

Bayesian Generalized Additive Model Selection Including a Fast Variational Option

BY VIRGINIA X. HE AND MATT P. WAND

University of Technology Sydney

30th October, 2023

Abstract

We use Bayesian model selection paradigms, such as group least absolute shrinkage and selection operator priors, to facilitate generalized additive model selection. Our approach allows for the effects of continuous predictors to be categorized as either zero, linear or non-linear. Employment of carefully tailored auxiliary variables results in Gibbsian Markov chain Monte Carlo schemes for practical implementation of the approach. In addition, mean field variational algorithms with closed form updates are obtained. Whilst not as accurate, this fast variational option enhances scalability to very large data sets. A package in the R language aids use in practice.

Keywords: Markov chain Monte Carlo; mean field variational Bayes; nonparametric regression; R package; scalable methodology.

1 Introduction

Generalized additive models offer attractive solutions to the problem of obtaining parsimonious, flexible and interpretable regression fits when faced with, potentially, large numbers of candidate predictors (e.g. Hastie & Tibshirani, 1990; Wood, 2017). Generalized additive models methodology and software is into its fourth decade. Nevertheless, principled, scalable and reliable selection of a model still has room for improvement. The version of the problem treated here is that where each candidate predictor is categorized into one of three classes: having zero effect, having a linear effect or having a non-linear effect on the mean response. We provide new and effective solutions to the problem by employing recent developments in Bayesian model selection and Bayesian computing. An accompanying package in the R language (R Core Team, 2023) allows immediate use of our new methodology.

Several approaches to the three-category generalized additive model selection problem have been proposed, such those in Shively *et al.* (1999), Ravikumar *et al.* (2009), Reich *et al.* (2009), Scheipl *et al.* (2012) and Chouldechova & Hastie (2015). Our approach is inspired and closely tied to that of Chouldechova & Hastie (2015) which has the advantages of excellent scalability and an accompanying R package (Chouldechova & Hastie, 2018). Key features of the Chouldechova & Hastie (2015) approach are: use of the group least absolute shrinkage and selection operator (LASSO), Demmler-Reinsch spline bases, regularization paths and cross-validatory selection of the regularization parameter. Both Gaussian and binary response cases are supported. Instead of the path and cross-validation aspects, we embed their infrastructure into a Bayesian graphical model and invoke Bayesian principles for model selection. Simulation results point to superior three-category model selection. Other advantages of our Bayesian approaches are being able to traverse a bigger model space compared with the regularization path approach and avoiding the practical difficulties associated with finding cross-validation minima.

Once a Bayesian version of the Chouldechova & Hastie (2015) model is specified, a pertinent challenge is tractability of Markov chain Monte Carlo and mean field variational Bayes approaches to approximate inference. We achieve this via the introduction of appropriate aux-

iliary variables. The binary response case benefits from the Albert & Chib (1993) auxiliary variable approach for probit links. The resultant graphical models are such that all full conditional distributions have standard forms. As a consequence, Markov chain Monte Carlo sampling is Gibbsian and the mean field variational Bayes have closed forms – both of which depend only on sufficient statistics of the input data. Combined with the orthogonality advantages of Demmler-Reinsch spline bases, the resultant fitting and inference is relatively fast and scales well to large data sets.

A simulation study shows that the new Bayesian approaches offer improved performance in terms of classification of effect types as being either zero, linear or non-linear, compared with that of Chouldechova & Hastie (2015). They also shown to perform well in comparison with the Bayesian approach of Scheipl *et al.* (2012), but are considerably faster.

The R package that accompanies this article’s methodology is named `gamselBayes` (He & Wand, 2023). In Section 4 we compare its performance with two other R packages: `gamsel` (Chouldechova & Hastie, 2022) and `spikeSlabGAM` (Scheipl, 2022) which also provide three-category model selection for generalized additive models. Note that there are many other R packages concerned with generalized additive model analysis, some of which employ versions of the LASSO-type approach used by `gamselBayes`. Examples of such packages are `BayesX` (Umlauf, Kneib & Klein, 2023), `bamlss` (Umlauf *et al.*, 2023) and `bmrs` (Bürkner, 2022).

Descriptions of our models and their conversion to computation-friendly forms are given in Section 2. Algorithms for practical fitting and model selection are listed in Section 3. We also point to the R package, `gamselBayes`, that allows easy and immediate access to the new methodology for users of the R language. Section 4 assesses performance of the new approaches in comparison with existing approaches with having similar aims. Applications to actual data are illustrated in Section 5. We close with some concluding remarks in Section 6.

2 Model Description

The original input data are as follows:

$$(\overset{\circ}{\mathbf{x}}_i^{\text{orig}}, \overset{\bullet}{\mathbf{x}}_i^{\text{orig}}, y_i^{\text{orig}}), \quad 1 \leq i \leq n,$$

where, for each $1 \leq i \leq n$, $\overset{\circ}{\mathbf{x}}_i^{\text{orig}}$ denotes a $d_{\circ} \times 1$ vector of predictors that can only enter the model linearly (e.g. binary predictors) and $\overset{\bullet}{\mathbf{x}}_i^{\text{orig}}$ denotes a $d_{\bullet} \times 1$ vector of continuous predictors that can enter the model either linearly or non-linearly. For Bayesian fitting and inference we work with standardized versions of the data. This has advantages such as the methodology being independent of units of measurement for fixed hyperparameter settings and improved numerical stability. Algorithm 1 in Section 3.1 provides the operational details of the standardization process. The full data to be used for fitting and model selection are

$$(\overset{\circ}{\mathbf{x}}_i, \overset{\bullet}{\mathbf{x}}_i, y_i), \quad 1 \leq i \leq n,$$

where $\overset{\circ}{\mathbf{x}}_i$ and $\overset{\bullet}{\mathbf{x}}_i$ are standardized data versions of $\overset{\circ}{\mathbf{x}}_i^{\text{orig}}$ and $\overset{\bullet}{\mathbf{x}}_i^{\text{orig}}$. Also, for the continuous response case the y_i are the standardized response data. In the binary response case the y_i are not pre-processed and remain as values in $\{0, 1\}$. For each $1 \leq i \leq n$ let

$$\overset{\circ}{x}_{ji} \equiv \text{the } j\text{th entry of } \overset{\circ}{\mathbf{x}}_i, \quad 1 \leq j \leq d_{\circ} \quad \text{and} \quad \overset{\bullet}{x}_{ji} \equiv \text{the } j\text{th entry of } \overset{\bullet}{\mathbf{x}}_i, \quad 1 \leq j \leq d_{\bullet}.$$

Generalized additive models involve linear predictors η_i , $1 \leq i \leq n$, having the generic forms

$$\eta_i \equiv \beta_0 + \sum_{j=1}^{d_{\circ}} \beta_j \overset{\circ}{x}_{ji} + \sum_{j=1}^{d_{\bullet}} f_j(\overset{\bullet}{x}_{ji}), \quad (1)$$

where $\beta_0, \dots, \beta_{d_{\circ}}$ are the coefficients of linear components and the f_j are smooth real-valued functions over an interval containing the $\overset{\bullet}{x}_{ji}$ data.

2.1 Matrix Notation

For any column vector \mathbf{a} we let $\|\mathbf{a}\| \equiv (\mathbf{a}^T \mathbf{a})^{1/2}$ denote the Euclidean norm of \mathbf{a} and \mathbf{a}_{-j} denote the column vector with the j th entry of \mathbf{a} omitted. If \mathbf{b} is a column vector having the same number of rows as \mathbf{a} then $\mathbf{a} \odot \mathbf{b}$ and \mathbf{a}/\mathbf{b} are, respectively, the column vectors formed from \mathbf{a} and \mathbf{b} by obtaining element-wise products and quotients. For any square matrix \mathbf{A} we let $\text{diagonal}(\mathbf{A})$ denote the column vector containing the diagonal entries of \mathbf{A} .

2.2 Distributional Definitions

Table 1 lists all distributions used in this article. In particular, the parametrizations of the corresponding density functions and probability functions are provided. In this table, and throughout this article, $\Gamma(x) \equiv \int_0^\infty t^{x-1} e^{-t} dt$ is the gamma function and Φ denotes the $N(0, 1)$ cumulative distribution function.

distribution	density/probability function in x	abbreviation
Bernoulli	$\wp^x (1 - \wp)^{1-x}; \quad x = 0, 1; 0 < \wp < 1$	Bernoulli(\wp)
Multivariate Normal	$ 2\pi\Sigma ^{-1/2} \exp\{-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \Sigma^{-1}(\mathbf{x} - \boldsymbol{\mu})\}$	$N(\boldsymbol{\mu}, \Sigma)$
Inverse Gamma	$\frac{\lambda^\kappa x^{-\kappa-1} e^{-\lambda/x}}{\Gamma(\kappa)}; \quad x > 0; \kappa, \lambda > 0$	Inverse-Gamma(κ, λ)
Inverse Gaussian	$\frac{\sqrt{\lambda} \exp\left\{\frac{-\lambda(x - \mu)^2}{2\mu^2 x}\right\}}{\sqrt{2\pi x^3}}; \quad x > 0; \mu, \lambda > 0$	Inverse-Gaussian(μ, λ)
Beta	$\frac{\Gamma(\alpha + \beta)x^{\alpha-1}(1-x)^{\beta-1}}{\Gamma(\alpha)\Gamma(\beta)}; \quad 0 < x < 1; \alpha, \beta > 0$	Beta(α, β)
Half-Cauchy	$\frac{2}{\pi\sigma((x/\sigma)^2 + 1)}; \quad x > 0; \sigma > 0$	Half-Cauchy(σ)
Truncated-Normal ₊	$\frac{\exp\{-(x - \mu)^2/(2\sigma^2)\}}{\Phi(\mu/\sigma)\sqrt{2\pi\sigma^2}}; \quad x > 0; \sigma > 0$	Truncated-Normal ₊ (μ, σ^2)

Table 1: Distributions used in this article and their corresponding density/probability functions.

2.3 Model for a Smooth Function

Let $\dot{x}_1, \dots, \dot{x}_n$ be a typical continuous predictor data sample. The corresponding smooth function model takes the form

$$f(\dot{x}_i) \equiv \beta \dot{x}_i + \sum_{k=1}^K u_k z_k(\dot{x}_i), \quad 1 \leq i \leq n, \quad (2)$$

for coefficients β and $\mathbf{u} \equiv (u_1, \dots, u_K)$. Here $\{z_k(\cdot) : 1 \leq k \leq K\}$ is an appropriate spline basis over an interval containing the \dot{x}_i data. In accordance with the set-up of Chouldechova

& Hastie (2015), we choose the spline basis to have orthogonality properties and lead to computational speed-ups. These properties can be explained succinctly in matrix algebraic terms. Define

$\dot{\mathbf{x}} \equiv$ the $n \times 1$ vector with i th entry \dot{x}_i and $\mathbf{Z} \equiv$ the $n \times K$ matrix having (i, k) entry $z_k(\dot{x}_i)$.

Then we construct \mathbf{Z} to satisfy

$$\mathbf{Z}^T \mathbf{1}_n = \mathbf{Z}^T \dot{\mathbf{x}} = \mathbf{0}_K \quad \text{and} \quad \mathbf{Z}^T \mathbf{Z} \text{ is a diagonal matrix.} \quad (3)$$

Spline bases satisfying (3) are referred to as having a *Demmler-Reinsch* form. In addition, we scale the columns of \mathbf{Z} so that the right-hand side of (2) has mixed model representations of the form

$$\dot{\mathbf{x}}\beta + \mathbf{Z}\mathbf{u} \quad \text{where } \mathbf{u} \text{ is a random vector having density function } \mathfrak{p}(\mathbf{u}) = h(\|\mathbf{u}\|) \quad (4)$$

for some scalar-valued function h . In other words, we apply linear transformations to ensure that the distribution of \mathbf{u} has spherical, rather than ellipsoidal, contours. For ordinary generalized additive model fitting, as opposed to selection, the most common choice of h is $h(x) = (2\pi\sigma_u^2)^{-K/2} \exp\{-x^2/(2\sigma_u^2)\}$, for some $\sigma_u > 0$, which corresponds to the spline coefficients model taking the form

$$\mathbf{u}|\sigma_u^2 \sim N(0, \sigma_u^2 \mathbf{I}). \quad (5)$$

For the generalized additive model selection, (5) should be replaced by an appropriate sparse signal prior distribution. Section 2.5 provides full details on this modelling aspect.

There are various ways in which $\{z_k(\cdot) : 1 \leq k \leq K\}$ can be set up so that (3) and (4) are satisfied. In this article we follow the constructions laid out in Section 4 of Wand & Ormerod (2008) and Algorithm 1 of Ngo & Wand (2004). The full details are given in Section S.1. of the supplement. We use the descriptor *canonical Demmler-Reinsch basis* for this type of spline basis.

2.4 Model for a Linear Coefficient

Let β denote a generic linear coefficient. We impose the following family of distributions on β :

$$\mathfrak{p}(\beta|\sigma_\beta, \rho_\beta) = \rho_\beta(2\sigma_\beta)^{-1} \exp(-|\beta|/\sigma_\beta) + (1 - \rho_\beta)\delta_0(\beta) \quad (6)$$

for parameters $\sigma_\beta > 0$ and $0 \leq \rho_\beta \leq 1$. Here δ_0 denotes the Dirac delta function at zero. We call (6) the *Laplace-Zero* family of distributions, since it is a “spike-and-slab” mixture of a Laplace density function and a point mass at zero. (e.g. Mitchell & Beauchamp, 1988).

The $\rho_\beta = 1$ version of (6) corresponds to the Bayesian Lasso approach of Park & Casella (2008). However, as pointed out there, Bayes estimation does not lead to sparse fits for the purely Laplace prior situation. The addition of a point mass at zero has the attraction of posterior distributions also having this feature and sparse Bayes-type fits. This aspect is exploited in Section 3.5 for principled model selection strategies.

The scale parameter in (6) has the prior distribution:

$$\sigma_\beta \sim \text{Half-Cauchy}(s_\beta)$$

for a hyperparameter $s_\beta > 0$. Gelman (2006) provides justification for the imposition of a Half Cauchy prior on scale parameters such as σ_β . The mixture parameter ρ_β is treated as a hyperparameter.

Many alternatives to (6) for Bayesian model selection have been proposed and studied. The overarching goal is the achievement of sparse solutions, as is the case for frequentist LASSO-type approaches, according to Bayesian fitting paradigms. The most common approach is to use “spike-and-slab” priors, for which (6) is a special case, and involves mixing a symmetric zero mean continuous random variable with either a point mass at zero or another continuous

random variable that is highly concentrated around zero. Key references include Lempers (1971), Mitchell & Beauchamp (1988), George & McCulloch (1993) and Ishwaran & Rao (2005). Alternative approaches involve a single continuous distributional form, rather than a mixture, that is sharply peaked at the origin and heavy-tailed. Examples include Park & Casella (2008), Carvalho *et al.* (2010) and Griffin & Brown (2011). Bhadra *et al.* (2019) compare and contrast both types of approaches.

2.5 Model for a Spline Coefficients Vector

Let \mathbf{u} denote a $K \times 1$ spline coefficient vector. We impose the following family of distributions on \mathbf{u} :

$$\mathbf{p}(\mathbf{u}|\sigma_u, \rho_u) = \rho_u (C_K \sigma_u)^{-1} \exp(-\|\mathbf{u}\|/\sigma_u) + (1 - \rho_u) \delta_0(\mathbf{u}) \quad (7)$$

for parameters $\sigma_u > 0$ and $0 \leq \rho_u \leq 1$ and with $C_K \equiv 2^K \pi^{(K-1)/2} \Gamma(\frac{1}{2}(K+1))$. Here δ_0 denotes the K -variate Dirac delta function at $\mathbf{0}_K$, the $K \times 1$ vector of zeroes.

Kyung *et al.* (2010) use the phrase *group lasso* for the family of priors defined by (7) in the $\rho_u = 1$ special case. This naming is due to the group LASSO methodology of Yuan & Lin (2006). The essence of Yuan & Lin's (2006) extension of the ordinary LASSO is that particular vectors coefficients, $\boldsymbol{\theta}$ say, are treated together as an entity and penalty terms of the form $\lambda \|\boldsymbol{\theta}\|$, for some $\lambda > 0$, allow for all entries of $\boldsymbol{\theta}$ to be estimated as exactly zero. In their frequentist approach to generalized additive model selection Chouldechova & Hastie (2015) apply this idea to vectors of spline coefficients, denoted in this section by \mathbf{u} . This allows for smooth function effects to be categorized as either linear or non-linear depending on whether $\hat{\mathbf{u}} = \mathbf{0}$ or $\hat{\mathbf{u}} \neq \mathbf{0}$, where $\hat{\mathbf{u}}$ is an estimate of \mathbf{u} . In keeping with (6), we extend the group lasso distribution to a K -variate "spike-and-slab" form. Note that (7) has a point mass at $\mathbf{0}_K$, the K -vector of zeroes.

The scale parameter has the following prior distributions:

$$\sigma_u \sim \text{Half-Cauchy}(s_u)$$

for hyperparameter $s_u > 0$. The prior distribution justification given at the end of Section 2.4 also applies here. The mixture parameter ρ_u is a user-specified hyperparameter.

2.6 Hyperparameter Default Values

The standardization of the input data invokes scale invariance and justifies setting the hyperparameters to fixed constant values. With noninformativity in mind, our recommended default values of the hyperparameters are:

$$\sigma_{\beta_0} = 10^5, s_\beta = s_\varepsilon = s_u = 1000, \rho_\beta = \rho_u = \frac{1}{2}.$$

These values are used in the upcoming numerical studies and examples.

2.7 Auxiliary Variable Representations

Distributional specifications such as (6) and (7) are not amenable to Markov chain Monte Carlo and mean field variational Bayes fitting algorithms due to their non-standard full conditional distributions. In this subsection we re-express them using auxiliary variables, which are tailored so that all full conditional distributions have standard forms.

First, note that $\sigma \sim \text{Half-Cauchy}(s)$ is equivalent to

$$\sigma^2 | a \sim \text{Inverse-Gamma}(\frac{1}{2}, 1/a), \quad a \sim \text{Inverse-Gamma}(\frac{1}{2}, 1/s).$$

For the case of (6) we introduce auxiliary variables $\gamma_\beta, \tilde{\beta}$ and b_β and re-define β such that

$$\beta = \gamma_\beta \tilde{\beta}, \quad \gamma_\beta \sim \text{Bernoulli}(\rho_\beta), \quad \tilde{\beta} | b_\beta, \sigma_\beta^2 \sim N(0, \sigma_\beta^2 / b_\beta) \quad \text{and} \quad b_\beta \sim \text{Inverse-Gamma}(1, \frac{1}{2}). \quad (8)$$

Then standard distributional manipulations can be used to show that (8) is equivalent to (6). Similarly, with the introduction of the random variable γ_u , (7) is equivalent to

$\mathbf{u} = \gamma_u \tilde{\mathbf{u}}$, $\gamma_u \sim \text{Bernoulli}(\rho_u)$, $\tilde{\mathbf{u}}|b_u, \sigma_u^2 \sim N(0, \sigma_u^2 \mathbf{I}/b_u)$ and $b_u \sim \text{Inverse-Gamma}(\frac{K+1}{2}, \frac{1}{2})$ courtesy of a result provided in Section 3.1 of Kyung *et al.* (2010) for the $\rho_u = 1$ case.

2.8 The Full Gaussian Response Model

Consider, first, the case where Gaussianity of the y_i s is reasonably assumed. Suppose that we apply the modelling structures of Sections 2.3–2.5 across each of d_o entries of the $\hat{\mathbf{x}}_i$ and d_\bullet entries of $\hat{\mathbf{x}}_i$. Let $\boldsymbol{\beta}$ be the $(d_o + d_\bullet) \times 1$ vector containing all of the linear term coefficients and $\mathbf{u}_1, \dots, \mathbf{u}_{d_\bullet}$ be the full set of spline coefficient vectors, where \mathbf{u}_j has dimension $K_j \times 1$. Also, apply the auxiliary variable representations of Section 2.7. The resultant full model is:

$$\begin{aligned} \mathbf{y} | \beta_0, \boldsymbol{\gamma}_\beta, \tilde{\boldsymbol{\beta}}, \gamma_{u1}, \dots, \gamma_{ud_\bullet}, \tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_{d_\bullet}, \sigma_\varepsilon^2 &\sim \\ N \left(\mathbf{1}_n \beta_0 + \mathbf{X}(\boldsymbol{\gamma}_\beta \odot \tilde{\boldsymbol{\beta}}) + \sum_{j=1}^{d_\bullet} \mathbf{Z}_j(\gamma_{uj} \tilde{\mathbf{u}}_j), \sigma_\varepsilon^2 \mathbf{I}_n \right), & \beta_0 \sim N(0, \sigma_{\beta_0}^2), \\ \sigma_\varepsilon^2 | a_\varepsilon &\sim \text{Inverse-Gamma}(\frac{1}{2}, 1/a_\varepsilon), \quad a_\varepsilon \sim \text{Inverse-Gamma}(\frac{1}{2}, 1/s_\varepsilon^2), \\ \gamma_{\beta j} &\stackrel{\text{ind.}}{\sim} \text{Bernoulli}(\rho_\beta), \quad \tilde{\beta}_j | \sigma_\beta^2, b_{\beta j} \stackrel{\text{ind.}}{\sim} N(0, \sigma_\beta^2/b_{\beta j}), \quad 1 \leq j \leq d_o + d_\bullet, \\ b_{\beta j} &\stackrel{\text{ind.}}{\sim} \text{Inverse-Gamma}(1, \frac{1}{2}), \quad 1 \leq j \leq d_o + d_\bullet, \quad \gamma_{uj} \stackrel{\text{ind.}}{\sim} \text{Bernoulli}(\rho_u), \quad 1 \leq j \leq d_\bullet, \\ \tilde{\mathbf{u}}_j | \sigma_{uj}^2, b_{uj} &\stackrel{\text{ind.}}{\sim} N(\mathbf{0}, (\sigma_{uj}^2/b_{uj}) \mathbf{I}_{K_j}), \quad b_{uj} \stackrel{\text{ind.}}{\sim} \text{Inverse-Gamma}(\frac{1}{2}(K_j + 1), \frac{1}{2}), \quad 1 \leq j \leq d_\bullet, \\ \sigma_{\beta}^2 | a_\beta &\sim \text{Inverse-Gamma}(\frac{1}{2}, 1/a_\beta), \quad a_\beta \sim \text{Inverse-Gamma}(\frac{1}{2}, 1/s_\beta^2), \\ \sigma_{uj}^2 | a_{uj} &\stackrel{\text{ind.}}{\sim} \text{Inverse-Gamma}(\frac{1}{2}, 1/a_{uj}), \quad a_{uj} \stackrel{\text{ind.}}{\sim} \text{Inverse-Gamma}(\frac{1}{2}, 1/s_u^2), \quad 1 \leq j \leq d_\bullet. \end{aligned} \tag{9}$$

In (9) we have $\boldsymbol{\gamma}_\beta \equiv [\gamma_{\beta 1}, \dots, \gamma_{\beta, d_o + d_\bullet}]^T$. Here, and elsewhere, the notation $\stackrel{\text{ind.}}{\sim}$ is an abbreviation for “distributed independently as”.

The full set of hyperparameters in (9) is:

$$\sigma_{\beta_0}, s_\beta, s_\varepsilon, s_u > 0 \quad \text{and} \quad 0 \leq \rho_\beta, \rho_u \leq 1.$$

Figure 1 shows the directed acyclic graph corresponding to (9).

2.9 Adjustment for Binary Responses

Now suppose that the y_i values are binary rather than continuous. Then an appropriate adjustment to (9) is that where the likelihood is changed to

$$\begin{aligned} y_i | \beta_0, \boldsymbol{\gamma}_\beta, \tilde{\boldsymbol{\beta}}, \gamma_{u1}, \dots, \gamma_{ud_\bullet}, \tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_{d_\bullet} \\ \stackrel{\text{ind.}}{\sim} \text{Bernoulli} \left(\Phi \left(\beta_0 + \left(\mathbf{X}(\boldsymbol{\gamma}_\beta \odot \tilde{\boldsymbol{\beta}}) + \sum_{j=1}^{d_\bullet} \mathbf{Z}_j(\gamma_{uj} \tilde{\mathbf{u}}_j) \right)_i \right) \right). \end{aligned} \tag{10}$$

Following Albert & Chib (1993), we introduce auxiliary random variables c_1, \dots, c_n such that

$$y_i = 1 \quad \text{if and only if} \quad c_i \geq 0 \tag{11}$$

and impose the following conditional distribution on $\mathbf{c} \equiv (c_1, \dots, c_n)$:

$$\begin{aligned} \mathbf{c} | \beta_0, \boldsymbol{\gamma}_\beta, \tilde{\boldsymbol{\beta}}, \gamma_{u1}, \dots, \gamma_{ud_\bullet}, \tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_{d_\bullet} \sim \\ N \left(\mathbf{1}_n \beta_0 + \mathbf{X}(\boldsymbol{\gamma}_\beta \odot \tilde{\boldsymbol{\beta}}) + \sum_{j=1}^{d_\bullet} \mathbf{Z}_j(\gamma_{uj} \tilde{\mathbf{u}}_j), \mathbf{I}_n \right). \end{aligned} \tag{12}$$




Figure 1: Directed acyclic graph representation of Bayesian model (9). Random variables and vectors are shown as larger open circles, with shading indicating to the observed response data. The small closed circles are user-specified hyperparameters.




Figure 2: Sub-graph of the directed acyclic graph for the binary response adjustment to (9). This graph is the same as that shown in Figure 1 except for locations near the response variables node. The new graph has the following modifications: (1) the σ_ϵ^2 and a_ϵ nodes are absent, (2) a hidden node c corresponding to the Albert-Chib auxiliary variables is added to the position held by y in the Gaussian response graph and the binary response observed data node y is a child of c .

The prior distributions on $\beta_0, \gamma_\beta, \tilde{\beta}, \gamma_{u1}, \dots, \gamma_{ud.}$ and $\tilde{u}_1, \dots, \tilde{u}_{d.}$ are the same as in the Gaussian response case. The error variance variables σ_ϵ^2 and a_ϵ are not present for binary responses. Therefore, our binary response model is a modification of (9) for which the y distributional specification is replaced by (11) and (12). Figure 2 shows this modification in graphical terms.

3 Practical Fitting and Model Selection

Practical generalized additive model selection based on the models described in Section 2 requires approximation of the posterior distributions of each of the hidden nodes (unshaded circles) in Figures 1 and 2. The problem reduces to approximation conditional marginalization of directed acyclic graphs. The most accurate practical approach is Markov chain Monte Carlo (e.g. Gelfand & Smith, 1990). For the Gaussian response model (9) and its binary response adjustment described in Section 2.9, Section 3.3 provides full algorithmic details for Markov chain Monte Carlo-based approximate conditional marginalization. A faster, but less accurate, alternative is mean field variational Bayes (e.g. Wainwright & Jordan, 2008). To facilitate scalability to very large data sets, we also provide a variational approximate conditional marginalization algorithm in Section 3.4. Both approaches have steps that depend on the data only through particular sufficient statistic quantities. Therefore, there are considerable speed gains from computing and storing these quantities as part of a pre-preprocessing phase.

3.1 Pre-Processing and Storage of Key Matrices

Algorithm 1 is an important part of our overall strategy for fitting our Bayesian generalized additive models in a stable and efficient manner. The first steps involve standardizing the input data and storing the linear transformation parameters to allow conversion of the final results to the original units. Then design matrices denoted by \mathbf{X} and \mathbf{Z} are computed, with the latter containing all required spline basis functions of the transformed predictor data. Lastly, sufficient statistic matrices such as $\mathbf{X}^T \mathbf{y}$ and $\mathbf{Z}^T \mathbf{Z}$ are computed and stored – ready for use in the upcoming Markov chain Monte Carlo and variational algorithms.

In Algorithm 1 and the upcoming discussion and algorithms we use the identifiers:

$$\mathbf{XTy}, \mathbf{XTX}, \mathbf{XTy}, \mathbf{ZTy}, \mathbf{ZTX} \quad \text{and} \quad \mathbf{ZTZ}$$

for storage of the sufficient statistic quantities $\mathbf{X}^T \mathbf{y}$, $\mathbf{X}^T \mathbf{X}$, $\mathbf{X}^T \mathbf{y}$, $\mathbf{Z}^T \mathbf{y}$, $\mathbf{Z}^T \mathbf{X}$ and $\mathbf{Z}^T \mathbf{Z}$.

3.2 Notation Used in the Fitting Algorithms

For the main fitting algorithms it is useful to have the following definitions in place:

$$\begin{aligned} K_j &\equiv \text{the number of columns in } \mathbf{Z}_j, \quad 1 \leq j \leq d_\bullet, \\ \mathbf{c} &\text{ is the } (d_\bullet + 1) \times 1 \text{ vector with entries } \mathbf{c}_1 \equiv 0 \text{ and } \mathbf{c}_{j+1} \equiv \sum_{k=1}^j K_k, \quad 1 \leq j \leq d_\bullet, \\ \mathbf{ZTy}^{(j)} &\equiv \text{the sub-block of } \mathbf{ZTy} \text{ corresponding to rows } (\mathbf{c}_j + 1) \text{ to } \mathbf{c}_{j+1}, \quad 1 \leq j \leq d_\bullet, \\ \mathbf{ZTX}^{(j)} &\equiv \text{the sub-block of } \mathbf{ZTX} \text{ corresponding to rows } (\mathbf{c}_j + 1) \text{ to } \mathbf{c}_{j+1}, \quad 1 \leq j \leq d_\bullet, \\ \mathbf{ZTZ}^{(j,j')} &\equiv \text{the sub-block of } \mathbf{ZTZ} \text{ corresponding to rows } (\mathbf{c}_j + 1) \text{ to } \mathbf{c}_{j+1} \\ &\quad \text{and columns } (\mathbf{c}_{j'} + 1) \text{ to } \mathbf{c}_{j'+1}, \quad 1 \leq j, j' \leq d_\bullet. \end{aligned} \tag{13}$$

Note that, according to the notation in (13),

$$\mathbf{ZTX}^{(j)} = \mathbf{Z}_j^T \mathbf{X} \quad \text{and} \quad \mathbf{ZTZ}^{(j,j')} = \mathbf{Z}_j^T \mathbf{Z}_{j'}.$$

The updates in approximate inference iterative algorithms, presented in Sections 3.3 and 3.4, depend on particular columns and rows of the matrices listed in (13). These will be specified using the following notational convention: \mathbf{e}_r is a column vector of appropriate length with r th entry equal to 1 and zeroes elsewhere. For example, the j th column of \mathbf{XTX} is $\mathbf{XTX} \mathbf{e}_j$ where \mathbf{e}_j is the $(d_\circ + d_\bullet) \times 1$ vector with j th entry 1 and 0 elsewhere. Implementations of the upcoming algorithms normally would not require explicit calculation and storage of \mathbf{e}_r vectors and, instead, array subsetting code specific to the programming language can be used.

Algorithm 1 *Pre-processing of original data and creation of key matrices for input into Bayesian generalized additive model algorithms.*

Inputs: $\mathbf{y}^{\text{orig}}(n \times 1)$; $\hat{\mathbf{x}}_j^{\text{orig}}(n \times 1), 1 \leq j \leq d_o$; $\hat{\mathbf{x}}_j^{\text{orig}}(n \times 1), 1 \leq j \leq d_\bullet$
 $\text{mean}(\mathbf{y}^{\text{orig}}) \leftarrow$ sample mean of \mathbf{y}^{orig} ; $\text{st.dev}(\mathbf{y}^{\text{orig}}) \leftarrow$ sample standard dev'n of \mathbf{y}^{orig}
If \mathbf{y}^{orig} is continuous then $\mathbf{y} \leftarrow \{\mathbf{y}^{\text{orig}} - \text{mean}(\mathbf{y}^{\text{orig}})\mathbf{1}_n\}/\text{st.dev}(\mathbf{y}^{\text{orig}})$
If \mathbf{y}^{orig} is binary then $\mathbf{y} \leftarrow \mathbf{y}^{\text{orig}}$
For $j = 1, \dots, d_o$:
 $\text{mean}(\hat{\mathbf{x}}_j^{\text{orig}}) \leftarrow$ sample mean of $\hat{\mathbf{x}}_j^{\text{orig}}$; $\text{st.dev.}(\hat{\mathbf{x}}_j^{\text{orig}}) \leftarrow$ sample standard dev'n of $\hat{\mathbf{x}}_j^{\text{orig}}$
 $\hat{\mathbf{x}}_j \leftarrow \{\hat{\mathbf{x}}_j^{\text{orig}} - \text{mean}(\hat{\mathbf{x}}_j^{\text{orig}})\mathbf{1}_n\}/\text{st.dev.}(\hat{\mathbf{x}}_j^{\text{orig}})$
For $j = 1, \dots, d_\bullet$:
 $\text{mean}(\hat{\mathbf{x}}_j^{\text{orig}}) \leftarrow$ sample mean of $\hat{\mathbf{x}}_j^{\text{orig}}$; $\text{st.dev.}(\hat{\mathbf{x}}_j^{\text{orig}}) \leftarrow$ sample standard dev'n of $\hat{\mathbf{x}}_j^{\text{orig}}$
 $\hat{\mathbf{x}}_j \leftarrow \{\hat{\mathbf{x}}_j^{\text{orig}} - \text{mean}(\hat{\mathbf{x}}_j^{\text{orig}})\mathbf{1}_n\}/\text{st.dev.}(\hat{\mathbf{x}}_j^{\text{orig}})$
 $\mathbf{X} \leftarrow [\hat{\mathbf{x}}_1 \cdots \hat{\mathbf{x}}_{d_o} \hat{\mathbf{x}}_1 \cdots \hat{\mathbf{x}}_{d_\bullet}]$
For $j = 1, \dots, d_\bullet$:
 $\mathbf{Z}_j \leftarrow n \times K_j$ matrix containing the canonical Demmler-Reinsch basis for the predictor data vector $\hat{\mathbf{x}}_j$, using the construction described in Section S.1 of the supplement
 $\mathbf{Z} \leftarrow [\mathbf{Z}_1 \cdots \mathbf{Z}_{d_\bullet}]$; $\mathbf{XTy} \leftarrow \mathbf{X}^T \mathbf{y}$; $\mathbf{XTX} \leftarrow \mathbf{X}^T \mathbf{X}$; $\mathbf{ZTy} \leftarrow \mathbf{Z}^T \mathbf{y}$
 $\mathbf{ZTX} \leftarrow \mathbf{Z}^T \mathbf{X}$; $\mathbf{ZTZ} \leftarrow \mathbf{Z}^T \mathbf{Z}$
Outputs: $\mathbf{y}, \mathbf{X}, \mathbf{Z}_1, \dots, \mathbf{Z}_{d_\bullet}, \mathbf{XTy}, \mathbf{XTX}, \mathbf{ZTy}, \mathbf{ZTX}, \mathbf{ZTZ}, \text{mean}(\mathbf{y}^{\text{orig}}), \text{st.dev}(\mathbf{y}^{\text{orig}}),$
 $\left\{ (\text{mean}(\hat{\mathbf{x}}_j^{\text{orig}}), \text{st.dev.}(\hat{\mathbf{x}}_j^{\text{orig}})) : 1 \leq j \leq d_o \right\}, \left\{ (\text{mean}(\hat{\mathbf{x}}_j^{\text{orig}}), \text{st.dev.}(\hat{\mathbf{x}}_j^{\text{orig}})) : 1 \leq j \leq d_\bullet \right\}$

However, for algorithm listing use of the e_r notation has the advantage of avoiding further subscripting.

To allow the Gaussian and Bernoulli response cases to be handled together we also use the notation $\mathbf{yT1}_{\text{adj}}, \mathbf{XTy}_{\text{adj}}$ and $\mathbf{ZTy}_{\text{adj}}$. These are adjustments of $\mathbf{yT1}, \mathbf{XTy}$ and \mathbf{ZTy} in which the \mathbf{y} vector is replaced by \mathbf{c} : the Albert-Chib auxiliary variables vector that arises in the Bernoulli response case. The notation of (13) for extraction of sub-blocks of \mathbf{ZTy} also applies to $\mathbf{ZTy}_{\text{adj}}$.

The main algorithms also uses the following functions:

$$\text{logit}(x) \equiv \log \left(\frac{x}{1-x} \right), \quad \text{expit}(x) \equiv \text{logit}^{-1}(x) = \frac{1}{1 + \exp(-x)} \quad \text{and} \quad \zeta(x) = \log\{2\Phi(x)\}$$

where, as before, Φ is the $N(0, 1)$ cumulative distribution function. It follows that $\zeta'(x) = \phi(x)/\Phi(x)$, where ϕ is the $N(0, 1)$ density function, which arises in Algorithm 3. Stable computation of $\zeta'(x)$ when x is a large negative number is not straightforward. Azzalini (2023) and Wand & Ormerod (2012), for example, provide practical solutions to this problem. Lastly, an expression of the form $\zeta'(v)$, where v is a column vector, is such that function evaluation is element-wise.

3.3 Markov Chain Monte Carlo

For the Bayesian graphical model (9) and the binary response adjustments given in Section 2.9, determination of each of the full conditional distributions for Markov Chain Monte Carlo sampling is fairly straightforward. Virtually all of the full conditional distributions have standard forms such as Bernoulli, Beta, Inverse Gamma and Multivariate Normal distributions.

Possible exceptions are the Inverse Gaussian and Truncated Normal distributions, but are such that effective solutions are provided, respectively, by Michael *et al.* (1976) and Robert (1995). Therefore, Markov Chain Monte Carlo sampling essentially reduces to Gibbs sampling for the models at hand. Algorithm 2 lists the full set of steps needed to draw samples from the posterior distributions of the model parameters. The fact that most of the draws only require the sufficient statistic matrices from Algorithm 1 means that the sampling can be done quite rapidly regardless of sample size.

Algorithm 2 Markov chain Monte Carlo generation of samples from the posterior distributions of the parameters in (9).

Data Inputs: \mathbf{y} ($n \times 1$); \mathbf{X} ($n \times (d_o + d_\bullet)$); \mathbf{Z}_j ($n \times K_j$), $1 \leq j \leq d_\bullet$.

Response Type Input: **responseType** \in {Gaussian, Bernoulli}.

Sufficient Statistics Inputs: **XTy**, **XTX**, **ZTy**, **ZTX**, **ZTZ**.

Hyperparameter Inputs: $\sigma_{\beta_0}, s_\beta, s_\varepsilon, s_u > 0$, $0 \leq \rho_\beta, \rho_u \leq 1$.

Chain Length Inputs: N_{warm} and N_{kept} , both positive integers.

Initialize: $\gamma_\beta^{[0]} \leftarrow \frac{1}{2} \mathbf{1}_{d_o + d_\bullet}$; $\gamma_{uj}^{[0]} \leftarrow \frac{1}{2}$, $1 \leq j \leq d_\bullet$; $\tilde{\boldsymbol{\beta}}^{[0]} \leftarrow \mathbf{0}_{d_o + d_\bullet}$
 $\tilde{\mathbf{u}}_j^{[0]} \leftarrow \mathbf{0}_{K_j}$, $1 \leq j \leq d_\bullet$; $a_\varepsilon^{[0]} \leftarrow 1$; $(\sigma_\varepsilon^2)^{[0]} \leftarrow 1$; $(\sigma_\beta^2)^{[0]} \leftarrow 1$; $a_\beta^{[0]} \leftarrow 1$
 $\mathbf{b}_\beta^{[0]} \leftarrow \mathbf{1}_{d_o + d_\bullet}$; $b_{uj}^{[0]} \leftarrow 1$, $1 \leq j \leq d_\bullet$
 $a_{uj}^{[0]} \leftarrow 1$, $1 \leq j \leq d_\bullet$; $(\sigma_{uj}^2)^{[0]} \leftarrow 1$, $1 \leq j \leq d_\bullet$.

$\mathbf{yT1}_{\text{adj}} \leftarrow \mathbf{0}$; $\mathbf{XTy}_{\text{adj}} \leftarrow \mathbf{XTy}$; $\mathbf{ZTy}_{\text{adj}} \leftarrow \mathbf{ZTy}$

For $j = 1, \dots, d_\bullet$: $\mathbf{w}_{Z_j} \leftarrow \text{diagonal}(\mathbf{ZTZ}^{(j,j)})$

For $g = 1, \dots, N_{\text{warm}} + N_{\text{kept}}$:

$\omega_1 \leftarrow \mathbf{yT1}_{\text{adj}}$

$\omega_2 \leftarrow \{n/(\sigma_\varepsilon^2)^{[g-1]}\} + (1/\sigma_{\beta_0}^2)$; $\beta_0^{[g]} \sim N\left(\frac{\omega_1}{(\sigma_\varepsilon^2)^{[g-1]}\omega_2}, \frac{1}{\omega_2}\right)$

$\boldsymbol{\Omega} \leftarrow \left(\gamma_\beta^{[g-1]} \gamma_\beta^{[g-1]T}\right) \odot (\mathbf{XTX}) / (\sigma_\varepsilon^2)^{[g-1]} + \text{diag}\left(\mathbf{b}_\beta^{[g-1]}\right) / (\sigma_\beta^2)^{[g-1]}$

$\omega_3 \leftarrow \mathbf{XTy}_{\text{adj}} - \sum_{j=1}^{d_\bullet} \mathbf{ZTX}^{(j)T} (\gamma_{uj}^{[g-1]} \tilde{\mathbf{u}}_j^{[g-1]})$

Decompose $\boldsymbol{\Omega} = \mathbf{U}_\Omega \text{diag}(\mathbf{d}_\Omega) \mathbf{U}_\Omega^T$ where $\mathbf{U}_\Omega \mathbf{U}_\Omega^T = \mathbf{I}$

$\mathbf{z} \sim N(\mathbf{0}, \mathbf{I}) ((d_o + d_\bullet) \times 1)$; $\tilde{\boldsymbol{\beta}}^{[g]} \leftarrow \mathbf{U}_\Omega \left(\frac{\mathbf{U}_\Omega^T \mathbf{z}}{\sqrt{\mathbf{d}_\Omega}} + \frac{\mathbf{U}_\Omega^T (\gamma_\beta^{[g-1]} \odot \omega_3)}{\mathbf{d}_\Omega (\sigma_\varepsilon^2)^{[g-1]}} \right)$

$(\mathbf{b}_\beta^{[g]})_j \sim \text{Inverse-Gaussian}\left(\sigma_\beta^{[g-1]} / |(\tilde{\boldsymbol{\beta}}^{[g]})_j|\right)$, $1 \leq j \leq d_o + d_\bullet$

$(\sigma_\beta^2)^{[g]} \sim \text{Inverse-Gamma}\left(\frac{1}{2}(d_o + d_\bullet + 1), 1/a_\beta^{[g-1]} + \frac{1}{2} \tilde{\boldsymbol{\beta}}^{[g]T} \text{diag}(\mathbf{b}_\beta^{[g]}) \tilde{\boldsymbol{\beta}}^{[g]}\right)$

$a_\beta^{[g]} \sim \text{Inverse-Gamma}\left(1, \{1/(\sigma_\beta^2)^{[g]}\} + (1/s_\beta^2)\right)$; $\boldsymbol{\beta}^{\text{curr}} \leftarrow \gamma_\beta^{[g-1]} \odot \tilde{\boldsymbol{\beta}}^{[g]}$

For $j = 1, \dots, d_\bullet$: $\mathbf{u}_j^{\text{curr}} \leftarrow \gamma_{uj}^{[g-1]} \tilde{\mathbf{u}}_j^{[g-1]}$

For $j = 1, \dots, d_o + d_\bullet$:

$\omega_4 \leftarrow \mathbf{e}_j^T \mathbf{XTy}_{\text{adj}} - (\mathbf{XTX} \mathbf{e}_j)_{-j}^T (\boldsymbol{\beta}^{\text{curr}})_{-j} - \sum_{j'=1}^{d_\bullet} (\mathbf{ZTX}^{(j')})^T \mathbf{e}_j \mathbf{u}_{j'}^{\text{curr}}$

$\omega_5 \leftarrow \text{logit}(\rho_\beta) - \frac{1}{2} \left\{ (\tilde{\boldsymbol{\beta}}_j^{[g]})^2 \mathbf{e}_j^T \mathbf{XTX} \mathbf{e}_j - 2 \tilde{\boldsymbol{\beta}}_j^{[g]} \omega_4 \right\} / (\sigma_\varepsilon^2)^{[g-1]}$

$(\gamma_\beta^{[g]})_j \sim \text{Bernoulli}(\text{expit}(\omega_5))$

continued on a subsequent page ...

Algorithm 2 continued. This is a continuation of the description of this algorithm that commences on a preceding page.

$$\boldsymbol{\beta}^{\text{curr}} \leftarrow \boldsymbol{\gamma}_\beta^{[g]} \odot \tilde{\boldsymbol{\beta}}^{[g]} \quad ; \quad \text{For } j = 1, \dots, d_\bullet: \quad \tilde{\mathbf{u}}_j^{\text{curr}} \leftarrow \tilde{\mathbf{u}}_j^{[g-1]}$$

For $j = 1, \dots, d_\bullet$:

$$\boldsymbol{\omega}_6 \leftarrow \mathbf{Z}\mathbf{T}\mathbf{y}_{\text{adj}}^{(j)} - \mathbf{Z}\mathbf{T}\mathbf{X}^{(j)}\boldsymbol{\beta}^{\text{curr}} - \sum_{j' \neq j}^{d_\bullet} \mathbf{Z}\mathbf{T}\mathbf{Z}^{(j, j')} (\gamma_{u_{j'}}^{[g-1]} \tilde{\mathbf{u}}_{j'}^{\text{curr}})$$

$$\boldsymbol{\omega}_7 \leftarrow \left\{ \gamma_{u_j}^{[g-1]} \boldsymbol{\omega}_{z_j} / (\sigma_\varepsilon^2)^{[g-1]} \right\} + \left\{ b_{u_j}^{[g-1]} \mathbf{1}_{K_j} / (\sigma_{u_j}^2)^{[g-1]} \right\}$$

$$\mathbf{z} \sim N(\mathbf{0}, \mathbf{I}_{K_j}) \quad ; \quad \tilde{\mathbf{u}}_j^{\text{curr}} \leftarrow (\mathbf{z} / \sqrt{\boldsymbol{\omega}_7}) + [\gamma_{u_j}^{[g-1]} \boldsymbol{\omega}_6 / \{\boldsymbol{\omega}_7 (\sigma_\varepsilon^2)^{[g-1]}\}]$$

$$\text{For } j = 1, \dots, d_\bullet: \quad \tilde{\mathbf{u}}_j^{[g]} \leftarrow \tilde{\mathbf{u}}_j^{\text{curr}}$$

For $j = 1, \dots, d_\bullet$:

$$b_{u_j}^{[g]} \sim \text{Inverse-Gaussian} \left(\sigma_{u_j}^{[g-1]} / \|\tilde{\mathbf{u}}_j^{[g]}\|, 1 \right)$$

$$(\sigma_{u_j}^2)^{[g]} \sim \text{Inverse-Gamma} \left(\frac{1}{2}(K_j + 1), \{1/a_{u_j}^{[g-1]}\} + \frac{1}{2}\|\tilde{\mathbf{u}}_j^{[g]}\|^2 b_{u_j}^{[g]} \right)$$

$$a_{u_j}^{[g]} \sim \text{Inverse-Gamma} \left(1, \{1/(\sigma_{u_j}^2)^{[g]}\} + (1/s_u^2) \right)$$

$$\text{For } j = 1, \dots, d_\bullet: \quad \gamma_{u_j}^{\text{curr}} \leftarrow \gamma_{u_j}^{[g-1]}$$

For $j = 1, \dots, d_\bullet$:

$$\boldsymbol{\omega}_8 \leftarrow \mathbf{Z}\mathbf{T}\mathbf{y}_{\text{adj}}^{(j)} - \mathbf{Z}\mathbf{T}\mathbf{X}^{(j)}\boldsymbol{\beta}^{\text{curr}} - \sum_{j' \neq j}^{d_\bullet} \mathbf{Z}\mathbf{T}\mathbf{Z}^{(j, j')} (\gamma_{u_{j'}}^{\text{curr}} \tilde{\mathbf{u}}_{j'}^{[g]})$$

$$\boldsymbol{\omega}_9 \leftarrow \text{logit}(\rho_u) - \frac{1}{2} \left\{ \boldsymbol{\omega}_{z_j}^T (\tilde{\mathbf{u}}_j^{[g]} \odot \tilde{\mathbf{u}}_j^{[g]}) - 2(\tilde{\mathbf{u}}_j^{[g]})^T \boldsymbol{\omega}_8 \right\} / (\sigma_\varepsilon^2)^{[g-1]}$$

$$\gamma_{u_j}^{\text{curr}} \sim \text{Bernoulli}(\text{expit}(\boldsymbol{\omega}_9))$$

$$\text{For } j = 1, \dots, d_\bullet: \quad \gamma_{u_j}^{[g]} \leftarrow \gamma_{u_j}^{\text{curr}}$$

$$\boldsymbol{\omega}_{10} \leftarrow \mathbf{1}_n \beta_0^{[g]} + \mathbf{X} \left(\boldsymbol{\gamma}_\beta^{[g]} \odot \tilde{\boldsymbol{\beta}}^{[g]} \right) + \sum_{j=1}^{d_\bullet} \mathbf{Z}_j \left(\gamma_{u_j}^{[g]} \tilde{\mathbf{u}}_j^{[g]} \right)$$

If **responseType** is Gaussian then

$$(\sigma_\varepsilon^2)^{[g]} \sim \text{Inverse-Gamma} \left(\frac{1}{2}(n + 1), (1/a_\varepsilon^{[g-1]}) + \frac{1}{2}\|\mathbf{y} - \boldsymbol{\omega}_{10}\|^2 \right)$$

$$a_\varepsilon^{[g]} \sim \text{Inverse-Gamma} \left(1, \{1/(\sigma_\varepsilon^2)^{[g]}\} + (1/s_\varepsilon^2) \right)$$

If **responseType** is Bernoulli then

$$(\sigma_\varepsilon^2)^{[g]} \leftarrow 1$$

For $i = 1, \dots, n$:

$$\omega_{11} \sim \text{Truncated-Normal}_+((2y_i - 1)(\omega_{10})_i, 1) \quad ; \quad c_i \leftarrow (2y_i - 1)\omega_{11}$$

$$\mathbf{y}\mathbf{T}\mathbf{1}_{\text{adj}} \leftarrow \mathbf{1}^T \mathbf{c} \quad ; \quad \mathbf{X}\mathbf{T}\mathbf{y}_{\text{adj}} \leftarrow \mathbf{X}^T \mathbf{c} \quad ; \quad \mathbf{Z}\mathbf{T}\mathbf{y}_{\text{adj}} \leftarrow \mathbf{Z}^T \mathbf{c}$$

Outputs: All chains after omission of the first N_{warm} values.

3.4 Mean Field Variational Bayes

Mean field variational Bayes approximate fitting and inference for (9) involves approximation of the joint posterior density function of the model parameters by a product density form such as

$$\begin{aligned} \mathfrak{p}(\beta_0, \boldsymbol{\gamma}_\beta, \tilde{\boldsymbol{\beta}}, \boldsymbol{\gamma}_u, \tilde{\mathbf{u}}, \mathbf{b}_\beta, \sigma_\beta^2, a_\beta, \mathbf{b}_u, \sigma_u^2, \mathbf{a}_u, \sigma_\varepsilon^2, a_\varepsilon | \mathbf{y}) \\ \approx \mathfrak{q}(\beta_0) \mathfrak{q}(\boldsymbol{\gamma}_\beta) \mathfrak{q}(\tilde{\boldsymbol{\beta}}) \mathfrak{q}(\boldsymbol{\gamma}_u) \mathfrak{q}(\tilde{\mathbf{u}}) \mathfrak{q}(\mathbf{b}_\beta) \mathfrak{q}(\sigma_\beta^2) \mathfrak{q}(a_\beta) \mathfrak{q}(\mathbf{b}_u) \mathfrak{q}(\sigma_u^2) \mathfrak{q}(\mathbf{a}_u) \mathfrak{q}(\sigma_\varepsilon^2) \mathfrak{q}(a_\varepsilon) \end{aligned} \quad (14)$$

where, for example, $\tilde{\mathbf{u}} \equiv (\tilde{u}_1, \dots, \tilde{u}_{d_\bullet})$ and $\mathbf{b}_\beta \equiv (b_{\beta_1}, \dots, b_{\beta_{(d_\bullet + d_\bullet)}})$. There are numerous options for the stringency of the product restriction and the choice involves trade-offs concerning tractability, accuracy and speed. For example, one could contemplate replacing $\mathfrak{q}(\beta_0) \mathfrak{q}(\tilde{\boldsymbol{\beta}}) \mathfrak{q}(\tilde{\mathbf{u}})$

in (14) by $q(\beta_0, \tilde{\beta}, \tilde{\mathbf{u}})$ and improve the accuracy of approximation. However, the more stringent approximation is less tractable. In addition to the product restriction (14) we also impose the product density restrictions:

$$q(\gamma_\beta) = \prod_{j=1}^{d_o+d_\bullet} q(\gamma_{\beta_j}), \quad q(\tilde{\mathbf{u}}) = \prod_{j=1}^{d_\bullet} q(\tilde{\mathbf{u}}_j) \quad \text{and} \quad q(\gamma_u) = \prod_{j=1}^{d_\bullet} q(\gamma_{u_j}). \quad (15)$$

With the product density restrictions in place, we obtain the optimal q -densities by minimising the Kullback-Leibler divergence of the left-hand side of (14) from the right-hand side. The optimal q -density forms can be expressed in terms of the full conditional density functions as given by equation (6) of Ormerod & Wand (2010). The optimal q -density parameters can then be solved via a coordinate ascent algorithm. Since each of the full conditionals in the models at hand have standard forms, the optimal q -density functions are relatively simple and the coordinate ascent updates have closed forms.

The Bayesian graphical model for wavelet regression described in Section 3 of Wand & Ormerod (2011) is similar in nature to the generalized additive selection model (9). Hence, the relevant details on the requisite mean field variational Bayes calculations for (9) can be gleaned from the q -density derivations given in Appendix D of Wand & Ormerod (2011).

Some examples of the resulting optimal q -density forms are:

$$\begin{aligned} q(\tilde{\beta}) \text{ has a } N(\boldsymbol{\mu}_{q(\tilde{\beta})}, \boldsymbol{\Sigma}_{q(\tilde{\beta})}) \text{ density function, and} \\ q(\sigma_\varepsilon^2) \text{ has an Inverse-Gamma}(\kappa_{q(\sigma_\varepsilon^2)}, \lambda_{q(\sigma_\varepsilon^2)}) \text{ density function.} \end{aligned} \quad (16)$$

The optimal Inverse Gamma shape parameter $\kappa_{q(\sigma_\varepsilon^2)}$ has explicit solution $\frac{1}{2}(n+1)$. However, the equations for the optimal values of $\boldsymbol{\mu}_{q(\tilde{\beta})}$, $\boldsymbol{\Sigma}_{q(\tilde{\beta})}$ and $\lambda_{q(\sigma_\varepsilon^2)}$ are interdependent and iteration is required to obtain their optimal values. Algorithm 3 lists the full set of steps required to obtain all q -density parameters, with notation similar to that used in (16) for the other q -density parameters.

A final aspect of Algorithm 3 is determination of good stopping criteria for the coordinate ascent scheme. As is common in the mean field variational Bayes literature we monitor relative increases in the approximate marginal log-likelihood, also known as the evidence lower bound, which we denote by $\log p(\mathbf{y}; q)$. Section S.2 of the supplement contains an explicit expression for the approximate marginal log-likelihood for the Section 2 models under product restrictions (14)–(15).

3.5 Model Selection Strategies

Essential components of our Bayesian generalized additive model selection methodology are rules, based on the posterior distributions of relevant parameters, for deciding whether an effect is zero, linear or non-linear. In practice, either the Markov chain Monte Carlo samples or mean field variational Bayes q -densities are used for approximate posterior-based decision making. However, we will describe our strategies in terms of exact posterior distributions – starting with the zero versus linear effect decision.

3.5.1 Deciding Between an Effect Being Zero or Linear

Let β be a generic regression coefficient attached to one of the \tilde{x}_j predictors. According to our models, $\beta = \gamma_\beta \tilde{\beta}$ where γ_β is binary and $\tilde{\beta}$ is continuous. Therefore

$$P(\beta = 0 | \mathbf{y}) = P(\gamma_\beta = 0 | \mathbf{y}) = 1 - E(\gamma_\beta | \mathbf{y}),$$

and the posterior mean of γ_β can be used to decide between hypotheses $H_0 : \beta = 0$ and $H_1 : \beta \neq 0$. A natural rule is to accept H_0 if and only if

$$P(\beta = 0 | \mathbf{y}) > \frac{1}{2} \quad \text{which is equivalent to} \quad E(\gamma_\beta | \mathbf{y}) \leq \frac{1}{2}.$$

Algorithm 3 Iterative determination of the optimal parameters according to a mean field variational Bayes approximation of the posterior distributions for model (9).

Data Inputs: \mathbf{y} ($n \times 1$); \mathbf{X} ($n \times (d_o + d_\bullet)$); \mathbf{Z}_j ($n \times K_j$), $1 \leq j \leq d_\bullet$.

Response Type Input: **responseType** \in {Gaussian, Bernoulli}.

Sufficient Statistics Inputs: \mathbf{XTy} , \mathbf{XTX} , \mathbf{ZTy} , \mathbf{ZTX} , \mathbf{ZTZ}

Hyperparameter Inputs: $\sigma_{\beta_0}, s_{\beta}, s_{\varepsilon}, s_u > 0$, $0 \leq \rho_{\beta}, \rho_u \leq 1$.

Convergence Criterion Input: $\varepsilon_{\text{toler}}$: a small positive number such as 10^{-8} .

Initialize: $\boldsymbol{\mu}_{q(\gamma_{\beta})} \leftarrow \frac{1}{2} \mathbf{1}_{d_o + d_\bullet}$; $\boldsymbol{\mu}_{q(\tilde{\beta})} \leftarrow \mathbf{0}_{d_o + d_\bullet}$; $\mu_{q(1/a_{\varepsilon})} \leftarrow 1$; $\mu_{q(1/\sigma_{\varepsilon}^2)} \leftarrow 1$

$\mu_{q(1/a_{\beta})} \leftarrow 1$; $\mu_{q(1/\sigma_{\beta}^2)} \leftarrow 1$; $\kappa_{q(\sigma_{\beta}^2)} \leftarrow \frac{1}{2}(d_o + d_\bullet + 1)$; $\kappa_{q(a_{\beta})} \leftarrow 1$

$\kappa_{q(\sigma_{\varepsilon}^2)} \leftarrow \frac{1}{2}(n + 1)$; $\kappa_{q(a_{\varepsilon})} \leftarrow 1$; $\boldsymbol{\mu}_{q(\mathbf{b}_{\beta})} \leftarrow \mathbf{1}_{d_o + d_\bullet}$.

$\mathbf{yT1}_{\text{adj}} \leftarrow \mathbf{0}$; $\mathbf{XTy}_{\text{adj}} \leftarrow \mathbf{XTy}$; $\mathbf{ZTy}_{\text{adj}} \leftarrow \mathbf{ZTy}$

For $j = 1, \dots, d_\bullet$:

$\boldsymbol{\mu}_{q(\tilde{\mathbf{u}}_j)} \leftarrow \mathbf{0}_{K_j}$; $\boldsymbol{\sigma}_{q(\tilde{\mathbf{u}}_j)}^2 \leftarrow \mathbf{1}_{K_j}$; $\mu_{q(\gamma_{u_j})} \leftarrow \frac{1}{2}$

$\mu_{q(1/a_{u_j})} \leftarrow 1$; $\mu_{q(1/\sigma_{u_j}^2)} \leftarrow 1$; $\mu_{q(\mathbf{b}_{u_j})} \leftarrow \mathbf{1}$

$\kappa_{q(\sigma_{u_j}^2)} \leftarrow \frac{1}{2}(K_j + 1)$; $\kappa_{q(a_{u_j})} \leftarrow 1$; $\mathbf{w}_{\mathbf{Z}_j} \leftarrow \text{diagonal}(\mathbf{ZTZ}^{(j,j)})$

Cycle:

$\omega_{12} \leftarrow \mathbf{yT1}_{\text{adj}}$

$\sigma_{q(\beta_0)}^2 \leftarrow 1 / \{n\mu_{q(1/\sigma_{\varepsilon}^2)} + (1/\sigma_{\beta_0}^2)\}$; $\mu_{q(\beta_0)} \leftarrow \sigma_{q(\beta_0)}^2 \mu_{q(1/\sigma_{\varepsilon}^2)} \omega_{12}$

$\boldsymbol{\Omega}_{q(\gamma_{\beta})} \leftarrow \text{diag}\{\boldsymbol{\mu}_{q(\gamma_{\beta})} \odot (\mathbf{1} - \boldsymbol{\mu}_{q(\gamma_{\beta})})\} + \boldsymbol{\mu}_{q(\gamma_{\beta})} \boldsymbol{\mu}_{q(\gamma_{\beta})}^T$

$\boldsymbol{\Sigma}_{q(\tilde{\beta})} \leftarrow \left\{ \mu_{q(1/\sigma_{\varepsilon}^2)} \boldsymbol{\Omega}_{q(\gamma_{\beta})} \odot (\mathbf{X}^T \mathbf{X}) + \mu_{q(1/\sigma_{\beta}^2)} \text{diag}(\boldsymbol{\mu}_{q(\mathbf{b}_{\beta})}) \right\}^{-1}$

$\omega_{13} \leftarrow \mathbf{XTy}_{\text{adj}} - \sum_{j=1}^{d_\bullet} \mathbf{ZTX}^{(j)T} (\mu_{q(\gamma_{u_j})} \boldsymbol{\mu}_{q(\tilde{\mathbf{u}}_j)})$

$\boldsymbol{\mu}_{q(\tilde{\beta})} \leftarrow \mu_{q(1/\sigma_{\varepsilon}^2)} \boldsymbol{\Sigma}_{q(\tilde{\beta})} (\boldsymbol{\mu}_{q(\gamma_{\beta})} \odot \omega_{13})$

$\omega_{14} \leftarrow \boldsymbol{\mu}_{q(\tilde{\beta})} \odot \boldsymbol{\mu}_{q(\tilde{\beta})} + \text{diagonal}(\boldsymbol{\Sigma}_{q(\tilde{\beta})})$; $\boldsymbol{\mu}_{q(\mathbf{b}_{\beta})} \leftarrow \left(\mu_{q(1/\sigma_{\beta}^2)} \omega_{14} \right)^{-1/2}$

$\lambda_{q(\sigma_{\beta}^2)} \leftarrow \mu_{q(1/a_{\beta})} + \frac{1}{2} \boldsymbol{\mu}_{q(\mathbf{b}_{\beta})}^T \omega_{14}$; $\mu_{q(1/\sigma_{\beta}^2)} \leftarrow \kappa_{q(\sigma_{\beta}^2)} / \lambda_{q(\sigma_{\beta}^2)}$

$\lambda_{q(a_{\beta})} \leftarrow \mu_{q(1/\sigma_{\beta}^2)} + s_{\beta}^{-2}$; $\mu_{q(1/a_{\beta})} \leftarrow \kappa_{q(a_{\beta})} / \lambda_{q(a_{\beta})}$

For $j = 1, \dots, d_\bullet$: $\boldsymbol{\mu}_{q(u_j)} \leftarrow \mu_{q(\gamma_{u_j})} \boldsymbol{\mu}_{q(\tilde{\mathbf{u}}_j)}$

For $j = 1, \dots, d_o + d_\bullet$:

$\omega_{15} \leftarrow \mathbf{e}_j^T \mathbf{XTy}_{\text{adj}} - \sum_{j'=1}^{d_\bullet} (\mathbf{ZTX}^{(j')} \mathbf{e}_j)^T \boldsymbol{\mu}_{q(u_{j'})}$

$\omega_{15} \leftarrow \mu_{q(\tilde{\beta}_j)} \omega_{15} - (\mathbf{XTXe}_j)^T_{-j} \left[(\boldsymbol{\mu}_{q(\gamma_{\beta})})_{-j} \odot \left\{ (\boldsymbol{\Sigma}_{q(\tilde{\beta})} \mathbf{e}_j)_{-j} + \mu_{q(\tilde{\beta}_j)} (\boldsymbol{\mu}_{q(\tilde{\beta})})_{-j} \right\} \right]$

$\mu_{q(\gamma_{\beta_j})} \leftarrow \text{expit} \left(\text{logit}(\rho_{\beta}) - \frac{1}{2} \mu_{q(1/\sigma_{\varepsilon}^2)} \left\{ (\mu_{q(\tilde{\beta}_j)}^2 + \sigma_{q(\tilde{\beta}_j)}^2) \mathbf{e}_j^T \mathbf{XTXe}_j - 2\omega_{15} \right\} \right)$

continued on a subsequent page ...

However, in the interests of parsimony, less stringent rules are worth considering. Rather than exclusively thresholding $E(\gamma_{\beta} | \mathbf{y})$ at $\frac{1}{2}$, we also consider a family of rules indexed by a threshold parameter $\tau \in (0, 1)$. After fixing τ our strategy for deciding between an effect being zero or linear is

the effect is zero if $E(\gamma_{\beta} | \mathbf{y}) \leq 1 - \tau$, otherwise the effect is linear.

According to this definition of the threshold parameter, lower values of τ lead to sparser fits.

Algorithm 3 continued. This is a continuation of the description of this algorithm that commences on a preceding page.

For $j = 1, \dots, d_\bullet$: $\boldsymbol{\mu}_{q(u_j)} \leftarrow \mu_{q(\gamma_{u_j})} \boldsymbol{\mu}_{q(\tilde{\mathbf{u}}_j)}$

For $j = 1, \dots, d_\bullet$:

$$\boldsymbol{\omega}_{16} \leftarrow \mathbf{Z} \mathbf{T} \mathbf{y}_{\text{adj}}^{(j)} - \mathbf{Z} \mathbf{T} \mathbf{X}^{(j)} \left(\boldsymbol{\mu}_{q(\gamma_\beta)} \odot \boldsymbol{\mu}_{q(\tilde{\beta})} \right) - \sum_{j' \neq j}^{d_\bullet} \mathbf{Z} \mathbf{T} \mathbf{Z}^{(j, j')} \boldsymbol{\mu}_{q(u_{j'})}$$

$$\boldsymbol{\sigma}_{q(\tilde{\mathbf{u}}_j)}^2 \leftarrow \mathbf{1}_{K_j} / \left\{ \mu_{q(1/\sigma_\varepsilon^2)} \mu_{q(\gamma_{u_j})} \boldsymbol{\omega}_{Zj} + \mu_{q(1/\sigma_{u_j}^2)} \mu_{q(b_{u_j})} \mathbf{1}_{K_j} \right\}$$

$$\boldsymbol{\mu}_{q(\tilde{\mathbf{u}}_j)} \leftarrow \mu_{q(1/\sigma_\varepsilon^2)} \left(\mu_{q(\gamma_{u_j})} \boldsymbol{\omega}_{16} \right) \odot \boldsymbol{\sigma}_{q(\tilde{\mathbf{u}}_j)}^2$$

For $j = 1, \dots, d_\bullet$:

$$\boldsymbol{\omega}_{17} \leftarrow \|\boldsymbol{\mu}_{q(\tilde{\mathbf{u}}_j)}\|^2 + \mathbf{1}_{K_j}^T \boldsymbol{\sigma}_{q(\tilde{\mathbf{u}}_j)}^2 \quad ; \quad \mu_{q(b_{u_j})} \leftarrow \left(\mu_{q(1/\sigma_{u_j}^2)} \boldsymbol{\omega}_{17} \right)^{-1/2}$$

$$\lambda_{q(\sigma_{u_j}^2)} \leftarrow \mu_{q(1/a_{u_j})} + \frac{1}{2} \mu_{q(b_{u_j})} \boldsymbol{\omega}_{17} \quad ; \quad \mu_{q(1/\sigma_{u_j}^2)} \leftarrow \kappa_{q(\sigma_{u_j}^2)} / \lambda_{q(\sigma_{u_j}^2)}$$

$$\lambda_{q(a_{u_j})} \leftarrow \mu_{q(1/\sigma_{u_j}^2)} + (1/s_u^2) \quad ; \quad \mu_{q(1/a_{u_j})} \leftarrow \kappa_{q(a_{u_j})} / \lambda_{q(a_{u_j})}$$

For $j = 1, \dots, d_\bullet$: $\boldsymbol{\mu}_{q(u_j)} \leftarrow \mu_{q(\gamma_{u_j})} \boldsymbol{\mu}_{q(\tilde{\mathbf{u}}_j)}$

For $j = 1, \dots, d_\bullet$:

$$\boldsymbol{\omega}_{18} \leftarrow \mathbf{Z} \mathbf{T} \mathbf{y}_{\text{adj}}^{(j)} - \mathbf{Z} \mathbf{T} \mathbf{X}^{(j)} \left(\boldsymbol{\mu}_{q(\gamma_\beta)} \odot \boldsymbol{\mu}_{q(\tilde{\beta})} \right) - \sum_{j' \neq j}^{d_\bullet} \mathbf{Z} \mathbf{T} \mathbf{Z}^{(j, j')} \boldsymbol{\mu}_{q(u_{j'})}$$

$$\boldsymbol{\omega}_{19} \leftarrow \boldsymbol{w}_{Zj}^T \left(\boldsymbol{\mu}_{q(\tilde{\mathbf{u}}_j)} \odot \boldsymbol{\mu}_{q(\tilde{\mathbf{u}}_j)} + \boldsymbol{\sigma}_{q(\tilde{\mathbf{u}}_j)}^2 \right) - 2 \boldsymbol{\mu}_{q(\tilde{\mathbf{u}}_j)}^T \boldsymbol{\omega}_{18}$$

$$\mu_{q(\gamma_{u_j})} \leftarrow \text{expit} \left(\text{logit}(\rho_u) - \frac{1}{2} \mu_{q(1/\sigma_\varepsilon^2)} \boldsymbol{\omega}_{19} \right)$$

$$\boldsymbol{\omega}_{20} \leftarrow \mathbf{1}_n \mu_{q(\beta_0)} + \mathbf{X} \left(\boldsymbol{\mu}_{q(\gamma_\beta)} \odot \boldsymbol{\mu}_{q(\tilde{\beta})} \right) + \sum_{j=1}^{d_\bullet} \mathbf{Z}_j \left(\mu_{q(\gamma_{u_j})} \boldsymbol{\mu}_{q(\tilde{\mathbf{u}}_j)} \right)$$

If **responseType** is Gaussian then

$$\boldsymbol{\Omega}_{q(\gamma_\beta)} \leftarrow \text{diag} \left\{ \boldsymbol{\mu}_{q(\gamma_\beta)} \odot (\mathbf{1} - \boldsymbol{\mu}_{q(\gamma_\beta)}) \right\} + \boldsymbol{\mu}_{q(\gamma_\beta)} \boldsymbol{\mu}_{q(\gamma_\beta)}^T$$

$$\lambda_{q(\sigma_\varepsilon^2)} \leftarrow \mu_{q(1/a_\varepsilon)} + \frac{1}{2} \|\mathbf{y} - \boldsymbol{\omega}_{20}\|^2 + \frac{1}{2} n \sigma_{q(\beta_0)}^2$$

$$+ \frac{1}{2} \text{tr} \left[\mathbf{X}^T \mathbf{X} \left\{ \boldsymbol{\Omega}_{q(\gamma_\beta)} \odot \left(\boldsymbol{\Sigma}_{q(\tilde{\beta})} + \boldsymbol{\mu}_{q(\tilde{\beta})} \boldsymbol{\mu}_{q(\tilde{\beta})}^T \right) \right\} \right]$$

$$- \frac{1}{2} \text{tr} \left\{ \mathbf{X}^T \mathbf{X} \left(\boldsymbol{\mu}_{q(\gamma_\beta)} \odot \boldsymbol{\mu}_{q(\tilde{\beta})} \right) \left(\boldsymbol{\mu}_{q(\gamma_\beta)} \odot \boldsymbol{\mu}_{q(\tilde{\beta})} \right)^T \right\}$$

$$+ \frac{1}{2} \sum_{j=1}^{d_\bullet} \boldsymbol{w}_{Zj}^T \left(\mu_{q(\gamma_{u_j})} \left[\boldsymbol{\sigma}_{q(\tilde{\mathbf{u}}_j)}^2 + \{1 - \mu_{q(\gamma_{u_j})}\} \boldsymbol{\mu}_{q(\tilde{\mathbf{u}}_j)} \odot \boldsymbol{\mu}_{q(\tilde{\mathbf{u}}_j)} \right] \right)$$

$$\mu_{q(1/\sigma_\varepsilon^2)} \leftarrow \kappa_{q(\sigma_\varepsilon^2)} / \lambda_{q(\sigma_\varepsilon^2)} \quad ; \quad \lambda_{q(a_\varepsilon)} \leftarrow \mu_{q(1/\sigma_\varepsilon^2)} + (1/s_\varepsilon^2) \quad ; \quad \mu_{q(1/a_\varepsilon)} \leftarrow \kappa_{q(a_\varepsilon)} / \lambda_{q(a_\varepsilon)}$$

If **responseType** is Bernoulli then

$$\mu_{q(1/\sigma_\varepsilon^2)} \leftarrow 1 \quad ; \quad \boldsymbol{\mu}_{q(c)} \leftarrow \boldsymbol{\omega}_{20} + (2\mathbf{y} - \mathbf{1}_n) \odot \zeta'((2\mathbf{y} - \mathbf{1}_n) \odot \boldsymbol{\omega}_{20})$$

$$\mathbf{y} \mathbf{T} \mathbf{1}_{\text{adj}} \leftarrow \boldsymbol{\mu}_{q(c)}^T \mathbf{1}_n \quad ; \quad \mathbf{X} \mathbf{T} \mathbf{y}_{\text{adj}} \leftarrow \mathbf{X}^T \boldsymbol{\mu}_{q(c)} \quad ; \quad \mathbf{Z} \mathbf{T} \mathbf{y}_{\text{adj}} \leftarrow \mathbf{Z}^T \boldsymbol{\mu}_{q(c)}$$

until the relative change in the $\log p(\mathbf{y}; \mathbf{q})$ is below $\varepsilon_{\text{toler}}$.

Outputs: All q-density parameters.

3.5.2 Deciding Between an Effect Being Zero, Linear or Non-Linear

Now let β be a generic linear coefficient and \mathbf{u} be a generic $K \times 1$ spline coefficient vector attached to one of the $\dot{\mathbf{x}}_j$ predictors. Since $\mathbf{u} = \gamma_u \tilde{\mathbf{u}}$, where the entries of γ_u are binary and the entries of $\tilde{\mathbf{u}}$ are continuous,

$$P(\mathbf{u} = \mathbf{0} | \mathbf{y}) = P(\gamma_u = 0 | \mathbf{y}) = 1 - E(\gamma_u | \mathbf{y}).$$

Therefore, after fixing τ , our strategy for deciding between an effect being zero, linear or non-linear is:

- the effect is zero if $\max\{E(\gamma_\beta|\mathbf{y}), E(\gamma_u|\mathbf{y})\} \leq 1 - \tau$,
- the effect is linear if $E(\gamma_\beta|\mathbf{y}) > 1 - \tau$ and $E(\gamma_u|\mathbf{y}) \leq 1 - \tau$,
- otherwise the effect is non-linear.

It is apparent from these rules that the parameter $\tau \in (0, 1)$ controls the degree of sparsity in the selected model, with lower values of τ producing sparse fits. Hence, we refer to τ as the *sparsity threshold parameter*. In practice, various values of τ can be contemplated but for a completely automatic model selection a good default choice is desirable. We confront this problem in the next subsection.

3.5.3 Choice of Default Values for the Sparsity Threshold Parameter

Among the family of rules indexed by the sparsity threshold parameter $\tau \in (0, 1)$, an important practical question is that of recommending a default value for τ . To aid such a recommendation, we simulated data sets from both Gaussian and Bernoulli response generalized additive models with $d_\bullet = 30$ continuous predictors. Ten of the predictors had a zero effect, 10 had linear effects with random generated coefficients, and 10 had non-linear effects. Each of the predictors were generated from independent Normal distributions. The non-linear effects corresponded to quintic polynomials with randomly generated coefficients. Each replication involved the generation of new coefficients. The sample sizes varied over $n \in \{500, 1000, 2000\}$ and, for the Gaussian response case, the error standard deviations varied over $\sigma_\epsilon \in \{0.25, 0.5, 1, 2\}$. For each combination of sample size and error standard deviation 100 data sets were generated. Fitting was carried out using both Algorithm 2 with $N_{\text{warm}} = N_{\text{kept}} = 1000$ and Algorithm 3 with $\epsilon_{\text{toler.}} = 10^{-8}$. Model selection was applied according to the rules of Section 3.5 with $\tau \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$. The performance measure was misclassification rate for the 30 candidate predictors being classified into one of three classes: zero effect, linear effect and non-linear effect.

Figure 3 displays the misclassification rate data for Algorithms 2 from 100 simulation replications. Each panel corresponds to a different combination of sample size and error standard deviation. Within each panel, side-by-side boxplots of the misclassification rate are shown as a function of τ . For low noise levels there is not much of a difference, but for $\sigma_\epsilon \geq 1$ it is advantageous to have τ equal to the natural choice of 0.5. Note, however, that this recommendation is necessarily limited due to being based on a single simulation study.

The analogous results for the mean field variational Bayes approach of Algorithm 3 are shown in Figure 4. This time the boxplots indicate better performance for $\tau < 0.5$. We conjecture that mean field approximations have a detrimental effect on the $\tau = \frac{1}{2}$ decision rules and, for reasons yet to be understood, are somewhat remedied by setting τ to be a lower value such as 0.1. Additional checks, not shown here, indicate the the classification performance gets worse for τ smaller than 0.1 for this simulation set-up. Acknowledging the limitations of a single simulation study, our recommended default for τ in the mean field variational Bayes case is $\tau = 0.1$.

We also ran simulation studies for the Bernoulli response case, with a similar design to the Gaussian study. The recommendations of $\tau = 0.5$ for Markov chain Monte Carlo and $\tau = 0.1$ for mean field variational Bayes were also supported by that study.

Additional simulation studies, involving an alternative evaluation metric and hyperparameter sensitivity checks, are given in Section S.3. These studies do not alter any of our recommendations concerning the choice of τ .

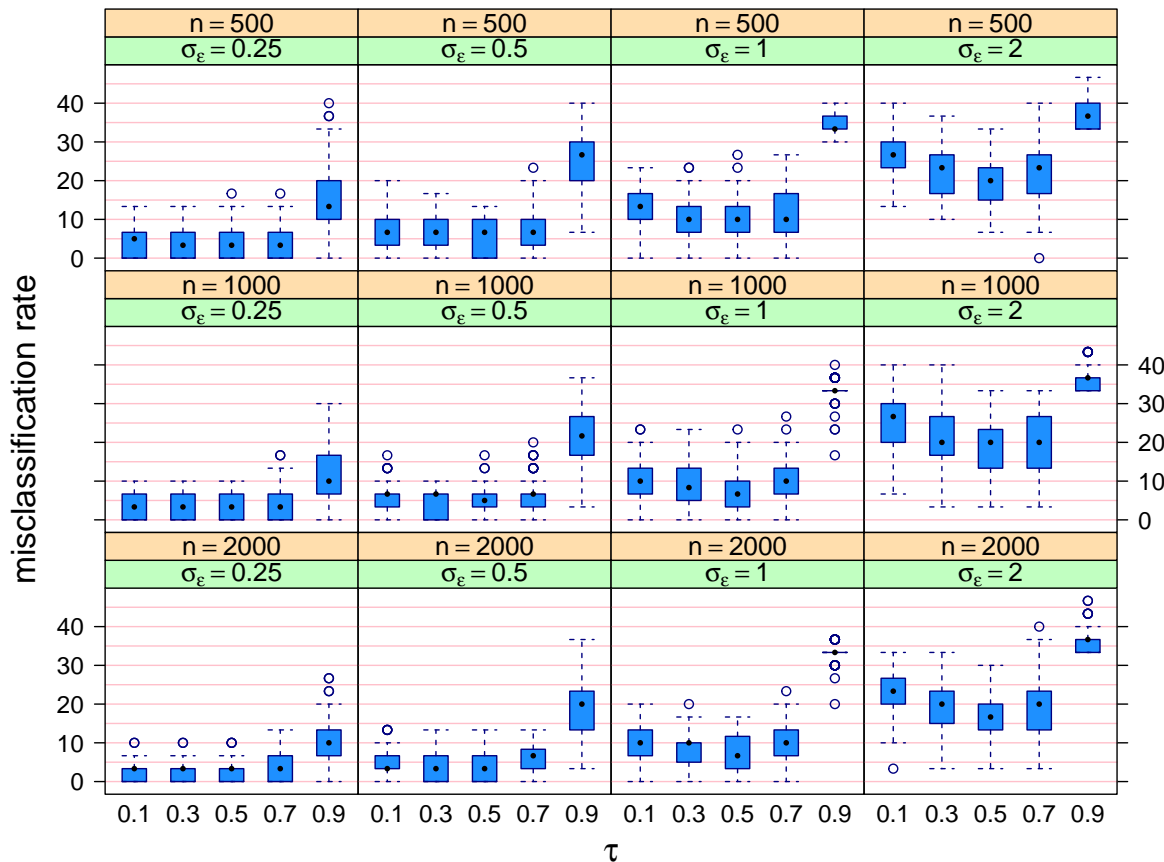


Figure 3: Side-by-side boxplots of the misclassification rates for the Markov chain Monte Carlo Algorithm 2 for the simulation study described in the text. Each panel corresponds to a different combination of sample size and error standard deviation. Within each panel, the side-by-side boxplots compare misclassification rate as a function of the threshold parameter τ .

3.6 Package in the R Language

The R package `gamselBayes` (He & Wand, 2023) implements Algorithms 2 and 3 and provides tabular and graphical summaries of selected generalized additive models. Speed is enhanced via C++ implementation of the loops in the two algorithms. The `gamselBayes` package is available on the Comprehensive R Archive Network (<https://www.R-project.org>). The `gamselBayes` package is accompanied by a vignette which provides fuller details on its use. The vignette PDF file is opened via the command `gamselBayesVignette()`.

4 Comparative Performance

We ran a second simulation study to assess comparative performance of the new methodology with respect to some of the other existing approaches to three-category generalized additive model selection. The simulation design was the same as that described in Section 3.5.3. In keeping with the findings of that section, in Algorithm 2 the threshold parameter was set to $\tau = 0.5$ and for Algorithm 3 it was set to $\tau = 0.1$.

The other approaches considered were those used by the R packages:

1. `spikeSlabGAM` (Scheipl, 2022), which is a Bayesian approach that is described in Scheipl *et al.* (2012). Details on use of the `spikeSlabGAM` package are given in Scheipl (2011).

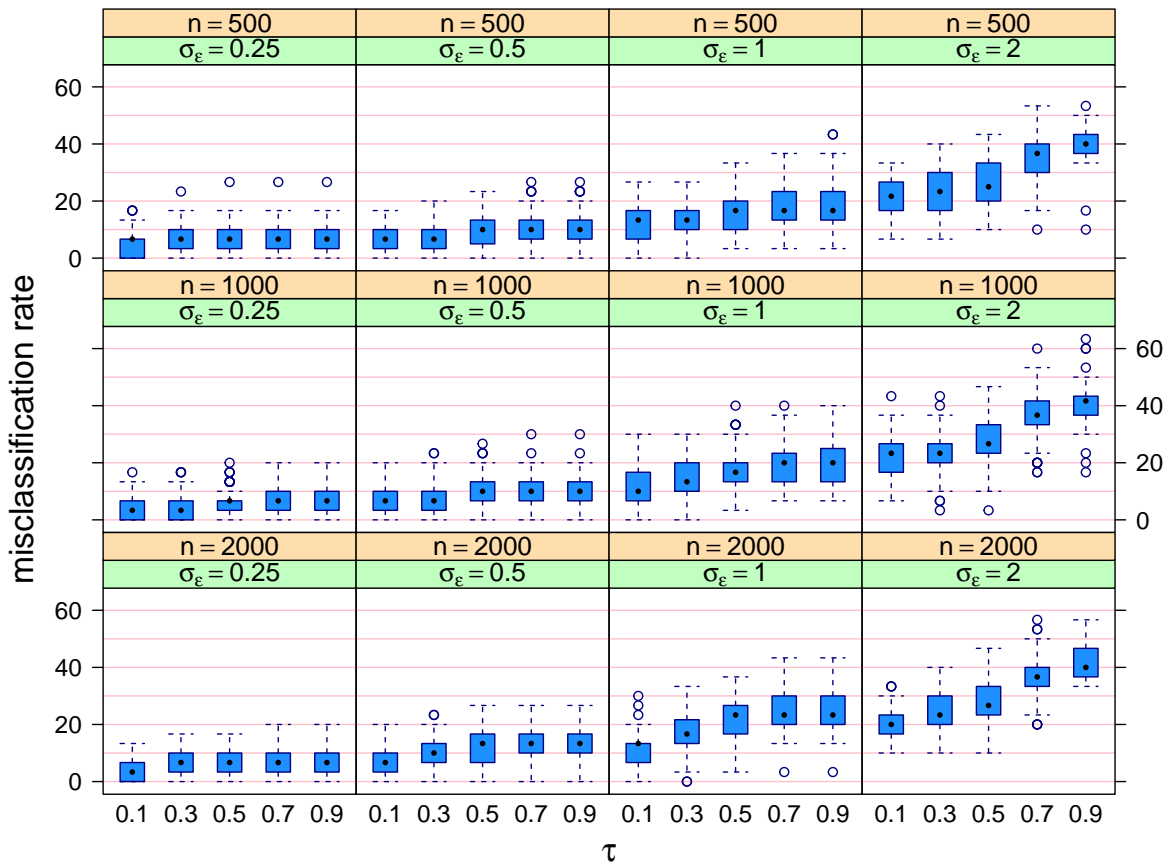


Figure 4: Side-by-side boxplots of the misclassification rates for the mean field variational Bayes Algorithm 3 for the simulation study described in the text. Each panel corresponds to a different combination of sample size and error standard deviation. Within each panel, the side-by-side boxplots compare misclassification rate as a function of the threshold parameter τ .

2. `gamsel` (Chouldechova & Hastie, 2022), which implements the frequentist approach described in Chouldechova & Hastie (2015). The package’s main function, `cv.gamsel()`, computes a family of generalized additive model fits over a grid of regularization parameter values. For selection of a single model, `cv.gamsel()` provides the option of minimizing a k -fold cross-validation function over the grid.

The Bernoulli response versions of these approaches involve the logit link function, rather than the probit link function used by Algorithms 2 and 3. This necessitated use of the appropriate inverse link transformation for the generation of binary response data in this simulation study.

In the case of `spikeSlabGAM`, we used the default call to its `spikeSlabGAM()` function. The model having highest posterior probability in the `spikeSlabGAM()` output object was selected. The essential difference between `spikeSlabGAM` and Algorithms 2 and 3 is the form of the prior distributions imposed on the coefficients for the linear and spline components. In the notation of Section 2.4, `spikeSlabGAM` replaces (6) by

$$p(\beta|\sigma_\beta, \rho_\beta) = \frac{\rho_\beta \exp\{-\beta^2/(2\sigma_\beta^2)\}}{(2\pi\sigma_\beta^2)^{1/2}} + \frac{(1-\rho_\beta) \exp[-\beta^2/\{2(v_0\sigma_\beta)^2\}]}{\{2\pi(v_0\sigma_\beta)^2\}^{1/2}} \quad \text{where } v_0 \ll 1 \quad (17)$$

with v_0 having a default value of 0.00025. Note (17) is an alternative to the “spike-and-slab” prior used by (6), with the “slab” being Gaussian rather than Laplacian and the default “spike” being a $N(0, 0.00025^2)$ mass rather than the point mass at zero. For spline coefficient vectors,

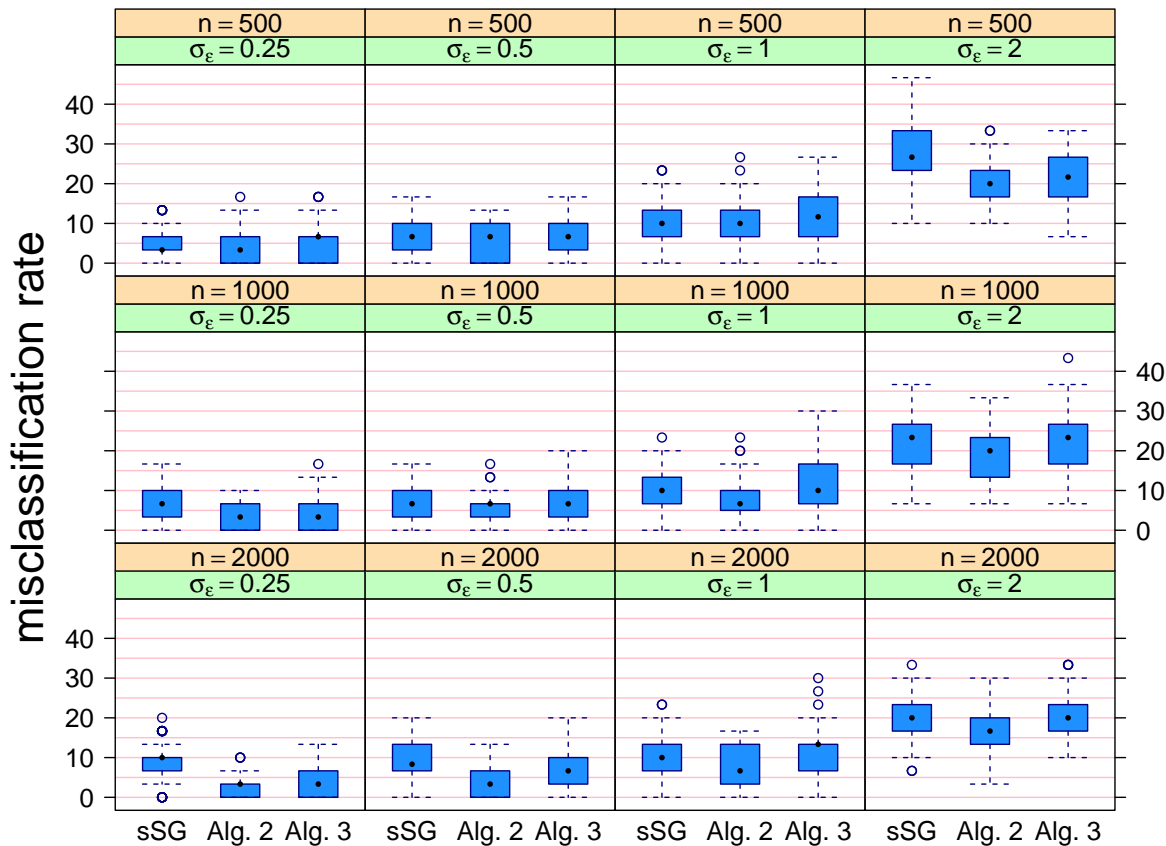


Figure 5: Side-by-side boxplots of the misclassification rates for the comparative performance simulation study described in the text in the case of the response variable being Gaussian. Each panel corresponds to a different combination of sample size and error standard deviation. Within each panel, the side-by-side boxplots compare misclassification rate across each of three methods: **spikeSlabGAM** with default settings (sSG), Algorithm 2 (Alg. 2) and Algorithm 3 (Alg. 3).

the alternative to (7) used by **spikeSlabGAM** is an extension of (17) that is described by Figure 1 of Scheipl (2011) and accompanying text. For Gaussian response models **spikeSlabGAM** uses Gibbs sampling, but requires Metropolis-Hastings sampling for non-Gaussian responses.

Preliminary checks revealed that default regularization grid used by `cv.gamsel()` did not lead to very good three-category classification performance, with the cross-validation mean function often being monotonic rather than U-shaped. To circumvent this apparent default grid problem, with respect to the three-category misclassification rate, we experimented with its choice and found that geometric sequence of size 50 between 0.01 and 2 usually lead to U-shaped cross-validation mean functions for the simulation settings. This regularization grid was used throughout the comparative performance simulation study with 10-fold cross-validation for model selection. Two cross-validation-based choices were considered: the regularization parameter matching the absolute minimum of the mean values, and largest regularization parameter value such that mean minus one standard deviation is below the absolute minimum. However, after running the simulation study it was found that the three-category misclassification rates for the `gamsel` approaches were considerably higher than the other approaches since it has a tendency to choose larger models. Given this poor performance for misclassification rate, relative to the other methods in the study, the `gamsel` results are excluded from the upcoming graphical summaries (Figures 5 and 6).

Figure 5 shows the misclassification rates for Algorithms 2 and 3 in comparison with the

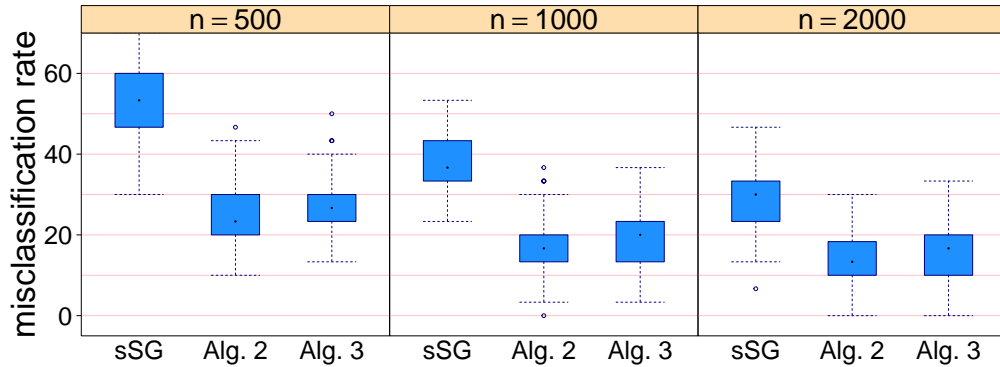


Figure 6: Side-by-side boxplots of the misclassification rates for the comparative performance simulation study described in the text in the case of the response variable being binary. Each panel corresponds to a different combination of sample size and error standard deviation. Within each panel, the side-by-side boxplots compare misclassification rate across each of three methods: spikeSlabGAM with default settings (sSG), Algorithm 2 (Alg. 2) and Algorithm 3 (Alg. 3).

	gamsel	spikeSlabGAM	Algorithm 2	Algorithm 3
10th percentile	3.69	84.9	1.78	0.326
50th percentile	8.12	167.0	2.14	0.466
90th percentile	17.50	339.0	3.03	0.768

Table 2: 10th, 50th and 90th percentiles for the number of seconds required for each generalized additive model selection approach across all settings and replications for the comparative performance simulation study.

default version of the spikeSlabGAM approach as side-by-side boxplots for the Gaussian response case. In the lower error standard deviation situations, all have similar performance. The fast variational approach of Algorithms 3 is seen to have lower accuracy when the noise level is higher. This degradation in performance needs to be mitigated against run time, which is addressed later in this section.

The binary response simulation results are shown in Figure 6. Algorithms 2 and 3 are seen to have better three-category classification performance compared with spikeSlabGAM for the binary response simulation study.

Lastly, we report on the computing times for the four approaches. Specifically, these are elapsed times in seconds for each generalized additive model selection on a MacBook Air laptop computer with 16 gigabytes of memory and a 3.2 gigahertz processor. Algorithms 2 and 3 were implemented using the Rcpp interface (Eddelbuettel & François, 2011) to the C++ language. The Markov chain Monte Carlo sample size values corresponded to the spikeSlabGAM and gamselBayes defaults of 1,500 and 2,000 respectively. Table 2 lists the 10th, 50th and 90th percentile number of seconds for each approach across all settings and replications.

It is apparent from Table 2 that, despite exhibiting very good classification, spikeSlabGAM is comparatively slow and does not scale well to large problems. Algorithm 2 took less than around 3 seconds for 90% of the fits in the simulation study. The faster variational approach of Algorithm 3 only required less than a second of computing time for most of the fits. Therefore, the new approaches have very good scalability for the generalized additive model selection problem.

The impact of sample size and number of candidate predictors on computing times for Algorithms 2 and 3 is investigated in Section S.3.2.

candidate predictor	est. type	candidate predictor	est. type
bad public credit record?	linear	credit score of 3?	zero
denied mortgage insurance?	linear	credit score of 4?	zero
applicant self-employed?	linear	credit score of 5?	zero
applicant single?	linear	mortgage credit score of 1?	zero
applicant black?	linear	mortgage credit score of 2?	zero
property a condominium?	zero	mortgage credit score of 3?	zero
unemploy. rate applic. indus.	zero	debt payments/income ratio	non-linear
credit score of 1?	linear	housing expenses/income ratio	zero
credit score of 2?	linear	loan size/property value ratio	non-linear

Table 3: Each of the candidate predictors for the Boston mortgage example and the estimated effect type from application of Algorithm 2 and effect type estimation rules of Section 3.5. The candidate predictors with question marks correspond to binary indicator variables. The abbreviation “unemploy. rate applic. indus.” stands for the unemployment rate of the industry corresponding to the applicant’s occupation.

5 Data Illustrations

We finish off with two illustrations for actual data. Both illustrations involve binary responses. The first one is a relative small problem, where Markov chain Monte Carlo fitting of the binary response adjustment of (9) is quick. The second example involves a much bigger data set, and mean field variational Bayes offers relatively fast model selection.

5.1 Application to Mortgage Applications Data

Data originating from the Federal Bank of Boston, U.S.A., has 2,380 records on mortgage applications, and is available in the R data package `Ecdat` (Croissant, 2022) as a data frame titled `Hmda`. The response variable is the indicator of whether the mortgage application was denied. After conversion of each of the categorical variables to indicator form there are 20 candidate predictors. Fourteen of these candidate predictors are binary, so can only be considered as having a zero or linear effect. The remaining four predictors are continuous, and three of these were considered as having zero, linear or non-linear effects. One of them, corresponding to the unemployment rate of the industry corresponding to the applicant’s occupation, has only 10 unique values and penalized spline models have borderline viability. Therefore, the effect of this predictor was restricted to zero versus linear.

Application of Algorithm 2 and the effect type estimation rules of Section 3.5 with $\tau = 0.5$ led to the estimated effect types listed in Table 3. Markov chain Monte Carlo sampling involved a warm-up of length 1,000 and 1,000 retained samples used for inference. Chain diagnostic graphics, including trace, lag-1 and autocorrelation function plots, indicated good convergence. The vignette attached to the `gamselBayes` package includes these diagnostic graphics.

As is apparent from Table 3, the selected model has 7 linear effects, 2 non-linear effects and 9 candidate predictors discarded. Table 4 provides estimation and inferential summaries for the linear effects.

Table 3 shows an applicant having bad public credit record is more likely to have their mortgage application denied, which is in keeping with financial commonsense. Of potential interest from a social justice standpoint is the significant effects on denial probability for applicants that are either black or single.

Figure 7 shows the two effects have non-linear effects in the selected model. The effect of debt payment to income ratio is quite a striking non-monotonic curve.

predictor	posterior mean	95% credible interval
indicator of bad public credit record	0.7350	(0.4926, 0.9848)
indicator of denied mortgage insurance	2.7620	(2.1426, 3.5172)
indicator of applicant being single	0.1370	(0.0000, 0.3417)
indicator of applicant being black	0.3461	(0.0842, 0.5404)
indicator of applicant being self-employed	0.1703	(0.0000, 0.4363)
indicator of credit score equalling 1	-0.6906	(-0.8980, -0.4513)
indicator of credit score equalling 2	-0.3238	(-0.5869, 0.0000)

Table 4: Approximate posterior means and approximate 95% credible intervals for the coefficients of each of the selected linear fits based on the Markov chain Monte Carlo samples generated from Algorithm 2 for the Boston mortgages example.




Figure 7: The two estimated non-linear effects for the Boston mortgage example from application of Algorithm 2 and effect type estimation rules of Section 3.5 with $\tau = 0.5$. Each curve is the slice of estimated probability of mortgage denial as a function of the predictor, with all other selected predictors set to their median values. The shaded region corresponds to pointwise approximate 95% credible intervals.

5.2 Application to Car Auction Data

During 2011-2012 the kaggle Internet platform (<https://www.kaggle.com>) hosted a classification competition involving training data consisting of 49 variables on 72,983 cars purchased at automobile auctions by automobile dealerships in U.S.A. The title of the competition was “Don’t Get Kicked!”. A version of the data in which all categorical variables have been converted to binary variable indicator form is stored in the data frame `carAuction` within the R package HRW (Harezlak *et al.*, 2021). The response variable is the indicator of whether the car purchased at auction by the dealership had serious problems that hinder or prevent it being sold. For short, we refer to such a car as a “bad buy”. Forty-four of the candidate predictors are binary. The other 5 candidate predictors are continuous. However, the age at sale variable has only 10 unique values. For the same reasons given for the unemployment rate variable considered in the Boston mortgages example, we exclude age at sale from having a non-linear effect.

Since this generalized additive model selection problem involves a relatively large sample

predictor	posterior mean	95% credible interval
indic. made in U.S.A.	-0.04605	(-0.06925, -0.02323)
age at sale (years)	0.09344	(0.08771, 0.09896)
indic. color is red	0.05130	(0.02498, 0.07713)
indic. make is Chevrolet	-0.1103	(-0.1358, -0.08426)
indic. make is Chrysler	0.09517	(0.06910, 0.1217)
indic. make is Dodge	-0.03093	(-0.05385, -0.006961)
indic. purchased online	-0.06229	(-0.1117, 0.0000)
acquisition price (U.S. dollars)	-6.014×10^{-6}	$(-9.352, -2.603) \times 10^{-6}$
indic. purchased in 2010	0.1085	(0.09290, 0.1242)
indic. purch. in Florida	-0.1173	(-0.1391, -0.09529)
indic. purch. in North Carolina	-0.1074	(-0.1326, -0.08190)
indic. purch. in Texas	0.09706	(0.07799, 0.1162)
indic. medium-sized vehicle	-0.07368	(-0.09095, -0.05622)
indic. sports utility vehicle	0.1899	(0.1646, 0.2153)
indic. manual transmission	-0.1574	(-0.1970, -0.1168)
indic. trim level is 'Bas'	0.05521	(0.0355, 0.07477)
indic. trim level is 'LS'	-0.06153	(-0.09180, -0.03176)
indic. has alloy wheels	-1.513	(-1.546, -1.480)
indic. has wheel covers	-1.585	(-1.619, -1.551)

Table 5: *Approximate posterior means and approximate 95% credible intervals for the coefficients of each of the selected linear fits based on the mean field variational Bayes optimal q -densities obtained from Algorithm 3 for the car auction example.*

size and number of candidate predictors, we use it to illustrate the fast variational approach corresponding to Algorithm 3. The mean field variational Bayes iterations described there were iterated until the relative change in the approximate marginal log-likelihood fell below 10^{-8} . On the second author's MacBook Air laptop, with a 3.2 gigahertz processor and 16 gigabytes of random access memory, mean field Bayes variational fitting took 11 seconds. The rules of Section 3.5 were applied with $\tau = 0.1$. This resulted in 19 predictors being selected as having a linear effect and 3 predictors having non-linear effects. Twenty-seven of the 49, or 55%, of candidate predictors were discarded.

Table 5 provides estimation and inferential summaries for the linear effects coefficients. Most of the predictor effects are intuitive, such as older cars being more likely to be a bad buy and presence of wheel covers lowering the bad buy probability. Some of them, such as the effect of cars being purchased in particular states, are more intriguing.

Figure 8 shows the three selected non-linear effects, which are the impacts of the probability of a bad buy as functions of the odometer reading in miles, acquisition cost paid for the car at the time of purchase in U.S. dollars and warranty cost in U.S. dollars. The middle panel of Figure 8 shows that a cost of about 10,000 U.S. dollars is best, and that the probability of bad buy increases when the cost deviates away from this amount. The shaded regions of Figure 8 corresponds to pointwise approximate 95% credible intervals. However, for a binary response model such as this, there is considerable mean field approximation error which tends to make the credible intervals overly narrow.

As a type of check, we also applied the Markov chain Monte Carlo Algorithm 2 to the same data set. This resulted in 14 of the 19 predictors in Table 5 being selected. Three predictors not selected by Algorithm 3, such as indicators of the auction provider, were selected by Algorithm 2. The odometer reading predictor was estimated to have a non-linear effect by Algorithm 3, but to have a linear effect by Algorithm 2. In summary, Algorithm 3 selected 22 predictors whilst Algorithm 2 selected 20 predictors with 17 predictors in common from the




Figure 8: The estimated non-linear effects for the car auction example from application of Algorithm 3 and effect type estimation rules of Section 3.5 with $\tau = 0.1$. The curves are slices of estimated probability of bad buy as a function of the predictor, with all other selected predictors set to their median values. The shaded regions correspond to pointwise approximate 95% credible intervals, but are subject to considerable mean field approximation error. The rug at the base of each plot is based on a random sample of 2,500 cars.

two approaches. This suggests reasonable accuracy of the faster variational approach for this example.

6 Concluding Remarks

The methodology conveyed by Algorithms 1–3 and the effect type classification rules given in Sections 3.5.1 and 3.5.2 represent a practical Bayesian alternative to the frequentist methodology of Chouldechova & Hastie (2015) for three-category generalized additive model selection. Both approaches are driven by the goals of speed and scalability to large data sets. The new Bayesian approach is the clear winner in terms of accuracy according to our simulation studies.

Our, admittedly limited, simulation studies indicate improved classification performance compared with default use of `spikeSlabGAM` in binary response situations. In the Gaussian response situations the performance of Algorithm 2 and `spikeSlabGAM` is similar, with Algorithm 3 falling behind for higher noise situations. This needs to be mitigated against the vastly improved speed and scalability, as indicated by Table 2, of this article’s new methodologies. The best approach in practice depends on data set size and time demands, with the simulation results of Sections 3 and 4 providing some guidance. Additional simulation results are given in Section S.3 of the supplement.

Acknowledgement

We are grateful to two reviewers for their comments that led to improvements. This research was supported by Australian Research Council grant DP180100597.

Disclosure Statement

The authors report that there are no competing interests to declare.

References

- Albert, J.H. & Chib, S. (1993). Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association*, **88**, 669–679.
- Azzalini, A. (2023). `sn 2.1.1`: The Skew-Normal and related distributions such as the Skew-t and the Unified Skew-Normal. R package.
<http://azzalini.stat.unipd.it/SN>
- Bhadra, A., Datta, J., Polson, N.G. & Willard, B. (2019). Lasso meets horseshoe: a survey. *Statistical Science*, **34**, 405–427.
- Bürkner, P.-C. (2022). `brms 2.18.0`: Bayesian regression models using Stan. R package.
<https://r-project.org>
- Carvalho, C.M., Polson, N.G. & Scott, J.G. (2010). The horseshoe estimator for sparse signals. *Biometrika*, **97**, 465–480.
- Chouldechova, A. & Hastie, T. (2015). Generalized additive model selection.
<https://arXiv.org/abs/1506.03850v2>
- Chouldechova, A. & Hastie, T. (2022). `gamsel 1.8`: Fit regularization path for generalized additive models. R package. <https://r-project.org>
- Croissant, Y. (2022). `Ecdat 0.4`: Data sets for econometrics. R package.
<https://r-project.org>
- Eddelbuettel, D. & François, R. (2011). `Rcpp`: Seamless R and C++ integration. *Journal of Statistical Software*, **40(8)**, 1–18.
- Gelfand, A.E. & Smith, A.F.M. (1990). Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association*, **85**, 398–409.
- Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models. *Bayesian Analysis*, **1**, 515–533.
- George & McCulloch, R.E. (1993). Variable selection via Gibbs sampling. *Journal of the American Statistical Association*, **88**, 881–889.
- Griffin, J.E. & Brown, P.J. (2011). Bayesian hyper-lassos with non-convex penalization. *Australian and New Zealand Journal of Statistics*, **53**, 423–442.
- Harezlak, J., Ruppert, D. & Wand, M.P. (2021). `HRW 1.0`: Datasets, functions and scripts for semiparametric regression supporting Harezlak, Ruppert & Wand (2018). R package.
<https://r-project.org>
- Hastie, T.J. & Tibshirani, R.J. (1990). *Generalized Additive Models*. New York: Chapman & Hall.
- He, V.X. & Wand, M.P. (2023). `gamselBayes`: Bayesian generalized additive model selection. R package version 2.0.
<http://cran.r-project.org>.
- Ishwaran, H. & Rao, J.S. (2005). Spike and slab variable selection: frequentist and Bayesian strategies. *The Annals of Statistics*, **33**, 730–733.

- Kyung, M., Gill, J., Ghosh, M. & Casella, G. (2010). Penalized regression, standard errors, and Bayesian lassos. *Bayesian Analysis*, **5**, 369–412.
- Lempers, F.B. (1971). *Posterior Probabilities of Alternative Linear Models*. Rotterdam: Rotterdam University Press.
- Michael, J.R., Schucany, W.R. & Haas, R.W. (1976). Generating random variates using transformations with multiple roots. *The American Statistician*, **30**, 88–90.
- Mitchell, T.J. & Beauchamp, J.J. (1988). Bayesian variable selection in linear regression. *Journal of the American Statistical Association*, **83**, 1023–1032.
- Ngo, L. and Wand, M.P. (2004). Smoothing with mixed model software. *Journal of Statistical Software*, **9**, Article 1, 1–54.
- Ormerod, J.T. and Wand, M.P. (2010). Explaining variational approximations. *The American Statistician*, **64**, 140–153.
- Park, T. & Casella, G. (2008). The Bayesian Lasso. *Journal of the American Statistical Association*, **103**, 681–686.
- Ravikumar, P., Lafferty, J., Liu, H. & Wasserman, L. (2009). Sparse additive models. *Journal of the Royal Statistical Society, Series B*, **71**, 1009–1030.
- R Core Team (2023). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. <https://www.r-project.org/>.
- Reich, B.J., Sorlie, C.B. & Bondell, H.D. (2009). Variable selection in smoothing spline ANOVA: application to deterministic computer codes. *Technometrics*, **51**, 110–120.
- Robert, C.P. (1995). Simulation of truncated normal variates. *Statistics and Computing*, **5**, 121–125.
- Scheipl, F. (2011). spikeSlabGAM: Bayesian variable selection, model choice and regularization for generalized additive mixed models in R. *Journal of Statistical Software*, **43**, Issue 14, 1–24.
- Scheipl, F. (2022). spikeSlabGAM 1.1: Bayesian variable selection and model choice for generalized additive mixed models. R package. <https://github.com/fabian-s/spikeSlabGAM>
- Scheipl, F., Fahrmeir, L. & Kneib, T. (2012). Spike-and-slab priors for function selection in structured additive regression models. *Journal of the American Statistical Association*, **107**, 1518–1532.
- Shively, T.S., Kohn, R. & Wood, S. (1999). Variable selection and function estimation in additive nonparametric regression using a data-based prior. *Journal of the American Statistical Association*, **94**, 777–794.
- Umlauf, N., Klein, N., Zeileis, A. & Simon, T. (2023). bamlss 1.2: Bayesian additive models for location, scale, and shape (and beyond). R package. <https://www.bamlss.org>
- Umlauf, N., Kneib, T. & Klein, N. (2023). BayesX 0.3: R utilities accompanying the software package BayesX. R package. <https://www.BayesX.org>

- Wainwright, M.J. & Jordan, M.I. (2008). Graphical models, exponential families and variational inference. *Foundations and Trends in Machine Learning*, **1**, 1–305.
- Wand, M.P. & Ormerod, J.T. (2008). On semiparametric regression with O’Sullivan penalized splines. *Australian and New Zealand Journal of Statistics*, **50**, 179–198.
- Wand, M.P. and Ormerod, J.T. (2011). Penalized wavelets: embedding wavelets into semiparametric regression. *Electronic Journal of Statistics*, **5**, 1654–1717.
- Wand, M.P. and Ormerod, J.T. (2012). Continued fraction enhancement of Bayesian computing. *Stat*, **1**, 31–41.
- Wood, S.N. (2017). *Generalized Additive Models: An Introduction with R, Second Edition*, Boca Raton, Florida: CRC Press.
- Yuan, M. & Lin, Y. (2006). Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society, Series B*, **68**, 49–67.

Supplement for:
**Bayesian Generalized Additive Model Selection
Including a Fast Variational Option**

BY VIRGINIA X. HE AND MATT P. WAND

University of Technology Sydney

S.1 The Canonical Demmler-Reinsch Spline Basis

Let $\mathbf{x} = (x_1, \dots, x_n)$ be a continuous univariate data set. In the context of this article, the x_i s correspond to values of a continuous candidate predictor. Let $[a, b]$ be an interval containing the x_i s. For an integer $K \leq n - 2$, let $\boldsymbol{\kappa}_{\text{inter}} \equiv (\kappa_1, \dots, \kappa_{K-2})$ be a set of so-called interior knots such that

$$a < \kappa_1 < \dots < \kappa_{K-2} < b.$$

A reasonable default value for K is around 30, or smaller values if the number of unique x_i s is lower. It is common to place the interior knots at sample quantiles of the x_i s.

We now list steps for construction of the matrix \mathbf{Z} containing canonical Demmler-Reinsch basis functions of the entries of \mathbf{x} . The justification for Steps (3)–(6) is given in Section 9.1.1 of Ngo & Wand (2004).

- (1) Use the steps described in Section 4 of Wand & Ormerod (2008) to obtain the matrix denoted by \mathbf{Z} in that section's equation (6), which contains canonical O'Sullivan spline basis functions. Denote this matrix by \mathbf{Z}_{OS} and note that it has dimension $n \times K$.
- (2) Form the matrix $\mathbf{C}_{\text{OS}} = [\mathbf{1}_n \ \mathbf{x} \ \mathbf{Z}_{\text{OS}}]$ and set $\mathbf{D} = \text{diag}(0, 0, \mathbf{1}_K)$.
- (3) Obtain the singular value decomposition of \mathbf{C}_{OS} :

$$\mathbf{C}_{\text{OS}} = \mathbf{U}_C \text{diag}(\mathbf{d}_C) \mathbf{V}_C^T \text{ where } \mathbf{U}_C \text{ is } n \times (K + 2) \text{ and } \mathbf{V}_C \text{ is } (K + 2) \times (K + 2)$$

$$\text{such that } \mathbf{U}_C^T \mathbf{U}_C = \mathbf{V}_C^T \mathbf{V}_C = \mathbf{I}_{K+2}.$$

- (4) Form the symmetric matrix $\text{diag}(\mathbf{1}/\mathbf{d}_C) \mathbf{V}_C^T \mathbf{D} \mathbf{V}_C \text{diag}(\mathbf{1}/\mathbf{d}_C)$ and obtain its singular value decomposition:

$$\text{diag}(\mathbf{1}/\mathbf{d}_C) \mathbf{V}_C^T \mathbf{D} \mathbf{V}_C \text{diag}(\mathbf{1}/\mathbf{d}_C) = \mathbf{U}_D \text{diag}(\mathbf{d}_D) \mathbf{V}_D^T \text{ where } \mathbf{U}_D \text{ is } (K + 2) \times (K + 2)$$

$$\text{and } \mathbf{V}_D \text{ is } (K + 2) \times (K + 2) \text{ such that } \mathbf{U}_D^T \mathbf{U}_D = \mathbf{V}_D^T \mathbf{V}_D = \mathbf{I}_{K+2}.$$

- (5) Set the full (non-canonical) Demmler-Reinsch matrix as follows: $\mathbf{C}_{\text{DR}} \leftarrow \mathbf{U}_C \mathbf{U}_D$.
- (6) The next steps assume that the singular value decompositions follow the convention that \mathbf{d}_D is a $(K + 2) \times 1$ vector with its entries in non-increasing order. Adjustments to the singular value decompositions are needed if this convention is not used.
- (7) Set the $(K + 2) \times 1$ vector \mathbf{s}_D as follows:

$$\omega_{21} \leftarrow \sqrt{K \text{th entry of } \mathbf{d}_D}, \quad ; \quad \mathbf{s}_D \leftarrow \omega_{21} \mathbf{1}_{K+2} / \sqrt{\mathbf{d}_D}$$

and then set the last two entries of \mathbf{s}_D to equal 1.

- (8) Set the full canonical Demmler-Reinsch design matrix as follows:

$$\mathbf{C}_{\text{cDR}} \leftarrow \mathbf{C}_{\text{DR}} \text{diag}(\mathbf{s}_D).$$

(9) Set the O’Sullivan to canonical Demmler-Reinsch transformation matrix as follows:

$$\mathbf{L}_{\text{OS.to.cDR}} \leftarrow \mathbf{V}_C \text{diag}(\mathbf{1}/d_C) \mathbf{U}_D \text{diag}(\mathbf{s}_D).$$

This $(K + 2) \times (K + 2)$ matrix has the following property:

$$\mathbf{C}_{\text{OS}} \mathbf{L}_{\text{OS.to.cDR}} = \mathbf{C}_{\text{cDR}}$$

and is useful for prediction and plotting purposes. This is because grid-wise analogues of \mathbf{C}_{OS} are readily computed using the structures described in Wand & Ormerod (2008) involving cubic B-spline basis functions.

(10) Reverse the order of the columns of \mathbf{C}_{cDR} . Reverse the order of the columns of $\mathbf{L}_{\text{OS.to.cDR}}$.

(11) The matrix containing canonical spline basis functions of the inputs \mathbf{x} and $\boldsymbol{\kappa}_{\text{inter}}$ is

$$\mathbf{Z} \leftarrow \text{the } n \times K \text{ matrix consisting of columns 3 to } K + 2 \text{ of } \mathbf{C}_{\text{cDR}}.$$

A function in the R language for computing \mathbf{Z} and $\mathbf{L}_{\text{OS.to.cDR}}$ for given \mathbf{x} and $\boldsymbol{\kappa}_{\text{inter}}$ can be accessed by downloading the accompanying `gamselBayes` package. Assuming that the `gamselBayes` package is installed, the relevant function is `gamselBayes:::ZcDR()`.

S.2 Approximate Marginal Log-Likelihood Expressions

The approximate marginal log-likelihood is

$$\log \underline{p}(\mathbf{y}; \mathbf{q}) = \begin{cases} \log \underline{p}(\mathbf{y}; \mathbf{q}, \mathbf{C}) + E_q[\log\{\mathbf{p}(\mathbf{y}|\beta_0, \gamma_\beta, \tilde{\boldsymbol{\beta}}, \gamma_u, \tilde{\mathbf{u}}, \sigma_\varepsilon^2)\}] \\ \quad + E_q[\log\{\mathbf{p}(\sigma_\varepsilon^2|a_\varepsilon)\}] - E_q[\log\{\mathbf{q}(\sigma_\varepsilon^2)\}] \\ \quad + E_q[\log\{\mathbf{p}(a_\varepsilon)\}] - E_q[\log\{\mathbf{q}(a_\varepsilon)\}] & \text{Gaussian response case,} \\ \log \underline{p}(\mathbf{y}; \mathbf{q}, \mathbf{C}) + E_q[\log\{\mathbf{p}(\mathbf{y}|\mathbf{c})\}] \\ \quad + E_q[\log\{\mathbf{p}(\mathbf{c}|\beta_0, \gamma_\beta, \tilde{\boldsymbol{\beta}}, \gamma_u, \tilde{\mathbf{u}})\}] - E_q[\log\{\mathbf{q}(\mathbf{c})\}] & \text{Bernoulli response case,} \end{cases}$$

where

$$\begin{aligned} \log \underline{p}(\mathbf{y}; \mathbf{q}, \mathbf{C}) &= E_q[\log\{\mathbf{p}(\beta_0)\}] - E_q[\log\{\mathbf{q}(\beta_0)\}] + E_q[\log\{\mathbf{p}(\gamma_\beta)\}] - E_q[\log\{\mathbf{q}(\gamma_\beta)\}] \\ &\quad + E_q[\log\{\mathbf{p}(\tilde{\boldsymbol{\beta}}|\mathbf{b}_\beta, \sigma_\beta^2)\}] - E_q[\log\{\mathbf{q}(\tilde{\boldsymbol{\beta}})\}] + E_q[\log\{\mathbf{p}(\mathbf{b}_\beta)\}] - E_q[\log\{\mathbf{q}(\mathbf{b}_\beta)\}] \\ &\quad + E_q[\log\{\mathbf{p}(\sigma_\beta^2|a_\beta)\}] - E_q[\log\{\mathbf{q}(\sigma_\beta^2)\}] + E_q[\log\{\mathbf{p}(a_\beta)\}] - E_q[\log\{\mathbf{q}(a_\beta)\}] \\ &\quad + E_q[\log\{\mathbf{p}(\gamma_u)\}] - E_q[\log\{\mathbf{q}(\gamma_u)\}] + E_q[\log\{\mathbf{p}(\tilde{\mathbf{u}}|\mathbf{b}_u, \sigma_u^2)\}] - E_q[\log\{\mathbf{q}(\tilde{\mathbf{u}})\}] \\ &\quad + E_q[\log\{\mathbf{p}(\mathbf{b}_u)\}] - E_q[\log\{\mathbf{q}(\mathbf{b}_u)\}] + E_q[\log\{\mathbf{p}(\sigma_u^2|\mathbf{a}_u)\}] - E_q[\log\{\mathbf{q}(\sigma_u^2)\}] \\ &\quad + E_q[\log\{\mathbf{p}(\mathbf{a}_u)\}] - E_q[\log\{\mathbf{q}(\mathbf{a}_u)\}]. \end{aligned} \tag{S.1}$$

Here “C” signifies the fact that (S.1) is common to both $\log \underline{p}(\mathbf{y}; \mathbf{q})$ expressions.

Explicit expressions for $\log \underline{p}(\mathbf{y}; \mathbf{q})$ can be obtained by simplifying each of the \mathbf{q} -density moment expressions. For example, the first term of (S.1) is

$$\begin{aligned} E_q[\log\{\mathbf{p}(\beta_0)\}] &= -\frac{1}{2} \log(2\pi) - \frac{1}{2} \log(\sigma_{\beta_0}^2) - \frac{1}{2} E_q(\beta_0^2)/\sigma_{\beta_0}^2 \\ &= -\frac{1}{2} \log(2\pi) - \frac{1}{2} \log(\sigma_{\beta_0}^2) - \frac{1}{2} \{\mu_{\mathbf{q}(\beta_0)}^2 + \sigma_{\mathbf{q}(\beta_0)}^2\}/\sigma_{\beta_0}^2. \end{aligned}$$

Also, since $\mathbf{q}(\beta_0)$ is the $N(\mu_{\mathbf{q}(\beta_0)}, \sigma_{\mathbf{q}(\beta_0)}^2)$ density function, the second term of (S.1) is

$$\begin{aligned} -E_q[\log\{\mathbf{q}(\beta_0)\}] &= \frac{1}{2} \log(2\pi) + \frac{1}{2} \log(\sigma_{\mathbf{q}(\beta_0)}^2) + \frac{1}{2} E_q\{(\beta_0 - \mu_{\mathbf{q}(\beta_0)})^2\}/\sigma_{\mathbf{q}(\beta_0)}^2 \\ &= \frac{1}{2} \{\log(2\pi) + 1\} + \frac{1}{2} \log(\sigma_{\mathbf{q}(\beta_0)}^2). \end{aligned}$$

Continuing in this fashion, and accounting for some cancellations, we obtain

$$\begin{aligned}
\log \underline{p}(\mathbf{y}; \mathbf{q}, \mathbf{C}) &= \text{const}_1 - \frac{1}{2} \{ \mu_{\mathbf{q}}^2(\beta_0) + \sigma_{\mathbf{q}}^2(\beta_0) \} / \sigma_{\beta_0}^2 + \frac{1}{2} \log(\sigma_{\mathbf{q}}^2(\beta_0)) + \text{logit}(\rho_{\beta}) \sum_{j=1}^{d_{\circ} + d_{\bullet}} \mu_{\mathbf{q}}(\gamma_{\beta j}) \\
&\quad - \sum_{j=1}^{d_{\circ} + d_{\bullet}} \left[\mu_{\mathbf{q}}(\gamma_{\beta j}) \log(\mu_{\mathbf{q}}(\gamma_{\beta j})) + \{1 - \mu_{\mathbf{q}}(\gamma_{\beta j})\} \log(1 - \mu_{\mathbf{q}}(\gamma_{\beta j})) \right] \\
&\quad - \frac{1}{2} \mu_{\mathbf{q}}(1/\sigma_{\beta}^2) \sum_{j=1}^{d_{\circ} + d_{\bullet}} \mu_{\mathbf{q}}(b_{\beta j}) (\mu_{\mathbf{q}}^2(\tilde{\beta}_j) + \sigma_{\mathbf{q}}^2(\tilde{\beta}_j)) + \frac{1}{2} \log |\Sigma_{\mathbf{q}}(\tilde{\beta})| \\
&\quad - \frac{1}{2} \sum_{j=1}^{d_{\circ} + d_{\bullet}} \{1/\mu_{\mathbf{q}}(b_{\beta j})\} - \mu_{\mathbf{q}}(1/a_{\beta}) \mu_{\mathbf{q}}(1/\sigma_{\beta}^2) - \frac{1}{2} (d_{\circ} + d_{\bullet} + 1) \log(\lambda_{\mathbf{q}}(\sigma_{\beta}^2)) \\
&\quad + \mu_{\mathbf{q}}(1/\sigma_{\beta}^2) \lambda_{\mathbf{q}}(\sigma_{\beta}^2) - \mu_{\mathbf{q}}(1/a_{\beta}) / s_{\beta}^2 + \lambda_{\mathbf{q}}(a_{\beta}) \mu_{\mathbf{q}}(1/a_{\beta}) - \log(\lambda_{\mathbf{q}}(a_{\beta})) \\
&\quad - \sum_{j=1}^{d_{\bullet}} \left[\mu_{\mathbf{q}}(\gamma_{u_j}) \log(\mu_{\mathbf{q}}(\gamma_{u_j})) + \{1 - \mu_{\mathbf{q}}(\gamma_{u_j})\} \log(1 - \mu_{\mathbf{q}}(\gamma_{u_j})) \right] \\
&\quad + \text{logit}(\rho_u) \sum_{j=1}^{d_{\bullet}} \mu_{\mathbf{q}}(\gamma_{u_j}) - \frac{1}{2} \sum_{j=1}^{d_{\bullet}} \mu_{\mathbf{q}}(1/\sigma_{u_j}^2) \mu_{\mathbf{q}}(b_{u_j}) \left(\|\boldsymbol{\mu}_{\mathbf{q}}(\tilde{\mathbf{u}}_j)\|^2 + \mathbf{1}_{K_j}^T \boldsymbol{\sigma}_{\mathbf{q}}^2(\tilde{\mathbf{u}}_j) \right) \\
&\quad + \frac{1}{2} \sum_{j=1}^{d_{\bullet}} \sum_{k=1}^{K_j} \log(\sigma_{\mathbf{q}}^2(\tilde{\mathbf{u}}_{jk})) - \frac{1}{2} \sum_{j=1}^{d_{\bullet}} \{1/\mu_{\mathbf{q}}(b_{u_j})\} - \sum_{j=1}^{d_{\bullet}} \mu_{\mathbf{q}}(1/a_{u_j}) \mu_{\mathbf{q}}(1/\sigma_{u_j}^2) \\
&\quad - \frac{1}{2} \sum_{j=1}^{d_{\bullet}} (K_j + 1) \log(\lambda_{\mathbf{q}}(\sigma_{u_j}^2)) + \sum_{j=1}^{d_{\bullet}} \mu_{\mathbf{q}}(1/\sigma_{u_j}^2) \lambda_{\mathbf{q}}(\sigma_{u_j}^2) - (1/s_u^2) \sum_{j=1}^{d_{\bullet}} \mu_{\mathbf{q}}(1/a_{u_j}) \\
&\quad + \sum_{j=1}^{d_{\bullet}} \{ \lambda_{\mathbf{q}}(a_{u_j}) \mu_{\mathbf{q}}(1/a_{u_j}) - \log(\lambda_{\mathbf{q}}(a_{u_j})) \}
\end{aligned}$$

where const_1 is a constant that does not depend on any \mathbf{q} -density parameters.

In the Gaussian response case, we have

$$\begin{aligned}
\log \underline{p}(\mathbf{y}; \mathbf{q}) &= \log \underline{p}(\mathbf{y}; \mathbf{q}, \mathbf{C}) - \frac{1}{2} (n + 1) \log(\lambda_{\mathbf{q}}(\sigma_{\varepsilon}^2)) - \mu_{\mathbf{q}}(1/a_{\varepsilon}) / s_{\varepsilon}^2 - \log(\lambda_{\mathbf{q}}(a_{\varepsilon})) + \lambda_{\mathbf{q}}(a_{\varepsilon}) \mu_{\mathbf{q}}(1/a_{\varepsilon}) \\
&\quad + \text{const}_2,
\end{aligned}$$

where const_2 is a constant that does not depend on any \mathbf{q} -density parameters. In the Bernoulli response case

$$\begin{aligned}
\log \underline{p}(\mathbf{y}; \mathbf{q}) &= \log \underline{p}(\mathbf{y}; \mathbf{q}, \mathbf{C}) + \sum_{i=1}^n \log \left\{ \Phi \left((2y_i - 1) \left(\mathbf{1}_n \mu_{\mathbf{q}}(\beta_0) + \mathbf{X}(\boldsymbol{\mu}_{\mathbf{q}}(\gamma_{\beta}) \odot \boldsymbol{\mu}_{\mathbf{q}}(\tilde{\beta})) \right. \right. \right. \\
&\quad \left. \left. \left. + \sum_{j=1}^{d_{\bullet}} \mathbf{Z}_j (\mu_{\mathbf{q}}(\gamma_{u_j}) \boldsymbol{\mu}_{\mathbf{q}}(\tilde{\mathbf{u}}_j)) \right) \right) \right\}.
\end{aligned}$$

S.3 Additional Simulation Results

We have conducted thorough simulation testing of Algorithms 2 and 3 and the model selection strategies given in Section 3.5. Space considerations are such that Sections 3 and 4 contain only our primary simulation results. Additional simulation results are conveyed in this section.




Figure S.1: Side-by-side boxplots of the logarithms, to base 10, of relative test error for the Markov chain Monte Carlo Algorithm 2 for the simulation study described in the text. Each panel corresponds to a different combination of sample size and error standard deviation. Within each panel, the side-by-side boxplots compare relative test error as a function of the threshold parameter τ .

S.3.1 Alternative Evaluation Metrics

The model selection recommendations of Section 3.5 are guided by the effect type misclassification rate since we believe this particular evaluation metric to be best aligned with the practical goal of achieving interpretable and parsimonious models. However effect type misclassification rate is just one of many possible evaluation metrics that could be used in simulation assessment, comparison and the guiding of tuning parameter choice. For example, the simulation studies of Hastie, Tibshirani & Tibshirani (2020), for a different regression-type setting, consider five evaluation metrics.

To see if and how our threshold parameter recommendations change if a different evaluation metric is used, we re-ran the Gaussian response simulation studies of Section 3.5 with effect type misclassification rate replaced by *relative test error*. For the situation where $d_{\circ} = 0$ and $d_{\bullet} \in \mathbb{N}$, suppose that the selected model based on the data set \mathcal{D} corresponds to \hat{f} for some additive function $\hat{f} : \mathbb{R}^{d_{\bullet}} \rightarrow \mathbb{R}$. If the true model corresponds to $f_{\text{true}} : \mathbb{R}^{d_{\bullet}} \rightarrow \mathbb{R}$ and the predictor $\mathbf{x} \in \mathbb{R}^{d_{\bullet}}$ is a random vector with density function $p(\mathbf{x})$ then the relative test error is

$$E[\{y - \hat{f}(\mathbf{x})\}^2 | \mathcal{D}] / \sigma_{\varepsilon}^2 \quad \text{where} \quad y \sim N(f_{\text{true}}(\mathbf{x}), \sigma_{\varepsilon}^2). \quad (\text{S.2})$$

Note that the expectation in (S.2) is over the predictor distribution corresponding to $p(\mathbf{x})$. The denominator in (S.2) is the *Bayes error*, corresponding to the situation where $\hat{f} = f_{\text{true}}$. Therefore, (S.2) is the test error relative to the Bayes error and is an evaluation metric with a lower bound of 1, and equals 1 when f_{true} is estimated perfectly.

Figures S.1 and S.2 are the analogues of Figures 3 and 4 with effect type misclassification rate replaced by relative test error. Monte Carlo approximations of the (S.2) numerator quantity based on 100,000 draws from the predictor distribution were used. To aid visualization the \log_{10} transformation is applied to the relative test error values.




Figure S.2: Side-by-side boxplots of the logarithms, to base 10, of relative test error for the mean field variational Bayes Algorithm 3 for the simulation study described in the text. Each panel corresponds to a different combination of sample size and error standard deviation. Within each panel, the side-by-side boxplots compare relative test error as a function of the threshold parameter τ .

From Figures S.1 and S.2 we see that the relative test errors are lower for higher sample sizes, as expected. Somewhat counter-intuitively the relative test errors are lower for higher noise levels. However, comparisons of relative test error across different values of σ_ϵ are not clear-cut when the estimators are subject to bias. In addition, relative test errors are barely affected by the choice of the thresholding parameter τ . Lastly, the relative test errors based on mean field variational Bayes approximate inference are similar to those based on Markov chain Monte Carlo. It is interesting that this particular evaluation metric is not affected very much by the choices between Algorithms 2 and 3 and the value of the threshold parameter τ .

S.3.2 Detailed Computing Time Results

We also conducted some more detailed involving computing times. One simulation study looked into the effect of sample size, whilst another one investigated how the number of candidate predictors impacts computing times. The results are presented in this section.

S.3.2.1 Assessment of the Effect of Sample Size

Our first detailed computing time simulation study was concerned with the effect of sample size. We fixed the candidate predictor dimensions to be $(d_o, d_\bullet) = (0, 10)$ and let the sample size n to range over the set

$$\{10^k : k = 2, 3, 4, 5, 6\}.$$

The data were generated in a manner similar to that for the simulation studies described in Sections 3 and 4, with 100 replications.




Figure S.3: Side-by-side box plots of computing time in seconds versus sample size for generalized additive model selection via Algorithms 2 and 3, for the first simulation study described in Section S.3.2. Each axis uses a \log_{10} scale.

Figure S.3 summarizes the results using side-by-side boxplots of the logarithmically transformed computing times, broken down according to sample size, response type and whether or not Algorithm 2 or Algorithm 3 was used. The relationships between the mean logarithmic number of seconds and logarithmic sample size are approximately linear, which suggests a simple power law relationship between computing time and sample size. Simple linear regression analyses suggest that the power is close to 1 and, hence, mean computing time is roughly proportional to sample size.

Figure S.3 also shows that use of Algorithm 3 leads to an approximately ten-fold reduction in computing time compared with Algorithm 2. For example, when $n = 100,000$ the mean computing time of Algorithm 2 for Bernoulli responses is about 100 seconds. For Algorithm 3 it is only about 10 seconds.

S.3.2.2 Assessment of the Effect of the Number of Candidate Predictors

We also ran a simulation study concerned with the effect of the number of candidate predictors on computing time. The sample size was fixed at 5,000 and d_{\bullet} , the number of candidate predictors that could enter the model non-linearly, varied over the set

$$\{2^k : k = 1, 2, 3, 4, 5, 6\}.$$

We generated the data in a manner similar to that for the simulation studies described in Sections 3 and 4 and, again, obtained 100 replications.

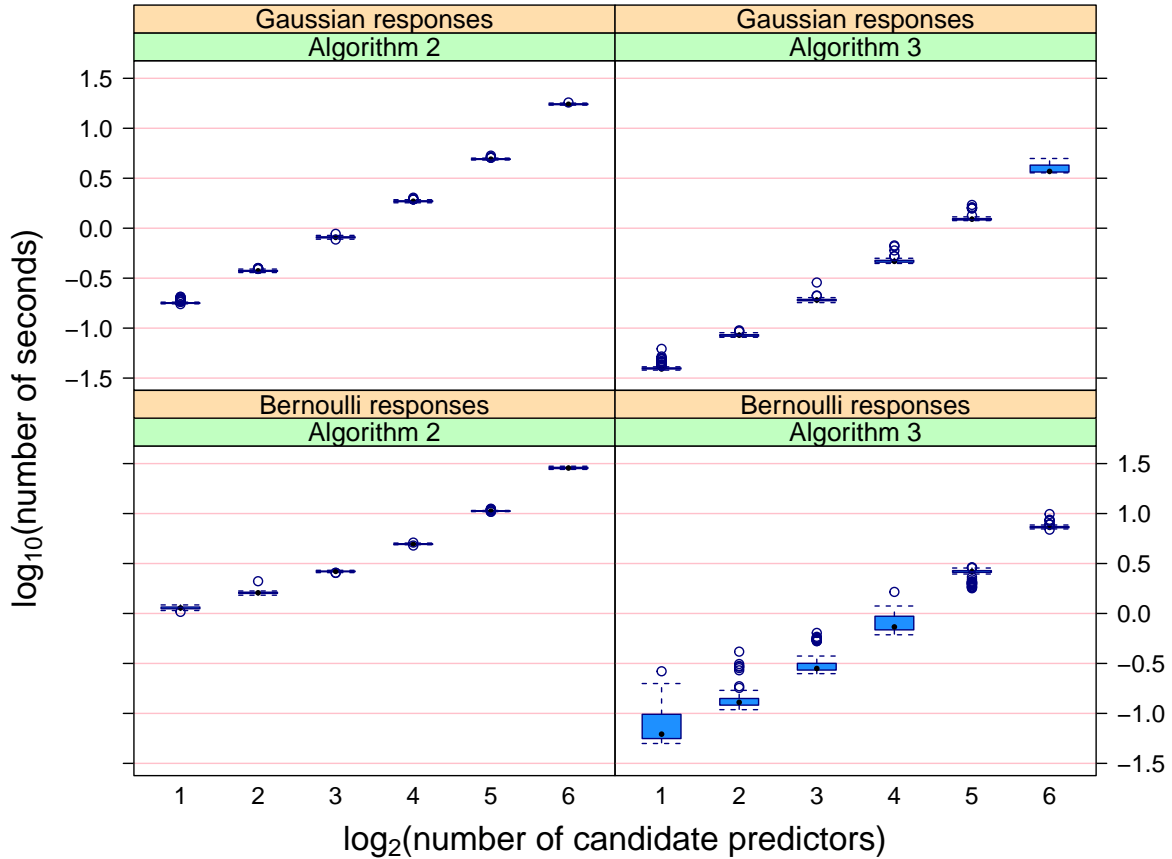


Figure S.4: Side-by-side box plots of computing time in seconds versus number of candidate predictors for generalized additive model selection via Algorithms 2 and 3, for the second simulation study described in Section S.3.2. The horizontal axis uses a \log_2 scale and the vertical axis uses a \log_{10} scale.

Figure S.4 summarises the results in similar way to Figure S.3. Once again, there is approximate linearity within each panel with logarithmic scales. Simple linear regression analyses of the data within each panel of Figure S.4 suggest that the mean computing time is approximately proportional to d_{\bullet}^{κ} , with κ dependent on the response distribution and fitting algorithm combination but within the interval (1.2, 1.5).

S.3.3 Hyperparameter Sensitivity Checks

Figure S.5 conveys the effect of the Half Cauchy distribution scale hyperparameters, denoted by s_{β} , s_{ε} and s_u in model (9), on the effect type misclassification rate. It is based on the simulation study set-up of Section 3.5 with the Markov chain Monte Carlo approach of Algorithm 2 and the threshold parameter τ set to our recommended default value of 0.5. The scale hyperparameters

ranged over the set

$$\{10^k : k = 1, 2, 3, 4\}.$$

Figure S.5 indicates that our default version of Algorithm 2 is not sensitive to the Half Cauchy distribution scale hyperparameter values.

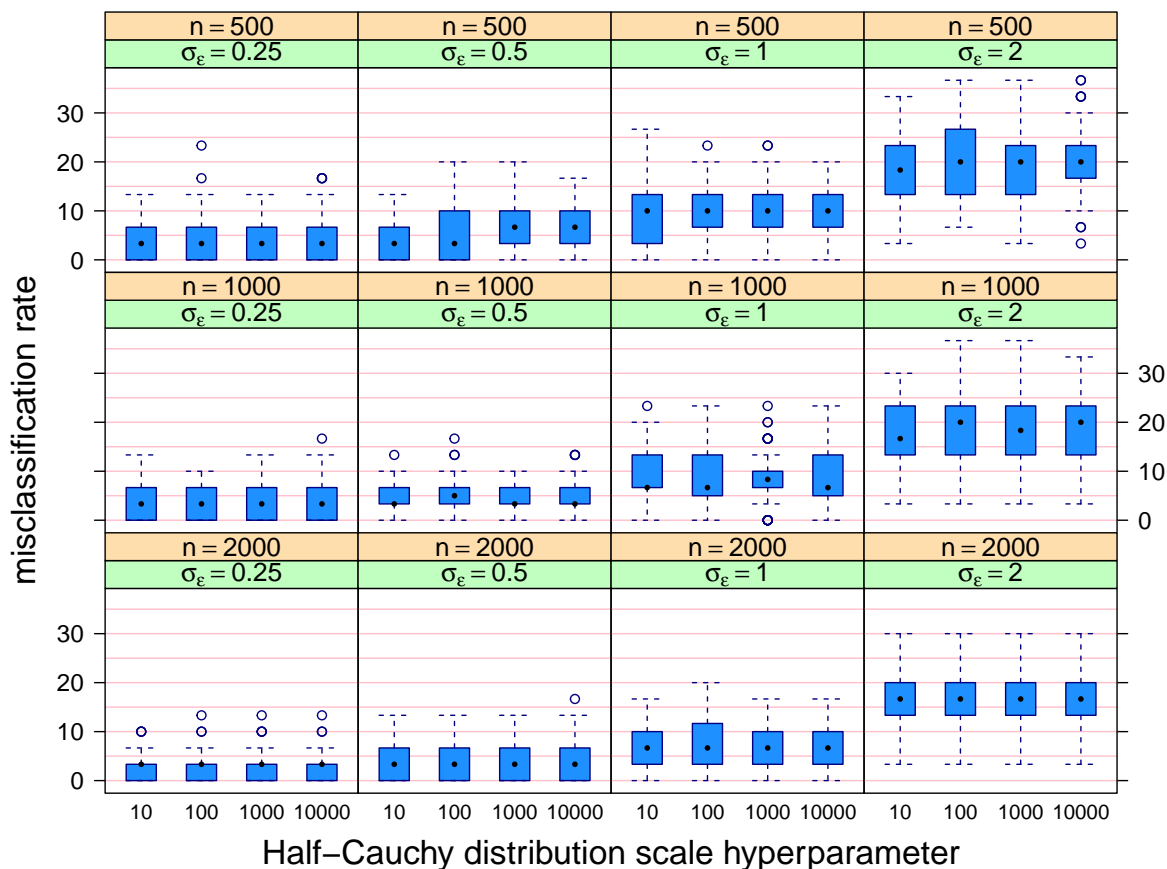


Figure S.5: Side-by-side boxplots of misclassification rate for varying values of the Half Cauchy distribution scale hyperparameter for the Gaussian response version of Algorithm 2, for the first simulation study described in Section S.3.3.

Figure S.6 is similar to Figure S.5, but is for the mean field variational Bayes approach used by Algorithm 3 with τ set to the default value of 0.1. Once again, low sensitivity to the Half Cauchy distribution scale hyperparameter values is exhibited.

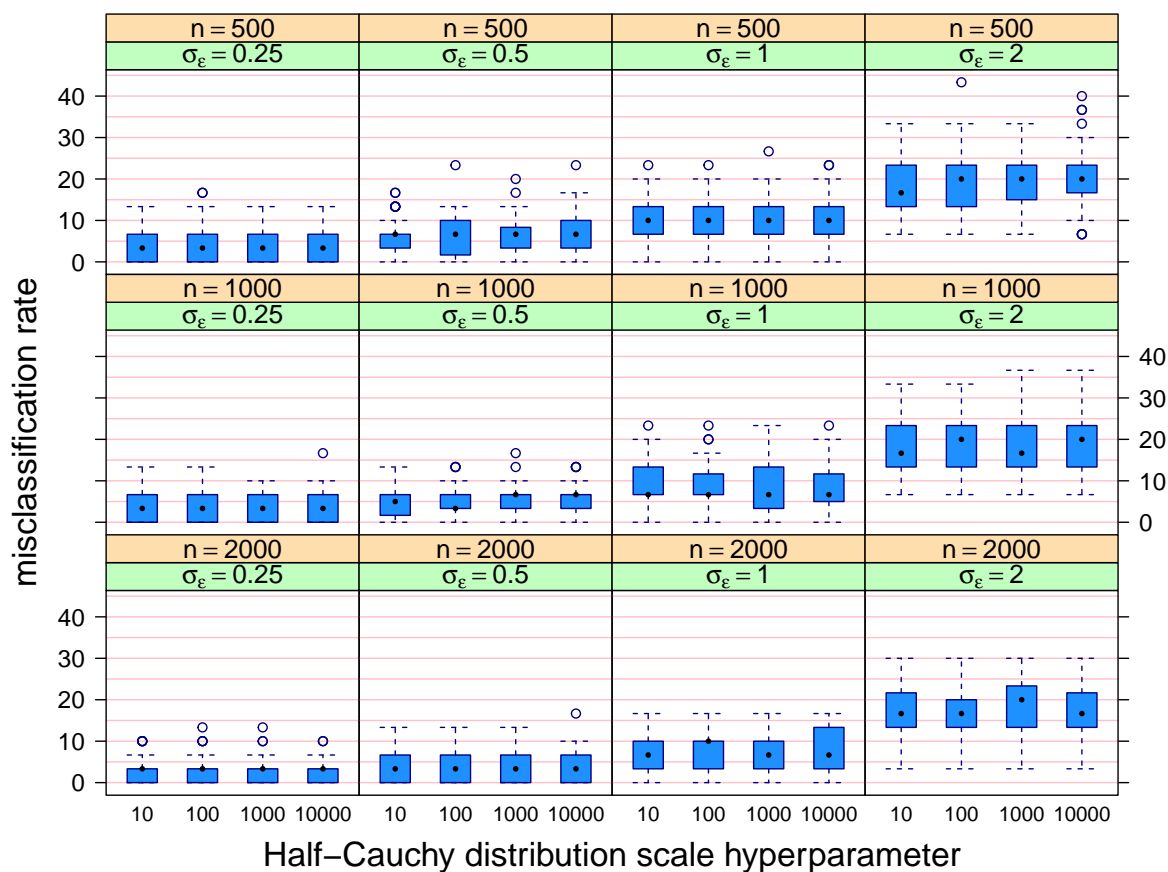


Figure S.6: Side-by-side boxplots of misclassification rate for varying values of the Half Cauchy distribution scale hyperparameter for the Gaussian response version of Algorithm 3, for the first simulation study described in Section S.3.3.

Reference

Hastie, T., Tibshirani, R. & Tibshirani, R. (2020). Best subset, forward stepwise of lasso? Analysis and recommendations based on extensive comparisons. *Statistical Science*, 35, 579–592.