

Bayesian Generalized Additive Model Selection Including a Fast Variational Option

BY VIRGINIA X. HE AND MATT P. WAND

University of Technology Sydney

25th January, 2022

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, 2021) 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 manual labor aspect of 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 auxiliary variables. The binary response case benefits from the Albert & Chib (1993) auxiliary vari-

able 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.

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, `ganselBayes`, 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.

2 Model Description

The original input data are as follows:

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

where, for each $1 \leq i \leq n$, $\overset{\circ}{\mathbf{x}}_i^{\text{orig}}$ denotes a $d_o \times 1$ vector of predictors that can only enter the model linearly (e.g. binary predictors) and $\dot{\mathbf{x}}_i^{\text{orig}}$ 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

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

where $\overset{\circ}{\mathbf{x}}_i$ and $\dot{\mathbf{x}}_i$ are standardized data versions of $\overset{\circ}{\mathbf{x}}_i^{\text{orig}}$ and $\dot{\mathbf{x}}_i^{\text{orig}}$. 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

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

Generalized additive models involve linear predictors having generic form

$$\beta_0 + \sum_{j=1}^{d_o} \beta_j \overset{\circ}{x}_{ji} + \sum_{j=1}^{d_\bullet} f_j(\dot{x}_{ji}) \tag{1}$$

for scalar parameters $\beta_0, \dots, \beta_{d_o}$ and the f_j are smooth real-valued functions over an interval containing the \dot{x}_{ji} data.

2.1 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, Φ denotes the $N(0, 1)$ cumulative distribution function.

distribution	density/probability function in x	abbreviation
Bernoulli	$\varphi^x(1 - \varphi)^{1-x}; \quad x = 0, 1; 0 < \varphi < 1$	Bernoulli(φ)
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)
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.

2.2 Model for the f_j Functions

Let $\dot{x}_1, \dots, \dot{x}_n$ be a typical continuous predictor data sample. Our smooth function models take 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

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

for some scalar-valued function h . In other words, we apply linearly 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 2.4 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 5.2. We use the descriptor *canonical Demmler-Reinsch basis* for this type of spline basis.

2.3 Model for a Linear Coefficient

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

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

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. Ishwarya & Rao, 2005).

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 Section 3.6 for principled model selection strategies.

The scale and mixture parameters in (6) have the following prior distributions:

$$\sigma_\beta \sim \text{Half-Cauchy}(s_\beta) \quad \text{and} \quad \rho_\beta \sim \text{Beta}(A_\beta, B_\beta)$$

for hyperparameters $s_\beta, A_\beta, B_\beta > 0$.

2.4 Model for a Spline Coefficients Vector

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

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

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 $\|\mathbf{u}\| \equiv (\mathbf{u}^T \mathbf{u})^{1/2}$ and 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 and mixture parameters have the following prior distributions:

$$\sigma_u \sim \text{Half-Cauchy}(s_u) \quad \text{and} \quad \rho_u \sim \text{Beta}(A_u, B_u)$$

for hyperparameters $s_u, A_u, B_u > 0$.

2.5 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 vector $\gamma_u = (\gamma_{u1}, \dots, \gamma_{uK})$, (7) is equivalent to

$$\mathbf{u} = \gamma_u \odot \tilde{\mathbf{u}}, \quad \gamma_{uj} \stackrel{\text{ind.}}{\sim} \text{Bernoulli}(\rho_u), \quad \tilde{\mathbf{u}}|b_u, \sigma_u^2 \sim N(0, \sigma_u^2 \mathbf{I}/b_u) \\ \text{and } b_u \sim \text{Inverse-Gamma}(1, \frac{K+1}{2})$$

courtesy of a result provided in Section 3.1 of Kyung *et al.* (2010) for the $\rho_u = 1$ case. Here, and elsewhere, the notation $\stackrel{\text{ind.}}{\sim}$ is an abbreviation for ‘‘distributed independently as’’.

2.6 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 2.2–2.4 across each of d_\circ entries of the $\hat{\mathbf{x}}_i$ and d_\bullet entries of $\hat{\mathbf{x}}_i$. Let β be the $(d_\circ + d_\bullet) \times 1$ 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 2.5. The resultant full model is:

$$\mathbf{y}|\beta_0, \gamma_\beta, \tilde{\beta}, \gamma_{u1}, \dots, \gamma_{ud_\bullet}, \tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_{d_\bullet}, \sigma_\varepsilon^2 \sim \\ N \left(\mathbf{1}_n \beta_0 + \mathbf{X}(\gamma_\beta \odot \tilde{\beta}) + \sum_{j=1}^{d_\bullet} \mathbf{Z}_j(\gamma_{uj} \odot \tilde{\mathbf{u}}_j), \sigma_\varepsilon^2 \mathbf{I}_n \right), \quad \beta_0 \sim N(0, \sigma_{\beta_0}^2), \\ \sigma_\varepsilon^2|a_\varepsilon \sim \text{Inverse-Gamma}(\frac{1}{2}, 1/a_\varepsilon), \quad a_\varepsilon \sim \text{Inverse-Gamma}(\frac{1}{2}, 1/s_\varepsilon^2), \\ \gamma_{\beta j}|\rho_\beta \stackrel{\text{ind.}}{\sim} \text{Bernoulli}(\rho_\beta), \quad \tilde{\beta}_j|\sigma_\beta^2, b_{\beta j} \stackrel{\text{ind.}}{\sim} N(0, \sigma_\beta^2/b_{\beta j}), \quad 1 \leq j \leq d_\circ + d_\bullet, \\ b_{\beta j} \stackrel{\text{ind.}}{\sim} \text{Inverse-Gamma}(1, \frac{1}{2}), \quad 1 \leq j \leq d_\circ + d_\bullet, \\ \gamma_{ujk}|\rho_{uj} \stackrel{\text{ind.}}{\sim} \text{Bernoulli}(\rho_{uj}), \quad 1 \leq j \leq d_\bullet, \quad 1 \leq k \leq K_j, \\ \tilde{\mathbf{u}}_j|\sigma_{uj}^2, b_{uj} \stackrel{\text{ind.}}{\sim} N(\mathbf{0}, (\sigma_{uj}^2/b_{uj}) \mathbf{I}_{K_j}), \quad b_{uj} \stackrel{\text{ind.}}{\sim} \text{Inverse-Gamma}(\frac{1}{2}(K_j + 1), \frac{1}{2}), \quad 1 \leq j \leq d_\bullet, \\ \sigma_\beta^2|a_\beta \sim \text{Inverse-Gamma}(\frac{1}{2}, 1/a_\beta), \quad a_\beta \sim \text{Inverse-Gamma}(\frac{1}{2}, 1/s_\beta^2), \\ \sigma_{uj}^2|a_{uj} \stackrel{\text{ind.}}{\sim} \text{Inverse-Gamma}(\frac{1}{2}, 1/a_{uj}), \quad a_{uj} \stackrel{\text{ind.}}{\sim} \text{Inverse-Gamma}(\frac{1}{2}, 1/s_u^2), \quad 1 \leq j \leq d_\bullet, \\ \rho_\beta \sim \text{Beta}(A_\beta, B_\beta) \quad \text{and} \quad \rho_{uj} \stackrel{\text{ind.}}{\sim} \text{Beta}(A_u, B_u), \quad 1 \leq j \leq d_\bullet. \quad (9)$$

The full set of hyperparameters in (9) is:

$$\sigma_{\beta_0}, s_\beta, s_\varepsilon, s_u, A_\beta, B_\beta, A_u, B_u > 0.$$

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

2.7 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

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




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.

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

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

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

$$\mathbf{c} | \beta_0, \gamma_\beta, \tilde{\boldsymbol{\beta}}, \gamma_{u1}, \dots, \gamma_{ud}, \tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_d \sim N \left(\mathbf{1}_n \beta_0 + \mathbf{X}(\gamma_\beta \odot \tilde{\boldsymbol{\beta}}) + \sum_{j=1}^d \mathbf{Z}_j(\gamma_{uj} \odot \tilde{\mathbf{u}}_j), \mathbf{I}_n \right). \quad (12)$$

The prior distributions on $\beta_0, \gamma_\beta, \tilde{\boldsymbol{\beta}}, \gamma_{u1}, \dots, \gamma_{ud}$ and $\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_d$ 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 \mathbf{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 2.7, 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




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 \mathbf{c} corresponding to the Albert-Chib auxiliary variables is added to the position held by \mathbf{y} in the Gaussian response graph and the binary response observed data node \mathbf{y} a parent of \mathbf{c} .

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 by 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 \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.

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

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 canonical Demmler-Reinsch basis for the predictor
data vector $\hat{\mathbf{x}}_j$, using the construction described in Appendix A
 $\mathbf{Z} \leftarrow [\mathbf{Z}_1 \cdots \mathbf{Z}_{d_\bullet}]$; $\mathbf{XTy} \leftarrow \mathbf{X}^T \mathbf{y}$; $\mathbf{XTX} \leftarrow \mathbf{X}^T \mathbf{X}$; $\mathbf{ZTX} \leftarrow \mathbf{Z}^T \mathbf{X}$; $\mathbf{ZTZ} \leftarrow \mathbf{Z}^T \mathbf{Z}$
Outputs: $\mathbf{y}, \mathbf{X}, \mathbf{Z}_1, \dots, \mathbf{Z}_{d_\bullet}, \mathbf{XTy}, \mathbf{XTX}, \mathbf{ZTX}, \mathbf{ZTZ}, \text{mean}(\mathbf{y}^{\text{orig}}), \text{st.dev}(\mathbf{y}^{\text{orig}}),$
 $\left\{ (\text{mean}(\hat{\mathbf{x}}_j^{\text{orig}}), \text{st.dev.}(\hat{\mathbf{x}}_j^{\text{orig}})) : 1 \leq j \leq d_o \right\}, \left\{ (\text{mean}(\hat{\mathbf{x}}_j^{\text{orig}}), \text{st.dev.}(\hat{\mathbf{x}}_j^{\text{orig}})) : 1 \leq j \leq d_\bullet \right\}$

using the following notational convention: e_r is a column vector of appropriate length with r th entry equal to 1 and zeroes elsewhere. For example, the

k th row of $\mathbf{ZTX}^{(j)}$ is $e_k^T \mathbf{ZTX}^{(j)}$ where e_k is the $K_j \times 1$ vector with k th entry 1 and 0 elsewhere.

Implementations of the upcoming algorithms normally would not require explicit calculation and storage of e_r vectors and, instead, array subsetting code specific to the programming language can be used. However, for algorithm listing 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 $\mathbf{ZTy}_{\text{adj}}$.

For any column vector \mathbf{a} we let \mathbf{a}_{-j} denote the column vector with the j entry of \mathbf{a} omitted. As a reminder, $\|\mathbf{a}\| \equiv (\mathbf{a}^T \mathbf{a})^{1/2}$ denotes the Euclidean norm of \mathbf{a} . 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} .

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 (2021) 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 2.7, 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.

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

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

where, for example, $\tilde{\mathbf{u}} \equiv (\tilde{u}_1, \dots, \tilde{u}_{d_\bullet})$ and $\mathbf{b}_\beta \equiv (b_{\beta 1}, \dots, b_{\beta(d_o + d_\bullet)})$. 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 $q(\beta_0)q(\tilde{\beta})q(\tilde{\mathbf{u}})$ 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 restriction:

$$q(\gamma_\beta) = \prod_{j=1}^{d_o + d_\bullet} q(\gamma_{\beta j}). \quad (15)$$

In addition, we have

$$q(\tilde{\mathbf{u}}) = \prod_{j=1}^{d_\bullet} q(\tilde{u}_j) \quad \text{and} \quad q(\gamma_u) = \prod_{j=1}^{d_\bullet} q(\gamma_{uj}) = \prod_{j=1}^{d_\bullet} \prod_{k=1}^{K_j} q(\gamma_{ujk}) \quad (16)$$

although some of these product density restrictions are *induced* rather than imposed. This subtlety is explained in Section 10.2.5 of Bishop (2006).

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 left-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

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: \mathbf{XTy} , \mathbf{XTX} , \mathbf{ZTy} , \mathbf{ZTX} , \mathbf{ZTZ} .

Hyperparameter Inputs: σ_{β_0} , s_β , s_ε , s_u , A_β , B_β , A_u , $B_u > 0$.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

$\rho_\beta^{[g]} \sim \text{Beta}\left(A_\beta + \mathbf{1}_{d_o + d_\bullet}^T \gamma_\beta^{[g]}, B_\beta + d_o + d_\bullet - \mathbf{1}_{d_o + d_\bullet}^T \gamma_\beta^{[g]}\right)$

continued on a subsequent page ...

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:

$$\begin{aligned} \mathbf{q}(\tilde{\beta}) \text{ has a } N(\boldsymbol{\mu}_{\mathbf{q}(\tilde{\beta})}, \boldsymbol{\Sigma}_{\mathbf{q}(\tilde{\beta})}) \text{ density function, and} \\ \mathbf{q}(\sigma_\varepsilon^2) \text{ has an Inverse-Gamma}(\kappa_{\mathbf{q}(\sigma_\varepsilon^2)}, \lambda_{\mathbf{q}(\sigma_\varepsilon^2)}) \text{ density function.} \end{aligned} \quad (17)$$

Algorithm 2 continued. This is a continuation of the description of this algorithm that continues 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 $j = 1, \dots, d_\bullet$:

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

$$\boldsymbol{\omega}_7 \leftarrow \left\{ \boldsymbol{\gamma}_{\mathbf{u}_j}^{[g-1]} \odot \mathbf{w}_{\mathbf{z}_j} / (\sigma_\varepsilon^2)^{[g-1]} \right\} + \left\{ b_{\mathbf{u}_j}^{[g-1]} \mathbf{1}_{K_j} / (\sigma_{\mathbf{u}_j}^2)^{[g-1]} \right\}$$

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

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

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

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

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

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

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

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

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

For $k = 1, \dots, K_j$:

$$\omega_9 \leftarrow \text{logit} \left(\rho_{\mathbf{u}_j}^{[g-1]} \right) - \frac{1}{2} \left\{ (\tilde{u}_{jk}^{[g]})^2 \mathbf{e}_k^T \mathbf{w}_{\mathbf{z}_j} - 2\tilde{u}_{jk}^{[g]} \mathbf{e}_k^T \boldsymbol{\omega}_8 \right\} / (\sigma_\varepsilon^2)^{[g-1]}$$

$$(\boldsymbol{\gamma}_{\mathbf{u}_j}^{\text{curr}})_k \sim \text{Bernoulli}(\text{expit}(\omega_9))$$

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

For $j = 1, \dots, d_\bullet$: $\rho_{\mathbf{u}_j}^{[g]} \sim \text{Beta} \left(A_u + \mathbf{1}_{K_j}^T \boldsymbol{\gamma}_{\mathbf{u}_j}^{[g]}, B_u + K_j - \mathbf{1}_{K_j}^T \boldsymbol{\gamma}_{\mathbf{u}_j}^{[g]} \right)$

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

If **responseType** is Gaussian then

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

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

If **responseType** is Bernoulli then

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

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

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

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

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

The optimal Inverse Gamma shape parameter $\kappa_{\mathbf{q}(\sigma_\varepsilon^2)}$ has explicit solution $\frac{1}{2}(n + 1)$. However, the equations for the optimal values of $\boldsymbol{\mu}_{\mathbf{q}(\tilde{\boldsymbol{\beta}})}$, $\boldsymbol{\Sigma}_{\mathbf{q}(\tilde{\boldsymbol{\beta}})}$ and $\lambda_{\mathbf{q}(\sigma_\varepsilon^2)}$ are interdependent and iteration is required to obtain their optimal values. Algorithm 3 lists the full set of steps required to obtain all \mathbf{q} -density parameters, with notation similar to that used in (17) for the other \mathbf{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 $\log \underline{p}(\mathbf{y}; \mathbf{q})$. Section 5.2 in the appendix contains an explicit expression for the approximate marginal log-likelihood for the Section 2 models under product restrictions (14)–(16).

3.5 Hyperparameter Default Values

With all input data standardized, the fitting algorithms are scale invariant and the hyperparameters can be set to fixed constants. With noninformativity in mind, the default values of the hyperparameters are as follows:

$$\sigma_{\beta_0} = 10^5, \quad s_\beta = s_\varepsilon = s_u = 1000, \quad A_\beta = B_\beta = A_u = B_u = 1.$$

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

3.6 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.

3.6.1 Deciding Between an Effect Being Zero or Linear

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

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

and the posterior distribution 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 $P(\gamma_\beta = 0 | \mathbf{y}) > \frac{1}{2}$. However, in the interests of parsimony, less stringent rules are worth considering. Rather than thresholding $P(\gamma_\beta = 0 | \mathbf{y})$ at $\frac{1}{2}$, we will consider a family of rules indexed by a threshold parameter $\tau \in [0, 1)$. For the case of τ being zero or very small it is useful in practice to have an “effective zero” parameter, $\varepsilon_{\text{eff.zero}}$, which is set to a small positive number such as 0.00001. Note that $\tau = 0$ corresponds to maximum a posteriori estimation of β . After fixing τ and $\varepsilon_{\text{eff.zero}}$, our strategy for deciding between an effect being zero or linear is

the effect is zero if $P(\gamma_\beta = 0 | \mathbf{y}) > \max(\tau, \varepsilon_{\text{eff.zero}})$, otherwise the effect is linear.

3.6.2 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 $\hat{\mathbf{x}}_j$ predictors. Since $\mathbf{u} = \gamma_u \odot \tilde{\mathbf{u}}$, where the entries of γ_u are binary and the entries of $\tilde{\mathbf{u}}$ are continuous,

$$P(\mathbf{u} = \mathbf{0} | \mathbf{y}) = P(\gamma_{u1} = 0, \dots, \gamma_{uK} = 0 | \mathbf{y}).$$

Therefore, after fixing τ and $\varepsilon_{\text{eff.zero}}$, our strategy for deciding between an effect being zero, linear or non-linear is:

the effect is zero if $\min\{P(\gamma_\beta = 0 | \mathbf{y}), P(\gamma_{u1} = 0 | \mathbf{y}), \dots, P(\gamma_{uK} = 0 | \mathbf{y})\} > \max(\tau, \varepsilon_{\text{eff.zero}})$,

the effect is linear if $\begin{cases} P(\gamma_\beta = 0 | \mathbf{y}) \leq \max(\tau, \varepsilon_{\text{eff.zero}}) \text{ and} \\ \min\{P(\gamma_{u1} = 0 | \mathbf{y}), \dots, P(\gamma_{uK} = 0 | \mathbf{y})\} > \max(\tau, \varepsilon_{\text{eff.zero}}), \end{cases}$

otherwise the effect is non-linear.

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: **XTy**, **XTX**, **ZTy**, **ZTX**, **ZTZ**

Hyperparameter Inputs: σ_{β_0} , s_β , s_ε , s_u , A_β , B_β , A_u , $B_u > 0$.

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

Initialize: $\boldsymbol{\mu}_{\text{q}(\gamma_\beta)} \leftarrow \frac{1}{2} \mathbf{1}_{d_o + d_\bullet}$; $\boldsymbol{\mu}_{\text{q}(\tilde{\beta})} \leftarrow \mathbf{0}_{d_o + d_\bullet}$; $\boldsymbol{\mu}_{\text{q}(\tilde{u}_j)} \leftarrow \mathbf{0}_{K_j}$
 $\mu_{\text{q}(1/a_\varepsilon)} \leftarrow 1$; $\mu_{\text{q}(1/\sigma_\varepsilon^2)} \leftarrow 1$;
 $\boldsymbol{\mu}_{\text{q}(\gamma_{u_j})} \leftarrow \frac{1}{2} \mathbf{1}_{K_j}$; $\boldsymbol{\mu}_{\text{q}(\tilde{u}_j)} \leftarrow \mathbf{0}_{K_j}$, $\sigma_{\text{q}(\tilde{u}_j)}^2 \leftarrow \mathbf{1}_{K_j}$ $1 \leq j \leq d_\bullet$;
 $\mu_{\text{q}(1/a_\beta)} \leftarrow 1$; $\mu_{\text{q}(1/\sigma_\beta^2)} \leftarrow 1$; $\mu_{\text{q}(1/a_{u_j})}$, $\mu_{\text{q}(1/\sigma_{u_j}^2)}$, $1 \leq j \leq d_\bullet$;
 $\mu_{\text{q}(b_\beta)} \leftarrow \mathbf{1}_{d_o + d_\bullet}$; $\mu_{\text{q}(b_{u_j})} \leftarrow 1$, $1 \leq j \leq d_\bullet$
 $\kappa_{\text{q}(\sigma_\beta^2)} \leftarrow \frac{1}{2}(d_o + d_\bullet + 1)$; $\kappa_{\text{q}(a_\beta)} \leftarrow 1$,
 $\kappa_{\text{q}(\sigma_{u_j}^2)} \leftarrow \frac{1}{2}(K_j + 1)$, $\kappa_{\text{q}(a_{u_j})} \leftarrow 1$, $1 \leq j \leq d_\bullet$;
 $\kappa_{\text{q}(\sigma_\varepsilon^2)} \leftarrow \frac{1}{2}(n + 1)$; $\kappa_{\text{q}(a_\varepsilon)} \leftarrow 1$

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

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

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

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

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

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

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

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

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

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

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

$$\mu_{\text{q}(\gamma_{\beta_\bullet})} \leftarrow \sum_{j=1}^{d_o + d_\bullet} \mu_{\text{q}(\gamma_{\beta_j})} ; A_{\text{q}(\rho_\beta)} \leftarrow A_\beta + \mu_{\text{q}(\gamma_{\beta_\bullet})} ; B_{\text{q}(\rho_\beta)} \leftarrow B_\beta + d_o + d_\bullet - \mu_{\text{q}(\gamma_{\beta_\bullet})}$$

$$\mu_{\text{q}(\text{logit}(\rho_\beta))} \leftarrow \text{digamma}(A_{\text{q}(\rho_\beta)}) - \text{digamma}(B_{\text{q}(\rho_\beta)})$$

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

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

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

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

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

continued on a subsequent page ...

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

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

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

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

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

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

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

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

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

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

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

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

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

$$\mu_{q(\gamma_{u_j \bullet})} \leftarrow \sum_{k=1}^{K_j} \mu_{q(\gamma_{u_{jk}})} \quad ; \quad A_{q(\rho_{u_j})} \leftarrow A_u + \mu_{q(\gamma_{u_j \bullet})} \quad ; \quad B_{q(\rho_{u_j})} \leftarrow B_u + K_j - \mu_{q(\gamma_{u_j \bullet})}$$

$$\mu_{q(\text{logit}(\rho_{u_j}))} \leftarrow \text{digamma}(A_{q(\rho_{u_j})}) - \text{digamma}(B_{q(\rho_{u_j})})$$

For $k = 1, \dots, K_j$:

$$\boldsymbol{\omega}_{19} \leftarrow \left(\mu_{q(\tilde{u}_{jk})}^2 + \boldsymbol{\sigma}_{q(\tilde{u}_{jk})}^2 \right) \mathbf{e}_k^T \boldsymbol{w}_{Zj} - 2 \mu_{q(\tilde{u}_{jk})} \mathbf{e}_k^T \boldsymbol{\omega}_{18}$$

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

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

If **responseType** is Gaussian then

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

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

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

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

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

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

If **responseType** is Bernoulli then

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

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

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

Outputs: All q -density parameters.

It is apparent from these rules that the parameter $\tau \in [0, 1)$ controls the degree of sparsity in the selected model. 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.

3.6.3 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 . We have made inroads into this problem by running numerous simulation studies, all of which point to good defaults for τ being considerably lower than the natural choice of $\tau = \frac{1}{2}$. We present the results of one such study here.

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. The non-linear effects corresponded to quintic polynomials with randomly generated coefficients. The sample sizes varied over $n \in \{500, 1000, 2000\}$ and, for the Gaussian response case, the error standard deviations varied over $\sigma_\varepsilon \in \{0.25, 0.5, 1, 2\}$. 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 $\varepsilon_{\text{toler.}} = 10^{-8}$. Model selection was applied according to the rules of Section 3.6 with $\tau \in \{0, 0.05, 0.1, 0.25, 0.5\}$. 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 $\sigma_\varepsilon \geq 1$ it is advantageous to have τ equal to a value around 0.1, and that $\tau = 0.5$ leads to worse performance. 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. The difference between $\tau = 0$ and $\tau > 0$ is quite striking. We conjecture that mean field approximations have a detrimental affect on the $\tau > 0$ situations, and for $\tau = 0$ this effect is much less. Acknowledging the limitations of a single simulation study, a recommended default for τ in the mean field variational Bayes case is $\tau = 0$, corresponding to maximum a posteriori estimation.

We also ran simulation studies for the Bernoulli response case, with a similar design to the Gaussian study. The recommendations of $\tau = 0.1$ for Markov chain Monte Carlo and $\tau = 0$ for mean field variational Bayes were also supported by that study.

3.7 Package in the R Language

The R package `gamselBayes` (He & Wand, 2022) 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 3.6.3. In keeping with the findings of that section, in Algorithm 2 the threshold parameter was set to




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 .

$\tau = 0.1$ and for Algorithm 3 it was set to $\tau = 0$. An effective zero value of $\varepsilon_{\text{eff.zero}} = 0.00001$ was used.

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

1. `spikeSlabGAM` (Scheipl, 2020), 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).
2. `gamsel` (Chouldechova & Hastie, 2018), 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.

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.

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




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 .

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` approachers 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 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

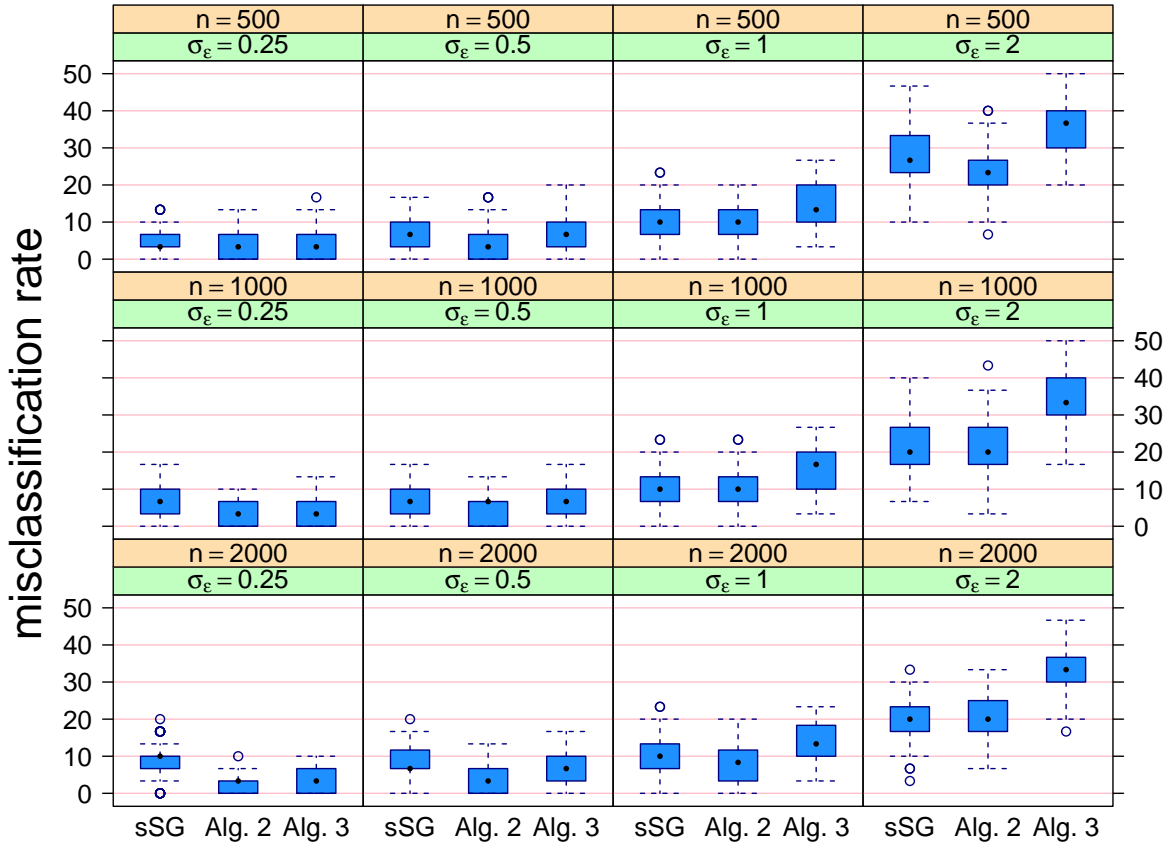


Figure 5: Side-by-side boxplots of the misclassification rates for the comparative performance simulation study described in the text in the case of the response variable being Gaussian. Each panel corresponds to a different combination of sample size and error standard deviation. Within each panel, 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).




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).

computer with 8 gigabytes of memory and a 2.2. gigahertz processor. Algorithms 2 and 3 were

	gamsel	spikeSlabGAM	Algorithm 2	Algorithm 3
10th percentile	5.56	119	2.55	0.69
50th percentile	11.30	243	3.36	1.05
90th percentile	25.40	526	5.23	2.21

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.

implemented in a version of the C++ language. 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 5 seconds for 90% of the fits in the simulation study. The faster variational approach of Algorithm 3 only required 1–2 seconds of computing time for most of the fits. Therefore, the new approaches within the `gamselBayes` R package have very scalability for the generalized additive model selection problem.

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, 2020) 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 3.6 with $\tau = 0.1$ 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 indicated good convergence.

As is apparent from Table 3, the selected model has 6 linear effects, 2 non-linear effects and 10 candidate 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.

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?	zero	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 3.6. 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.

predictor	posterior mean	95% credible interval
indicator of bad public credit record	0.7328	(0.4806, 0.9738)
indicator of denied mortgage insurance	2.7223	(2.1075, 3.4302)
indicator of applicant being single	0.1822	(0.0000, 0.3346)
indicator of applicant being black	0.3507	(0.1329, 0.5565)
indicator of credit score equalling 1	−0.7023	(−0.9210, −0.4683)
indicator of credit score equalling 2	−0.3196	(−0.5991, −0.0421)

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.

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 restrict exclude age at sale from having a non-linear effect.

Since this generalized additive model selection problem involves a relatively large sample 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 2.2 gigahertz processor and 8 gigabytes of random access memory, mean field Bayes variational fitting took 13 seconds. The rules of Section 3.6 were applied with $\tau = 0$, corresponding to maximum a posteriori estimation of the γ_β and γ_{u_j} indicator variables. This resulted in 14 predictors being selected as having a linear effect and one predictor, the acquisition cost paid for the car at the time of purchase, having a




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 3.6 with $\tau = 0.1$. 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.

predictor	posterior mean	95% credible interval
indicator purchased in 2010	0.1018	(0.08636, 0.1172)
age at sale (years)	0.09641	(0.09095, 0.1019)
indic. make is Chevrolet	-0.1410	(-0.1622, -0.1206)
indic. make is Chrysler	0.08772	(0.06115, 0.1141)
indic. trim is 'Bas'	0.06660	(0.04691, 0.08613)
indic. manual transmission	-0.1596	(-0.1993, -0.1196)
indic. has alloy wheels	-1.511	(-1.544, -1.477)
indic. has wheel covers	-1.586	(-1.620, -1.552)
indic. medium-sized vehicle	-0.0721	(-0.08925, -0.05489)
indic. sports utility vehicle	0.1835	(0.1583, 0.2092)
indic. purch. in Texas	0.09932	(0.07983, 0.1187)
indic. purch. in Florida	-0.1144	(-0.1365, -0.09267)
indic. purch. in North Carolina	-0.1069	(-0.1322, -0.08128)
odometer reading (miles)	3.699×10^{-6}	$(3.148, 4.248) \times 10^{-6}$

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 obtain from Algorithm 3 for the car auction example.

non-linear effect. Thirty-four of the 49, or 69%, of candidate predictors were discarded, leading to a selected model that is quite parsimonious.

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 only selected non-linear effect, which is the impact of the probability of a bad buy as a function of the acquisition cost paid for the car at the time of purchase in U.S.




Figure 8: The only estimated non-linear effects for the car auction example from application of Algorithm 3 and effect type estimation rules of Section 3.6 with $\tau = 0$. The curve is the slice of estimated probability of bad buy 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.

dollars. It 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.

As a type of check, we also applied the Markov chain Monte Carlo Algorithm 2 to the same data set. This resulted in 13 of the 14 predictors in Table 5 being selected. The indicator of the trim level being ‘Bas’ was not selected by this application of Algorithm 2. Also, an additional non-linear effect, corresponding to the warranty cost of the vehicle, was selected. In summary, each approach selected 14 predictors but had two predictors and their effect types differing. This suggests reasonable accuracy of the faster variational approach for this example.

Acknowledgement

This research was supported by Australian Research Council grant DP180100597.

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. (2021). The R package ‘sn’: The Skew-Normal and related distributions such as the Skew-t and the Unified Skew-Normal (version 2.0.0). <http://azzalini.stat.unipd.it/SN>
- Bishop, C.M. (2006). *Pattern Recognition and Machine Learning*. New York: Springer.
- Chouldechova, A. & Hastie, T. (2015). Generalized additive model selection. <https://arXiv.org/abs/1506.03850v2>

- Chouldechova, A. & Hastie, T. (2018). `gamsel 1.8`: Fit regularization path for generalized additive models. R package. <https://r-project.org>
- Croissant, Y. (2020). `Ecdat 0.3`: Data sets for econometrics. R package. <https://r-project.org>
- Gelfand, A.E. & Smith, A.F.M. (1990). Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association*, **85**, 398–409.
- 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. (2022). `gamselBayes`: Bayesian generalized additive model selection. R package version 1.0. <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.
- Michael, J.R., Schucany, W.R. & Haas, R.W. (1976). Generating random variates using transformations with multiple roots. *The American Statistician*, **30**, 88–90.
- 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 (2021). 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. (2020). `spikeSlabGAM 1.1`: Bayesian variable selection and model choice for generalized additive mixed models. R package.

- 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.
- 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.

Appendix A: The Canonical Demmler-Reinsch Spline Basis

Let $\boldsymbol{x} = (x_1, \dots, x_n)$ 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 \boldsymbol{Z} containing canonical Demmler-Reinsch basis functions of the entries of \boldsymbol{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 \boldsymbol{Z} in that section’s equation (6), which contains canonical O’Sullivan spline basis functions. Denote this matrix by $\boldsymbol{Z}_{\text{OS}}$ and note that it has dimension $n \times K$.
- (2) Form the matrix $\boldsymbol{C}_{\text{OS}} = [\mathbf{1}_n \ \boldsymbol{x} \ \boldsymbol{Z}_{\text{OS}}]$ and set $\boldsymbol{D} = \text{diag}(0, 0, \mathbf{1}_K)$.
- (3) Obtain the singular value decomposition of $\boldsymbol{C}_{\text{OS}}$:

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

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

- (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 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:

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

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 κ_{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 κ_{inter} can be accessed by downloading the accompanying `gamselBayes` package. Assuming that the `gamselBayes` package is installed, the relevant function is `gamselBayes:::ZcDR()`.

Appendix B: Approximate Marginal Log-Likelihood Expressions

First define

$$\begin{aligned} \log \underline{p}(\mathbf{y}; \mathbf{q}, \text{BASE}) &= E_{\mathbf{q}}[\log\{\mathbf{p}(\beta_0)\}] - E_{\mathbf{q}}[\log\{\mathbf{q}(\beta_0)\}] \\ &+ E_{\mathbf{q}}[\log\{\mathbf{p}(\gamma_\beta|\rho_\beta)\}] - E_{\mathbf{q}}[\log\{\mathbf{q}(\gamma_\beta)\}] + E_{\mathbf{q}}[\log\{\mathbf{p}(\rho_\beta)\}] - E_{\mathbf{q}}[\log\{\mathbf{q}(\rho_\beta)\}] \\ &+ E_{\mathbf{q}}[\log\{\mathbf{p}(\tilde{\beta}|\mathbf{b}_\beta, \sigma_\beta^2)\}] - E_{\mathbf{q}}[\log\{\mathbf{q}(\tilde{\beta})\}] + E_{\mathbf{q}}[\log\{\mathbf{p}(\mathbf{b}_\beta)\}] - E_{\mathbf{q}}[\log\{\mathbf{q}(\mathbf{b}_\beta)\}] \\ &+ E_{\mathbf{q}}[\log\{\mathbf{p}(\sigma_\beta^2|a_\beta)\}] - E_{\mathbf{q}}[\log\{\mathbf{q}(\sigma_\beta^2)\}] + E_{\mathbf{q}}[\log\{\mathbf{p}(a_\beta)\}] - E_{\mathbf{q}}[\log\{\mathbf{q}(a_\beta)\}] \\ &+ E_{\mathbf{q}}[\log\{\mathbf{p}(\gamma_u|\rho_u)\}] - E_{\mathbf{q}}[\log\{\mathbf{q}(\gamma_u)\}] + E_{\mathbf{q}}[\log\{\mathbf{p}(\rho_u)\}] - E_{\mathbf{q}}[\log\{\mathbf{q}(\rho_u)\}] \\ &+ E_{\mathbf{q}}[\log\{\mathbf{p}(\tilde{\mathbf{u}}|\mathbf{b}_u, \sigma_u^2)\}] - E_{\mathbf{q}}[\log\{\mathbf{q}(\tilde{\mathbf{u}})\}] + E_{\mathbf{q}}[\log\{\mathbf{p}(\mathbf{b}_u)\}] - E_{\mathbf{q}}[\log\{\mathbf{q}(\mathbf{b}_u)\}] \\ &+ E_{\mathbf{q}}[\log\{\mathbf{p}(\sigma_u^2|a_u)\}] - E_{\mathbf{q}}[\log\{\mathbf{q}(\sigma_u^2)\}] + E_{\mathbf{q}}[\log\{\mathbf{p}(a_u)\}] - E_{\mathbf{q}}[\log\{\mathbf{q}(a_u)\}]. \end{aligned} \quad (18)$$

Then, in the Gaussian response case, the approximate marginal log-likelihood is

$$\begin{aligned}\log \underline{p}(\mathbf{y}; \mathbf{q}) &= \log \underline{p}(\mathbf{y}; \mathbf{q}, \text{BASE}) + E_{\mathbf{q}}[\log\{\mathbf{p}(\mathbf{y}|\beta_0, \gamma_\beta, \tilde{\boldsymbol{\beta}}, \gamma_u, \tilde{\mathbf{u}}, \sigma_\varepsilon^2)\}] \\ &\quad + E_{\mathbf{q}}[\log\{\mathbf{p}(\sigma_\varepsilon^2|a_\varepsilon)\}] - E_{\mathbf{q}}[\log\{\mathbf{q}(\sigma_\varepsilon^2)\}] \\ &\quad + E_{\mathbf{q}}[\log\{\mathbf{p}(a_\varepsilon)\}] - E_{\mathbf{q}}[\log\{\mathbf{q}(a_\varepsilon)\}]\end{aligned}$$

whilst, in the Bernoulli response case, it is

$$\begin{aligned}\log \underline{p}(\mathbf{y}; \mathbf{q}) &= \log \underline{p}(\mathbf{y}; \mathbf{q}, \text{BASE}) + E_{\mathbf{q}}[\log\{\mathbf{p}(\mathbf{y}|\mathbf{c})\}] \\ &\quad + E_{\mathbf{q}}[\log\{\mathbf{p}(\mathbf{c}|\beta_0, \gamma_\beta, \tilde{\boldsymbol{\beta}}, \gamma_u, \tilde{\mathbf{u}})\}] - E_{\mathbf{q}}[\log\{\mathbf{q}(\mathbf{c})\}].\end{aligned}$$

Explicit expressions for $\log \underline{p}(\mathbf{y}; \mathbf{q})$ in each case can be obtained by simplifying each of the \mathbf{q} -density moment expressions. For example, the first term of (18) is

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

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 (18) is

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

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

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

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

In the Gaussian response case, we have

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

and in the Bernoulli response case

$$\begin{aligned} \log p(\underline{\mathbf{y}}; \mathbf{q}) &= \log p(\underline{\mathbf{y}}; \mathbf{q}, \text{BASE}) + \sum_{i=1}^n \log \left\{ \Phi \left((2y_i - 1) \left(\mathbf{1}_n \mu_{q(\beta_0)} + \mathbf{X} \left(\mu_{q(\gamma_\beta)} \odot \mu_{q(\tilde{\beta})} \right) \right. \right. \right. \\ &\quad \left. \left. \left. + \sum_{j=1}^{d_\bullet} \mathbf{Z}_j \left(\mu_{q(\gamma_{u_j})} \odot \mu_{q(\tilde{\mathbf{u}}_j)} \right) \right) \right) \right\} + \text{const}_3 \end{aligned}$$

where const_2 and const_3 are constants that do not depend on any q -density parameters.