \leq 2e^{-ab} (ab)^{k-1}$$

for $ab \geq 2(k-1)$; see inequality (3.2) in [14]. This gives the bound that was to be proven.

References

1. Anderson, T.W.: An Introduction to Multivariate Statistical Analysis. Wiley Series in Probability and Statistics, 3rd edn. Wiley-Interscience, Hoboken, NJ (2003)
2. Barber, R.F., Drton, M.: High-dimensional Ising model selection with Bayesian information criteria. Electron. J. Stat. **9**, 567–607 (2015)

3. Bishop, C.M.: Pattern Recognition and Machine Learning. Information Science and Statistics. Springer, New York (2006)
4. Bogdan, M., Ghosh, J.K., Doerge, R.: Modifying the Schwarz Bayesian information criterion to locate multiple interacting quantitative trait loci. *Genetics* **167**(2), 989–999 (2004)
5. Chen, J., Chen, Z.: Extended Bayesian information criteria for model selection with large model spaces. *Biometrika* **95**(3), 759–771 (2008)
6. Chen, J., Chen, Z.: Extended BIC for small- n -large- P sparse GLM. *Stat. Sinica* **22**(2), 555–574 (2012)
7. Friel, N., Wyse, J.: Estimating the evidence—a review. *Statistica Neerlandica* **66**(3), 288–308 (2012)
8. Frommlet, F., Ruhaltinger, F., Twaróg, P., Bogdan, M.: Modified versions of Bayesian information criterion for genome-wide association studies. *Comput. Stat. Data Anal.* **56**(5), 1038–1051 (2012)
9. Houghton, D.M.A.: On the choice of a model to fit data from an exponential family. *Ann. Stat.* **16**(1), 342–355 (1988)
10. Laurent, B., Massart, P.: Adaptive estimation of a quadratic functional by model selection. *Ann. Stat.* **28**(5), 1302–1338 (2000)
11. Luo, S., Chen, Z.: Selection consistency of EBIC for GLIM with non-canonical links and diverging number of parameters. *Stat. Interface* **6**(2), 275–284 (2013)
12. Luo, S., Xu, J., Chen, Z.: Extended Bayesian information criterion in the Cox model with a high-dimensional feature space. *Ann. Inst. Stat. Math.* **67**(2), 287–311 (2015)
13. McCullagh, P., Nelder, J.A.: Generalized Linear Models. Monographs on Statistics and Applied Probability, 2nd edn. Chapman & Hall, London (1989)
14. Natalini, P., Palumbo, B.: Inequalities for the incomplete gamma function. *Math. Inequal. Appl.* **3**(1), 69–77 (2000)
15. Schwarz, G.: Estimating the dimension of a model. *Ann. Stat.* **6**(2), 461–464 (1978)
16. Scott, J.G., Berger, J.O.: Bayes and empirical-Bayes multiplicity adjustment in the variable-selection problem. *Ann. Stat.* **38**(5), 2587–2619 (2010)
17. Żak-Szatkowska, M., Bogdan, M.: Modified versions of the Bayesian information criterion for sparse generalized linear models. *Comput. Stat. Data Anal.* **55**(11), 2908–2924 (2011)
18. Zhang, F. (ed.): The Schur Complement and Its Applications. Numerical Methods and Algorithms, vol. 4. Springer, New York (2005)

Statistical Analysis for High-Dimensional Data

The Abel Symposium 2014

Frigessi, A.; Bühlmann, P.; Glad, I.K.; Langaas, M.;

Richardson, S.; Vannucci, M. (Eds.)

2016, XII, 306 p. 65 illus., 46 illus. in color., Hardcover

ISBN: 978-3-319-27097-5