



Bayesian Model Selection in Additive Partial Linear Models Via Locally Adaptive Splines

Seonghyun Jeong, Taeyoung Park & David A. van Dyk

To cite this article: Seonghyun Jeong, Taeyoung Park & David A. van Dyk (2022) Bayesian Model Selection in Additive Partial Linear Models Via Locally Adaptive Splines, Journal of Computational and Graphical Statistics, 31:2, 324-336, DOI: [10.1080/10618600.2021.1999827](https://doi.org/10.1080/10618600.2021.1999827)

To link to this article: <https://doi.org/10.1080/10618600.2021.1999827>



© 2021 The Author(s). Published with license by Taylor & Francis Group, LLC.



Published online: 16 Dec 2021.



Submit your article to this journal [↗](#)



Article views: 979



View related articles [↗](#)



View Crossmark data [↗](#)

Bayesian Model Selection in Additive Partial Linear Models Via Locally Adaptive Splines

Seonghyun Jeong^{a,b}, Taeyoung Park^{a,b}, and David A. van Dyk^c

^aDepartment of Applied Statistics, Yonsei University, Seoul, Korea; ^bDepartment of Statistics and Data Science, Yonsei University, Seoul, Korea; ^cStatistics Section, Department of Mathematics, Imperial College London, London, UK

ABSTRACT

We provide a flexible framework for selecting among a class of additive partial linear models that allows both linear and nonlinear additive components. In practice, it is challenging to determine which additive components should be excluded from the model while simultaneously determining whether nonzero additive components should be represented as linear or nonlinear components in the final model. In this article, we propose a Bayesian model selection method that is facilitated by a carefully specified class of models, including the choice of a prior distribution and the nonparametric model used for the nonlinear additive components. We employ a series of latent variables that determine the effect of each variable among the three possibilities (no effect, linear effect, and nonlinear effect) and that simultaneously determine the knots of each spline for a suitable penalization of smooth functions. The use of a pseudo-prior distribution along with a collapsing scheme enables us to deploy well-behaved Markov chain Monte Carlo samplers, both for model selection and for fitting the preferred model. Our method and algorithm are deployed on a suite of numerical studies and are applied to a nutritional epidemiology study. The numerical results show that the proposed methodology outperforms previously available methods in terms of effective sample sizes of the Markov chain samplers and the overall misclassification rates.

ARTICLE HISTORY

Received August 2020
Revised September 2021

KEYWORDS

Bayesian adaptive regression;
Function estimation;
Knot-selection; Mixtures of
g-priors; Nonparametric
regression

1. Introduction

An additive partial linear model (APLM) generalizes both a standard linear model and a nonparametric additive model by combining both linear and nonlinear additive components in the framework of an additive regression model. Let $Y \in \mathbb{R}$ be a response variable, let $\mathcal{X} = (X_1, \dots, X_p)^\top \in \mathbb{R}^p$ be a p -dimensional continuous predictor vector, and let $\mathcal{Z} = (Z_1, \dots, Z_q)^\top \in \mathbb{R}^q$ be a q -dimensional predictor vector that may be either continuous or discrete. An APLM is written as

$$Y = \alpha + \sum_{j=1}^p f_j(X_j) + \sum_{k=1}^q \beta_k Z_k + \epsilon, \quad (1)$$

where α is an intercept term, $f_1, \dots, f_p$ are unknown functions corresponding to $X_1, \dots, X_p$, $\beta_1, \dots, \beta_q$ are coefficients corresponding to $Z_1, \dots, Z_q$, and ϵ is a Gaussian random error independent of the predictors $(\mathcal{X}, \mathcal{Z})$ with mean zero and finite variance σ^2 . The identifiability of $f_1, \dots, f_p$ in Equation (1) is assured by the restriction $E[f_j(X_j)] = 0$ for $j = 1, \dots, p$ (Xue 2009; Huang, Horowitz, and Wei 2010). The APLM in (1) is a useful compromise between the simplicity of a linear model and the flexibility of a nonparametric additive model, especially when the model involves categorical predictors. (A fully nonparametric additive structure for dummy variables associated with the categorical predictors reduces to a linear model.) Suppose

that some additive components, f_j , are zero, linear, or nonlinear, while some coefficients, β_k , are zero or not. In this article, we consider model selection for the APLM in (1), where nonzero additive components are distinguished from zero components while it is simultaneously determined whether each f_j can be further simplified as a strictly linear function via locally adaptive estimation.

Whereas there is a rich literature on both frequentist and Bayesian model selection methods that identify zero and nonlinear effects for the additive components in nonparametric additive models (Lin and Zhang 2006; Ravikumar et al. 2009; Xue 2009; Huang, Horowitz, and Wei 2010; Raskutti, Wainwright, and Yu 2012; Banerjee and Ghosal 2014; Curtis, Banerjee, and Ghosal 2014), the more general idea, which simultaneously identifies zero, linear, and nonlinear effects, has been sparsely considered. Frequentist approaches to this general problem typically involve many tuning parameters or multi-stage tuning procedures, which is challenging in practice (Lian, Liang, and Ruppert 2015; Wu and Stefanski 2015). On the other hand, Bayesian approaches provide a more natural solution that involves the comparison of posterior probabilities between competing models, but the problem of controlling the smoothness of unknown functions in the nonlinear additive components remains. Taking a Bayesian perspective, Rossell and Rubio (2019) addressed the above-mentioned general model selection problem for survival models

by introducing ternary latent variables. Although they provide a solid theoretical understanding of their method, smooth function estimation via an appropriate penalization is not considered in their work. Scheipl, Fahrmeir, and Kneib (2012) and Sabanés Bové, Held, and Kauermann (2015) dealt with the model selection problem together with smooth function estimation, but their approaches to the model selection are less intuitive, and their methods for function estimation do not catch the locally varying smoothness (see the next paragraph for more details). This study aims to address these issues.

It still remains to discuss how the nonparametric additive components can be estimated simultaneously with model selection. Several Bayesian strategies for estimating smooth functions have been developed, including Gaussian process regression, kernel regression, and spline regression. Here, we discuss two popular spline-based methods. First, Bayesian penalized splines can be used to estimate unknown functions with fixed knot locations and a prior distribution on the smoothing parameters (Ruppert, Wand, and Carroll 2003; Lang and Brezger 2004). For general model selection in the APLM that we consider, Scheipl, Fahrmeir, and Kneib (2012) and Sabanés Bové, Held, and Kauermann (2015) adopted the Bayesian penalized splines to estimate unknown functions in nonlinear additive components. However, this approach cannot capture the locally varying curvature of smooth functions without significant modifications (Ruppert, Wand, and Carroll 2003). Although several extensions to local adaptation have been proposed (Baladandayuthapani, Mallick, and Carroll 2005; Jullion and Lambert 2007; Scheipl and Kneib 2009), they are not readily available for model selection in an APLM because they require additional tuning parameters that complicate the calculations needed for model comparison. Second, Bayesian regression splines can be used to estimate unknown functions with basis (or knot) selection (Smith and Kohn 1996; Denison, Mallick, and Smith 1998; Kohn, Smith, and Chan 2001; DiMatteo, Genovese, and Kass 2001). By considering candidates of possible basis functions (depending on knot locations) and selecting optimal ones based on data, unknown functions can be estimated with locally adaptive smoothing and without overfitting, thereby overcoming the primary challenge of the Bayesian penalized splines.

An even more important advantage of the basis-selection approach is that model selection is accomplished as an automatic byproduct of the estimation of the nonlinear additive components in the APLM. Specifically, if no basis functions are selected in the regression spline, the corresponding nonlinear additive component is zero, whereas if only a linear basis function is selected, then it is linear. Despite this apparent advantage, the basis-selection approach has not been widely used for identifying function classes and selecting the set of nonlinear additive components to include in an APLM in this way. This is because doing so would require simultaneously selecting basis functions for each regression spline of a possibly large number of nonlinear additive components, leading to time-consuming computational challenges (Scheipl, Fahrmeir, and Kneib 2012). The primary contribution of this article is to show how a careful choice of prior distributions combined with an efficient Markov chain Monte Carlo (MCMC) posterior sam-

pling technique allows us to deploy the basis-selection approach for model selection.

We refer to our proposed method as Model Selection via Basis Selection (MSBS) in APLMs. MSBS allows the user to specify which variables can be considered only as linear predictors and which can be considered as either linear or nonlinear. It then automatically identifies predictors with zero or linear effects among the former and automatically identifies zero, linear, and nonlinear effects among the latter group. This, for example, allows dummy variables to be introduced without being confused as possible nonlinear predictors. We adopt a data-driven strategy to control both the number of basis functions and to select which basis functions to include from a large set. This allows us to flexibly and adaptively estimate the smooth predictor functions that additively enter the response function of the APLM.

Similar to Rossell and Rubio (2019), we introduce a set of latent indicator variables that identify which predictor variables have no effect, which have linear effects, and which have nonlinear effects. Another set of latent variables allows for locally adaptive smoothness of the predictor functions via their basis selection. The proposed model employs semi-conjugate prior distributions such as Zellner's g -prior distribution (Zellner 1986), which allows us to analytically compute the marginal posterior distributions of certain subsets of the parameters. By iteratively sampling from the conditional distributions of either the joint posterior distribution or its marginal posterior distribution, we construct an efficient partially collapsed Gibbs sampler (van Dyk and Park 2008). Convergence of the sampler is further improved via the "pseudo-prior" method developed by Carlin and Chib (1995). Finally, we introduce efficient MCMC samplers to fit the model chosen by the model selection procedure. Using a suite of simulations and applied examples, we illustrate how our proposed method outperforms existing Bayesian methods in terms of both computational efficiency and model selection performance.

The article is organized as follows. Section 2 introduces our formulation of the APLM, including the nonparametric models that we use for the predictor functions and the latent variables that we use to control model selection and smoothness of the predictor functions. For model selection and estimation, appropriate prior distributions are devised in Section 3 and computational methods for posterior inference in Section 4. Section 5 validates our overall proposed MSBS method through a suite of numerical examples and simulation studies. Section 6 illustrates its application to nutritional epidemiology data and Section 7 concludes with discussion. The R package for MSBS is currently available at the first author's github and described in Appendix A. Proofs of theoretical results are collected in Appendices B and C.

2. Formulation of Additive Partial Linear Models

Under certain smoothness assumptions, the nonlinear additive components, f_j , in (1) can be well approximated by regression splines defined as linear combinations of certain spline basis functions, although the appropriate choice of basis

functions depends on the smoothness assumptions. Let $\mathfrak{D} = \{(y_i, \mathbf{x}_i, \mathbf{z}_i)\}_{i=1}^n$ denote n i.i.d. realizations of $(Y, \mathcal{X}, \mathcal{Z})$ and let

$$\mathcal{H}_j^{(n)} = \left\{ h_j : h_j(\cdot) = \mathbf{b}_j(\cdot)^\top \boldsymbol{\phi}_j, \boldsymbol{\phi}_j \in \mathbb{R}^{M_j}, \sum_{i=1}^n h_j(x_{ij}) = 0 \right\} \quad (2)$$

denote a spline space spanned by M_j -dimensional basis functions $\mathbf{b}_j : \mathbb{R} \mapsto \mathbb{R}^{M_j}$ such that $\sum_{i=1}^n h_j(x_{ij}) = 0$, where x_{ij} is element j of $\mathbf{x}_i$. We consider a spline approximation $f_j^* \in \mathcal{H}_j^{(n)}$ to the true function f_j . The restriction $\sum_{i=1}^n h_j(x_{ij}) = 0$ is imposed on the spline space $\mathcal{H}_j^{(n)}$ in order to replace the identifiability restriction $E[f_j(X_j)] = 0$ in an empirical sense, that is, $\sum_{i=1}^n f_j^*(x_{ij}) = 0$. This can be achieved by using the centered basis functions defined by $\mathbf{b}_j^*(u) = \mathbf{b}_j(u) - n^{-1} \sum_{i=1}^n \mathbf{b}_j^*(x_{ij})$, where $\mathbf{b}_j^*(u)$ denotes a set of unrestricted basis functions.

Many types of basis functions, including B-splines, Bernstein polynomials, Fourier series expansion, and wavelets, are available to estimate a nonlinear function represented by a regression spline, depending on smoothness assumptions. For fast and numerically stable computation, we use the cubic radial basis functions (Kohn, Smith, and Chan 2001),

$$\mathbf{b}_j^*(u) = (u, |u - t_{j1}|^3, \dots, |u - t_{jL_j}|^3)^\top, \quad (3)$$

for a given $L_j > 0$. The abscissae, $\min_{1 \leq i \leq n} x_{ij} < t_{j1} < \dots < t_{jL_j} < \max_{1 \leq i \leq n} x_{ij}$, are candidates for knot locations, only some of which are chosen to achieve acceptable smoothness and curvature of f_j^* . This construction of basis functions implies that we have $M_j = L_j + 1$ for the spline space $\mathcal{H}_j^{(n)}$ in (2). Since only a few knots are used among the L_j candidates, the results are not sensitive to a value of L_j unless it is too small. It is sufficient that L_j be large enough to ensure that local features of the function can be detected. We recommend using a value between 20 and 30 for L_j unless the function is known to be extremely wiggly.

We introduce two sets of latent variables into the APLM, one set for controlling the smoothness of nonlinear predictor functions and the other for model selection. First, we let $\boldsymbol{\delta}_j = (\delta_{jm})_{m=0}^{L_j}$ denote a set of latent variables for smoothness control, where $\delta_{jm} = 1$ if basis function $(m+1)$ of the ordered list in (3) is used to approximate the nonlinear predictor function f_j and is 0 otherwise. Throughout the article, we use the cardinality notation $|\cdot|$ to denote the number of nonzero elements of a nonnegative vector. Thus, $|\boldsymbol{\delta}_j|$ denotes the number of basis functions used for f_j^* , i.e., $|\boldsymbol{\delta}_j| = \sum_{m=0}^{L_j} \delta_{jm}$. Let $\mathbf{b}_{\boldsymbol{\delta}_j}(\cdot)$ denote the $|\boldsymbol{\delta}_j| \times 1$ vector of basis functions selected by $\boldsymbol{\delta}_j$ and let $\boldsymbol{\phi}_{\boldsymbol{\delta}_j}$ denote the coefficients corresponding to the basis functions in $\mathbf{b}_{\boldsymbol{\delta}_j}(\cdot)$. Given the latent variables $\boldsymbol{\delta}_j$, f_j is approximated by the nonlinear function f_j^* ,

$$f_j^*(\cdot) = \mathbf{b}_{\boldsymbol{\delta}_j}(\cdot)^\top \boldsymbol{\phi}_{\boldsymbol{\delta}_j}. \quad (4)$$

Because there exists $\boldsymbol{\phi}_j$ such that $\mathbf{b}_{\boldsymbol{\delta}_j}(\cdot)^\top \boldsymbol{\phi}_{\boldsymbol{\delta}_j} = \mathbf{b}_j(\cdot)^\top \boldsymbol{\phi}_j$ and $\sum_{i=1}^n f_j^*(x_{ij}) = 0$ by the definition of the centered basis function $\mathbf{b}_j$, we know that $f_j^* \in \mathcal{H}_j^{(n)}$.

Turning to the set of latent variables introduced to facilitate model selection in the APLM, let $\boldsymbol{\gamma}^x = (\gamma_j^x)_{j=1}^p$ and $\boldsymbol{\gamma}^z = (\gamma_k^z)_{k=1}^q$ be defined by

$$\gamma_j^x = \begin{cases} 0 & \text{if } X_j \text{ has no effect on a response variable,} \\ 1 & \text{if } X_j \text{ has a linear effect on a response variable,} \\ 2 & \text{if } X_j \text{ has a nonlinear effect on a response variable,} \end{cases}$$

$$\gamma_k^z = \begin{cases} 0 & \text{if } Z_k \text{ has no effect on a response variable,} \\ 1 & \text{if } Z_k \text{ has a linear effect on a response variable,} \end{cases}$$

for $j = 1, \dots, p$ and $k = 1, \dots, q$. The construction of the ternary latent variable γ_j^x is similar to that of Rossell and Rubio (2019), but we allow more flexibility by also considering γ_k^z for the linear components. Introducing ternary latent variables makes it easy to assign prior information to each of no effect, linear effect, and nonlinear effect. Whereas the additive structure of $\mathcal{Z}$ is directly implied by $\boldsymbol{\gamma}^z$, determining that of $\mathcal{X}$ requires both $\boldsymbol{\gamma}^x$ and $\boldsymbol{\delta} = (\boldsymbol{\delta}_j)_{j=1}^p$. That is, all of the basis functions selected by $\boldsymbol{\delta}_j$ are included in the model if $\gamma_j^x = 2$, only the linear basis function is included if $\gamma_j^x = 1$, and none of the basis functions are included if $\gamma_j^x = 0$. Thus, we can define a set of new latent variables that indicates whether or not each basis function for X_j is included in a model. Specifically, let

$$\eta(\gamma_j^x, \boldsymbol{\delta}_j) = (\mathbb{1}(\gamma_j^x = 1) + \delta_{j0} \mathbb{1}(\gamma_j^x = 2), \delta_{j1} \mathbb{1}(\gamma_j^x = 2), \dots, \delta_{jL_j} \mathbb{1}(\gamma_j^x = 2)), \quad (5)$$

where $\mathbb{1}(\cdot)$ is the indicator function. That is, element m of $\eta(\gamma_j^x, \boldsymbol{\delta}_j)$ is equal to 1 if basis function m in $\mathbf{b}_j(\cdot)$ is included in a model and is 0 otherwise, so that $\eta(\gamma_j^x, \boldsymbol{\delta}_j)$ corresponds to a standard inclusion indicator vector for model selection. To ensure identifiability, we assume that if $\gamma_j^x = 2$, at least one $\delta_{jm} = 1$ with $m \geq 1$.

Using the latent variables $\boldsymbol{\gamma} = (\boldsymbol{\gamma}^x, \boldsymbol{\gamma}^z)$ and $\boldsymbol{\delta}$, the APLM with data $\mathfrak{D}$ can be written in matrix form as

$$\mathbf{y} = \alpha \mathbf{1}_n + \mathbf{B}_{\eta(\boldsymbol{\gamma}^x, \boldsymbol{\delta})} \boldsymbol{\phi}_{\eta(\boldsymbol{\gamma}^x, \boldsymbol{\delta})} + \mathbf{Z} \boldsymbol{\gamma}^z \boldsymbol{\beta}_{\boldsymbol{\gamma}^z} + \boldsymbol{\epsilon}, \quad (6)$$

where $\mathbf{y} = (y_1, \dots, y_n)^\top$ is a vector of response variables, $\mathbf{B}_{\eta(\boldsymbol{\gamma}^x, \boldsymbol{\delta})}$ is the submatrix of the matrix

$$\mathbf{B} = \begin{pmatrix} \mathbf{b}_1(x_{11})^\top & \cdots & \mathbf{b}_p(x_{1p})^\top \\ \vdots & \ddots & \vdots \\ \mathbf{b}_1(x_{n1})^\top & \cdots & \mathbf{b}_p(x_{np})^\top \end{pmatrix},$$

with columns selected according to $\boldsymbol{\eta}(\boldsymbol{\gamma}^x, \boldsymbol{\delta}) = (\eta(\gamma_j^x, \boldsymbol{\delta}_j))_{j=1}^p$, $\boldsymbol{\phi}_{\eta(\boldsymbol{\gamma}^x, \boldsymbol{\delta})}$ is the vector of coefficients corresponding to the columns of $\mathbf{B}_{\eta(\boldsymbol{\gamma}^x, \boldsymbol{\delta})}$, $\mathbf{Z} \boldsymbol{\gamma}^z$ is the submatrix of the matrix $\mathbf{Z} = (\mathbf{z}_1, \dots, \mathbf{z}_n)^\top$ with columns selected according to $\boldsymbol{\gamma}^z$, $\boldsymbol{\beta}_{\boldsymbol{\gamma}^z}$ is a vector of coefficients corresponding to the columns of $\mathbf{Z} \boldsymbol{\gamma}^z$, and $\boldsymbol{\epsilon} \sim N(\mathbf{0}_n, \sigma^2 \mathbf{I}_n)$.

From a practical perspective, how the available predictors are divided into $\mathcal{X}$ and $\mathcal{Z}$ is of great interest. We recommend including all continuous variable in $\mathcal{X}$ and including the dummy variables corresponding to all of the categorical variables as $\mathcal{Z}$. If there is a strong belief that some continuous variables do not have nonlinear effects, they can also be treated as a part of $\mathcal{Z}$.

3. Prior Specification

Our ultimate goal is to fit model (6) in a Bayesian framework. Doing this in a computationally efficient manner while maintaining desirable statistical properties requires carefully devised prior distributions. To facilitate formulation of these prior distributions, we let $\mathbf{W}_{[\mathbf{y}, \delta]} = (\mathbf{B}_{\eta(\mathbf{y}^x, \delta)}, \mathbf{Z}_{\mathbf{y}^z})$ and $\boldsymbol{\theta}_{[\mathbf{y}, \delta]} = (\boldsymbol{\phi}_{\eta(\mathbf{y}^x, \delta)}^\top, \boldsymbol{\beta}_{\mathbf{y}^z}^\top)^\top$ be a design matrix and the corresponding regression coefficients, respectively, and let $J_{[\mathbf{y}, \delta]} = |\eta(\mathbf{y}^x, \delta)| + |\mathbf{y}^z|$ (defined by $\sum_{j=1}^p \{\mathbb{1}(\gamma_j^x = 1) + |\delta_j| \mathbb{1}(\gamma_j^x = 2)\} + \sum_{k=1}^q \gamma_k^z$) be the dimension of $\boldsymbol{\theta}_{[\mathbf{y}, \delta]}$. Without loss of generality, we assume that each column of $\mathbf{Z}$ is centered, so that $\mathbf{W}_{[\mathbf{y}, \delta]}^\top \mathbf{1}_n = \mathbf{0}$; recall that the columns of $\mathbf{B}$ are also centered by definition.

Given the latent variables, $\mathbf{y}$ and δ , model (6) takes the form of a Gaussian linear regression. Thus, the standard conjugate prior distributions for a Gaussian linear regression are *semi-conjugate* priors for our APLM since the latent variables are unknown. For model selection, the posterior distribution of $(\mathbf{y}, \delta)$ is of primary interest. Using semi-conjugate prior distributions for the regression parameters allows us to compute the marginal likelihood function given $(\mathbf{y}, \delta)$ analytically – thus facilitating model selection.

3.1. Prior Distributions for α, σ^2 , and $\boldsymbol{\theta}_{[\mathbf{y}, \delta]}$

We consider the intercept, α , separately from the other regression parameters because it is common to all models. Using an argument based on invariance to scale and location transformations for orthogonal designs (Berger, Pericchi, and Varshavsky 1998), we suggest the improper prior distribution for (α, σ^2) given by

$$\pi(\alpha, \sigma^2) \propto 1/\sigma^2.$$

The use of this improper prior is justified since α and σ^2 appear in all the models we consider, whereas a proper prior must be used for $\boldsymbol{\theta}_{[\mathbf{y}, \delta]}$ (Berger, Pericchi, and Varshavsky 1998; Bayarri et al. 2012). Following convention, we use Zellner's g -prior distribution for $\boldsymbol{\theta}_{[\mathbf{y}, \delta]}$ (Zellner 1986; George and Foster 2000; Fernandez, Ley, and Steel 2001). We make the mild assumption that there is no explicit functional relationship among the columns of $(\mathbf{X}, \mathbf{Z})$ such that $\mathbf{W}_{[\mathbf{y}, \delta]}$ is full column rank for every $(\mathbf{y}, \delta)$ with $J_{[\mathbf{y}, \delta]} \leq n - 2$. Then for all $J_{[\mathbf{y}, \delta]} \leq n - 2$, the g -prior is given by the normal distribution,

$$\boldsymbol{\theta}_{[\mathbf{y}, \delta]} | \mathbf{W}, \mathbf{y}, \delta, g, \sigma^2 \sim N(0, g\sigma^2(\mathbf{W}_{[\mathbf{y}, \delta]}^\top \mathbf{W}_{[\mathbf{y}, \delta]})^{-1}),$$

where g is a dispersion factor and $\mathbf{W} = (\mathbf{B}, \mathbf{Z})$. Although the g -prior is also defined for $J_{[\mathbf{y}, \delta]} = n - 1$ (note that $\mathbf{W}_{[\mathbf{y}, \delta]}^\top \mathbf{W}_{[\mathbf{y}, \delta]}$ is invertible only when $J_{[\mathbf{y}, \delta]} \leq n - 1$ since the columns of $\mathbf{W}$ are centered), we exclude this case to simplify calculation of the marginal likelihood (see (8) and the last paragraph of this subsection below). Our specification of the joint prior distribution on $(\alpha, \boldsymbol{\theta}, \sigma^2)$ yields a convenient form for the marginal likelihood $\mathcal{L}(\mathbf{y} | \mathbf{W}, \mathbf{y}, \delta, g)$.

While the g -prior guarantees a closed-form expression for the marginal likelihood, a fixed choice of g would be difficult to determine and may lead to issues such as Bartlett's paradox or the information paradox (Liang et al. 2008). Hence, we

assign a prior to g , specifically a beta-prime prior distribution (Maruyama and George 2011) with density given by

$$\pi(g; a, b) = \frac{g^b(1+g)^{-a-b-2}}{B(a+1, b+1)}, \quad g > 0, \quad (7)$$

for $a > -1$ and $b > -1$, where $B(\cdot, \cdot)$ is the beta function. Under this prior, $1/(1+g)$ has a beta prior distribution, $Beta(a+1, b+1)$. As suggested by Maruyama and George (2011), we set $a = -3/4$ and $b = (n - J_{[\mathbf{y}, \delta]})/2 - 7/4$. This choice provides an easy-to-evaluate closed-form expression of the marginalized likelihood (see equation (3.10) and Remark 3.1 of Maruyama and George (2011)); for every $\mathbf{y}$ and δ such that $J_{[\mathbf{y}, \delta]} \leq n - 2$,

$$\begin{aligned} \mathcal{L}(\mathbf{y} | \mathbf{W}, \mathbf{y}, \delta) &= \int \mathcal{L}(\mathbf{y} | \mathbf{W}, \mathbf{y}, \delta, g) \pi(g; -3/4, (n - J_{[\mathbf{y}, \delta]})/2 - 7/4) dg \\ &= \frac{K(n, \mathbf{y}) B(J_{[\mathbf{y}, \delta]}/2 + 1/4, (n - J_{[\mathbf{y}, \delta]})/2 - 3/4)}{(1 - R_{[\mathbf{y}, \delta]}^2)^{(n - J_{[\mathbf{y}, \delta]})/2 - 3/4} B(1/4, (n - J_{[\mathbf{y}, \delta]})/2 - 3/4)}, \end{aligned} \quad (8)$$

where $R_{[\mathbf{y}, \delta]}^2$ is the coefficient of determination of the regression model determined by $(\mathbf{y}, \delta)$, and $K(n, \mathbf{y}) = \pi^{-(n-1)/2} n^{1/2} \Gamma((n-1)/2) \|\mathbf{y} - \bar{\mathbf{y}}\mathbf{1}_n\|^{-(n-1)}$ with $\bar{\mathbf{y}} = n^{-1} \sum_{i=1}^n y_i$.

The prior in (7) is not proper if $J_{[\mathbf{y}, \delta]} = n - 1$ with our choice of a and b . Although it is still possible to calculate a marginal likelihood by using a proper prior in this particular case (Maruyama and George 2011, eq. (3.9)), we instead exclude the case $J_{[\mathbf{y}, \delta]} = n - 1$ for simplicity. Hence, we impose the restriction $J_{[\mathbf{y}, \delta]} \leq n - 2$ in the construction of the priors for $\mathbf{y}$ and δ throughout the following subsections. We close this subsection by noting that the posterior distribution is proper despite the use of the improper prior $\pi(\alpha, \sigma^2) \propto 1/\sigma^2$. This can be easily checked by inspection of the expression for the marginal likelihood in (8).

3.2. Prior Distribution for $\mathbf{y}$

In this and the next subsections, we specify the prior distribution on the latent variables $\mathbf{y}$ and δ . Throughout we impose the restriction $J_{[\mathbf{y}, \delta]} \leq n - 2$ as mentioned above, that is, we assume $\pi(J_{[\mathbf{y}, \delta]} \leq n - 2) = 1$.

A popular prior distribution for latent indicator variables in Bayesian variable selection is the so-called hierarchical uniform prior, a mixture of a uniform distribution and a product of Bernoulli distributions (Kohn, Smith, and Chan 2001; Cripps, Carter, and Kohn 2005). The probability mass function for a binary vector $\mathbf{v} = (v_1, \dots, v_r)^\top$ following the hierarchical uniform distribution is given by $\pi(\mathbf{v}) = B(\sum_{i=1}^r v_i + 1, r - \sum_{i=1}^r v_i + 1)$. The distribution is considered to be vaguely informative because it implies $\sum_{k=1}^r v_k$ is uniformly distributed on $0, 1, \dots, r$. In variable selection, the distribution implies a discrete uniform distribution on the number of variables, provides effective multiplicity adjustment (Scott and Berger 2010), and outperforms other prior distributions such as a product of Bernoulli distributions with fixed probability (Moreno, Girón, and Casella 2015; Wang and Maruyama 2018).

In our setting, however, the elements of $\mathbf{y}^x$ are ternary (rather than binary). To accommodate this while preserving the desirable properties of the hierarchical uniform prior distribution, we propose a prior distribution for $\mathbf{y}$, given by

$$\pi(\mathbf{y}) = \frac{1}{2^{|\mathbf{y}|}} B(|\mathbf{y}| + 1, p + q - |\mathbf{y}| + 1),$$

$$\mathbf{y} : 0 \leq |\mathbf{y}| \leq p + q, \quad (9)$$

where $|\mathbf{y}| = \sum_{j=1}^p \mathbb{1}(\gamma_j^x > 0) + \sum_{k=1}^q \gamma_k^z$ and $|\mathbf{y}^x| = \sum_{j=1}^p \mathbb{1}(\gamma_j^x > 0)$ (following the definition of $|\cdot|$ given in Section 2). As with the hierarchical uniform prior distribution, further insight can be obtained by considering the prior distribution of $|\mathbf{y}|$ induced by (9). This is explored in the following proposition.

Proposition 1. The prior distribution in (9) implies $\pi(|\mathbf{y}| = k) = 1/(p + q + 1)$ for $k = 0, 1, \dots, p + q$.

Proposition 1 states that the prior distribution in (9) assigns equal probabilities to the number of predictors in the model in the same way that the hierarchical uniform prior distribution does. Thus, the prior in (9) is considered to be weakly informative in the same sense that the hierarchical uniform prior distribution is. The use of the prior distribution in (9) can also be supported by examining its marginal prior distributions for γ_j^x and γ_k^z .

Proposition 2. The marginal prior distributions of γ_j^x and γ_k^z induced by the prior distribution in (9) satisfy $\pi(\gamma_j^x = 0) = 1/2$, $\pi(\gamma_j^x = 1) = \pi(\gamma_j^x = 2) = 1/4$ for $j = 1, \dots, p$, and $\pi(\gamma_k^z = 0) = \pi(\gamma_k^z = 1) = 1/2$ for $k = 1, \dots, q$, respectively.

Proposition 2 states that the prior distribution in (9) implies that all of the variables are equally likely to be excluded or included in the model, and when they are included, the potentially nonlinear variables are equally likely to have a linear or nonlinear effect. Proofs for Propositions 1 and 2 are provided in Appendices B and C, respectively.

3.3. Prior Distribution for δ

It is sensible to assume independent prior distributions for different predictor functions, and we do so conditional on the order of the predictor function (no, linear, or nonlinear effect) as recorded by $\mathbf{y}^x$. Specifically, we assume $\pi(\delta|\mathbf{y}^x) = \prod_{j=1}^p \pi(\delta_j|\gamma_j^x)$. Partitioning δ_j into $(\delta_{j0}, \delta_{j\setminus 0})$ with $\delta_{j\setminus 0} = (\delta_{jm})_{m=1}^{L_j}$, we assume that δ_{j0} and $\delta_{j\setminus 0}$ are a priori independent, i.e., $\pi(\delta_j|\gamma_j^x) = \pi(\delta_{j0}|\gamma_j^x) \pi(\delta_{j\setminus 0}|\gamma_j^x)$.

Because the basis functions for predictor function j only come into the model when $\gamma_j^x = 2$, the latent variables δ_j do not affect the model if $\gamma_j^x \neq 2$. Still we specify both of the prior distributions $\pi(\delta_j|\gamma_j^x \neq 2)$ and $\pi(\delta_j|\gamma_j^x = 2)$ to enable efficient sampling of the posterior distribution; see Section 3.3.2. Of course, regardless of the observed data $\mathcal{D}$, we have $\pi(\delta_j|\gamma_j^x \neq 2) = \pi(\delta_j|\gamma_j^x \neq 2, \mathcal{D})$. This type of a prior distribution is often called a “pseudo-prior distribution” (Carlin and Chib 1995; Dellaportas, Forster, and Ntzoufras 2002). Unless otherwise specified, when we refer to the prior distribution of δ_j , we mean $\pi(\delta_j|\gamma_j^x = 2)$.

3.3.1. Prior Distributions for Nonlinear Predictor Functions, $\pi(\delta_j|\gamma_j^x = 2)$

To begin with, we specify that a priori probability of each linear basis function being included for nonlinear predictor functions is equal to 50%, that is, we simply set $\pi(\delta_{j0}|\gamma_j^x = 2) = 1/2$ for each j . The situation becomes more complicated for nonlinear basis functions. Because $\gamma_j^x = 2$ indicates that predictor function j is nonlinear, its decomposition must include at least one nonlinear basis function; otherwise the model is not identifiable in view of (5). This demands $\pi(|\delta_{j\setminus 0}| > 0 | \gamma_j^x = 2) = 1$, where $|\delta_{j\setminus 0}| = \sum_{m=1}^{L_j} \delta_{jm}$. We achieved this via a zero-truncated prior distribution formed by restricting a distribution for $\delta_{j\setminus 0}$ to the set $\{\delta_{j\setminus 0} : 1 \leq |\delta_{j\setminus 0}| \leq L_j\}$.

The theory of random series prior distributions suggests that a prior on $|\delta_{j\setminus 0}|$ should decay to zero to ensure optimal posterior contraction (Belitser and Serra 2014; Shen and Ghosal 2015). Although we use a different basis, with this in mind, we set a geometric prior distribution on $|\delta_{j\setminus 0}|$, renormalized so that $1 \leq |\delta_{j\setminus 0}| \leq L_j$, and assume that all values of $\delta_{j\setminus 0}$ resulting in the same value of $|\delta_{j\setminus 0}|$ are equally likely. That is, we assume

$$\pi(\delta_{j\setminus 0}|\gamma_j^x = 2) = \left(\frac{L_j}{|\delta_{j\setminus 0}|} \right)^{-1} \frac{(1 - p_j)^{|\delta_{j\setminus 0}| - 1} p_j}{1 - (1 - p_j)^{L_j}},$$

$$\delta_{j\setminus 0} : 1 \leq |\delta_{j\setminus 0}| \leq L_j, \quad (10)$$

where $0 < p_j < 1$. The prior distribution of $|\delta_{j\setminus 0}|$ induced by (10) approaches a zero-truncated geometric distribution as L_j increases. Based on our experience, we recommend using a fixed value of p_j between 0.3 and 0.6 so that the prior on $|\delta_{j\setminus 0}|$ is sufficiently concentrated on values less than 10 to 15. We use a fixed choice $p_j = 0.4$ in all of the numerical analyses in this article.

3.3.2. Specifying the Pseudo-Prior Distributions $\pi(\delta_j|\gamma_j^x \neq 2)$

Recall that from (5), the marginal likelihood is invariant to δ_j if $\gamma_j^x \neq 2$ and thus the pseudo-prior distribution, $\pi(\delta_j|\gamma_j^x \neq 2)$, is irrelevant to the marginal posterior distribution of $\mathbf{y}$, that is,

$$\pi(\mathbf{y}|\mathcal{D}) \propto \pi(\mathbf{y}) \sum_{\delta} \mathcal{L}(\mathbf{y}|\mathbf{W}, \mathbf{y}, \delta) \pi(\delta|\mathbf{y}^x)$$

$$= \pi(\mathbf{y}) \sum_{\delta_j: \gamma_j^x \neq 2} \left\{ \mathcal{L}(\mathbf{y}|\mathbf{W}, \mathbf{y}, \delta) \prod_{j: \gamma_j^x = 2} \pi(\delta_j|\gamma_j^x = 2) \right\}.$$

Although the choice of a pseudo-prior distribution does not affect the posterior distribution, it does affect the convergence of our posterior sampling algorithms. In particular, if $\pi(\delta_{j\setminus 0}|\gamma_j^x \neq 2)$ is similar to $\pi(\delta_{j\setminus 0}|\gamma_j^x = 2, \mathcal{D})$, then γ_j^x jumps more frequently in our MCMC sampler without being stuck at $\gamma_j^x = 0$ or $\gamma_j^x = 1$. For this reason, it is often recommended that a pilot chain be run to examine the posterior distribution of model-specific parameters (Dellaportas, Forster, and Ntzoufras 2002; van Dyk et al. 2009). Following this advice, we examine $\pi(\delta|\mathbf{y}_F, \mathcal{D})$ by running a pilot chain using (8) and (10), where $\mathbf{y}_F$ denotes a full model with $\gamma_j^x = 2$, $j = 1, \dots, p$, and $\gamma_k^z = 1$, $k = 1, \dots, q$. We then set $\pi(\delta_j|\gamma_j^x \neq 2) =$

$\hat{\pi}(\delta_{j0}|\mathbf{y}_F, \mathcal{D}) \hat{\pi}(\delta_{j\setminus 0}|\mathbf{y}_F, \mathcal{D})$, where $\hat{\pi}$ denotes an estimated probability distribution obtained from the pilot chain. The first factor in the approximation can easily be computed as the relative frequency of $\delta_{j0} = 0$ and $\delta_{j0} = 1$, for each j , with the constraint that neither probability may be zero or one. (We truncate $\hat{\pi}(\delta_{j0}|\mathbf{y}_F, \mathcal{D})$ so that it is between 0.01 and 0.99.)

Due to the dimensionality of $\delta_{j\setminus 0}$, the pilot chain may not visit all areas of its parameter space, making approximating the joint posterior distribution $\pi(\delta_{j\setminus 0}|\mathbf{y}_F, \mathcal{D})$ challenging. (We must consider the joint distribution because adjacent nonlinear spline basis functions are typically highly correlated.) To circumvent this issue, we smooth the posterior distribution of $\delta_{j\setminus 0}$ obtained by the pilot chain. Specifically, we use a normal density function to approximate the posterior distribution of $|\delta_{j\setminus 0}|$, and then put equal weight on each value of $\delta_{j\setminus 0}$ that gives the same value of $|\delta_{j\setminus 0}|$. (Note that we smooth the discrete distribution using a continuous density.) This ensures a nonzero density on the entire parameter space. The resulting smoothed posterior distribution is given by

$$\hat{\pi}(\delta_{j\setminus 0}|\mathbf{y}_F, \mathcal{D}) = \left(\frac{L_j}{|\delta_{j\setminus 0}|} \right)^{-1} \frac{\exp\{-(|\delta_{j\setminus 0}| - \xi_j)^2/2v_j^2\}}{\sum_{\ell=1}^{L_j} \exp\{-(\ell - \xi_j)^2/2v_j^2\}},$$

$$\delta_{j\setminus 0} : 1 \leq |\delta_{j\setminus 0}| \leq L_j, \quad (11)$$

where the hyperparameters ξ_j and v_j^2 are set by solving nonlinear equations to match the posterior mean and variance of $|\delta_{j\setminus 0}|$ obtained by the pilot chain to those of the distribution in Equation (11). (If the posterior variance of $|\delta_{j\setminus 0}|$ obtained by the pilot chain is zero, then we adjust it to be strictly positive in matching the hyperparameters.) Smoothing in this way allows us to run a shorter pilot chain that may not have fully converged. Because $|\delta_{j\setminus 0}|$ is discrete and restricted on $1 \leq |\delta_{j\setminus 0}| \leq L_j$, the hyperparameters ξ_j and v_j^2 are not the (approximated) posterior mean and variance of $|\delta_{j\setminus 0}|$, but values that reproduce them using (11). Based on our experience, the density in (11) gives good empirical results compared with other parametric smoothing alternatives.

The above construction of the pseudo-prior requires $p + q \leq n - 2$ due to the use of $\mathbf{y}_F$. Even if $p + q > n - 2$, a pilot chain can still be run with a minor modification. This can be done, for example, by setting $\gamma_k^z = 0$, $k = 1, \dots, q$, for $\mathbf{y}_F$ and by ignoring the restriction $|\delta_{j\setminus 0}| > 0$ (there is no identifiability issue due to fixed $\mathbf{y}_F$). Our experience shows that this approach results in slightly poorer convergence, but not significantly so. Our R package deploys this strategy when $p + q > n - 2$.

4. Monte Carlo Sampling Algorithms

In this section, we present the three Monte Carlo sampling algorithms needed for MSBS: (i) the algorithm used to obtain the pilot chain needed to set the pseudo-prior distribution described in Section 3.3.2, (ii) the algorithms used to sample the posterior distribution of $\mathbf{y}$ for model selection, and (iii) the algorithm used to sample the model parameters of the selected model from the posterior distribution.

4.1. Obtaining the Pilot Chain Needed for the Pseudo-Prior Distribution

To specify the pseudo-prior distribution, $\pi(\delta_j|\gamma_j^x \neq 2)$, we require a pilot chain that is a realization of a Markov chain with stationary distribution $p(\delta|\mathbf{y}_F, \mathcal{D})$; see Section 3.3.2. We obtain the Markov chain via a Metropolis within blocked Gibbs sampler. In our experience, a small number of iterations (e.g., 100 to 200) is enough for a pilot chain to find a reasonable approximation for the pseudo-prior distribution.

Sampler 1: Metropolis within blocked Gibbs sampler for pilot chain

Given starting values $\delta^{(0)}$, iterate the following.

For iteration $t = 1, 2, \dots$:

1. Partition the $(p + \sum_{j=1}^p L_j)$ components of δ into blocks. First, separate each of $\delta_{10}, \dots, \delta_{p0}$ into its own block of size one; in this way the sampler updates the linear and nonlinear basis functions separately. Randomly partition the remaining L_j elements of each $\delta_{j\setminus 0}$ into blocks of size 2, 3, or 4. The procedure partitions δ into a set of blocks, $[\delta]_1, \dots, [\delta]_B$, where B is the number of blocks constructed in iteration t . An example of possible blocks is

$$(\delta_{10}), (\delta_{11}, \delta_{12}, \delta_{13}, \delta_{14}), (\delta_{15}, \delta_{16}), \dots, (\delta_{1L_1-2}, \delta_{1L_1-1}, \delta_{1L_1}),$$

$$\vdots$$

$$(\delta_{j0}), (\delta_{j1}, \delta_{j2}), (\delta_{j3}, \delta_{j4}, \delta_{j5}), \dots, (\delta_{jL_j-3}, \delta_{jL_j-2}, \delta_{jL_j-1}, \delta_{jL_j}),$$

$$\vdots$$

$$(\delta_{p0}), (\delta_{p1}, \delta_{p2}, \delta_{p3}), (\delta_{p4}, \delta_{p5}, \delta_{p6}), \dots, (\delta_{pL_p-2}, \delta_{pL_p-1}, \delta_{pL_p}).$$

Assign the current value in $\delta^{(t-1)}$ to each of the block to arrive at $[\delta^{(t-1)}]_1, \dots, [\delta^{(t-1)}]_B$.

2. For block $b = 1, \dots, B$:

- (a) Sample a proposal $[\delta']_b$ of $[\delta]_b$ from the conditional prior distribution induced by $\pi(\delta|\mathbf{y}_F)$, i.e., sample

$$[\delta']_b \sim \pi([\delta]_b | [\delta^{(t)}]_{-b}, \mathbf{y}_F) \propto \pi([\delta]_b, [\delta^{(t)}]_{-b} | \mathbf{y}_F),$$

where $[\delta^{(t)}]_{-b} = ([\delta^{(t)}]_1, \dots, [\delta^{(t)}]_{b-1}, [\delta^{(t-1)}]_{b+1}, \dots, [\delta^{(t-1)}]_B)$.

- (b) Compute the acceptance probability

$$r = \min \left\{ 1, \frac{\mathcal{L}(\mathbf{y} | \mathbf{W}, \mathbf{y}_F, [\delta']_b, [\delta^{(t)}]_{-b})}{\mathcal{L}(\mathbf{y} | \mathbf{W}, \mathbf{y}_F, [\delta^{(t)}]_b, [\delta^{(t)}]_{-b})} \right\}.$$

- (c) Set $[\delta^{(t)}]_b = \begin{cases} [\delta']_b & \text{with probability } r, \\ [\delta^{(t-1)}]_b & \text{otherwise.} \end{cases}$

Since adjacent nonlinear basis functions are highly correlated, individual updating algorithms such as the standard Gibbs sampler may not work very well. Thus, Sampler 1 uses a blocking scheme to enable the swapping of adjacent basis functions and avoid the chain being trapped in a local mode. On the other hand, the linear basis term is not directly associated with other nonlinear basis functions. Since it is not clear how they are correlated, each δ_{j0} is isolated by forming its own block of size one. The partition changes every iteration for more flexibility.

In a nonparametric regression with basis selection, a small number of basis functions is often sufficient to estimate an unknown function, and hence the current value of $[\delta^{(t)}]_b$ is likely to be sparse. We thus use the conditional prior with a decaying tail property as a proposal distribution for $[\delta]_b$, so that a zero (or sparse) vector is likely to be proposed for each block. This allows us to circumvent the time-consuming evaluation of the acceptance probability r when the proposal is identical to the current value of $[\delta^{(t)}]_b$. Using the conditional prior as a proposal distribution improves the convergence efficiency of Sampler 1 with respect to *per floating point operation* rather than *per iteration* (Kohn, Smith, and Chan 2001). Kohn, Smith, and Chan (2001) empirically presented the efficiency of a sampling scheme based on a similar idea.

4.2. Bayesian Model Selection

For model selection, we obtain a Monte Carlo sample from the marginal posterior distribution $\pi(\boldsymbol{\gamma}|\mathcal{D})$. We do so by using a Metropolis within (blocked) Gibbs sampler to obtain a Markov chain with stationary distribution equal to $\pi(\boldsymbol{\gamma}, \delta|\mathcal{D})$ and discarding the sample of δ . (Owing to the pseudo-prior distribution the sample of δ does not have an immediate interpretation and is only introduced as a mechanism to obtain the sample of $\boldsymbol{\gamma}$.) For better readability, the values of γ_j^x and γ_k^z at iteration t of the following sampler are written as $[\gamma_j^x]^{(t)}$ and $[\gamma_k^z]^{(t)}$, respectively.

Sampler 2: Generating a Markov chain with stationary distribution $\pi(\boldsymbol{\gamma}, \delta|\mathcal{D})$

Given starting values $[\boldsymbol{\gamma}^x]^{(0)} = ([\gamma_j^x]^{(0)})_{j=1}^p$, $[\boldsymbol{\gamma}^z]^{(0)} = ([\gamma_k^z]^{(0)})_{k=1}^q$ and $\delta^{(0)} = (\delta_j^{(0)})_{j=1}^p$, iterate the following.

For iteration $t = 1, 2, \dots$:

1. Partition the index set $\{1, \dots, p\}$ into two disjoint sets, $\mathcal{J}_1^{(t)} = \{j : [\gamma_j^x]^{(t-1)} \neq 2\}$ and $\mathcal{J}_2^{(t)} = \{j : [\gamma_j^x]^{(t-1)} = 2\}$.
2. Draw $\{\delta_j^{(t)} : j \in \mathcal{J}_1^{(t)}\}$ from the constructed pseudo-prior,

$$\prod_{j \in \mathcal{J}_1^{(t)}} \pi(\delta_j | \gamma_j^x \neq 2) = \prod_{j \in \mathcal{J}_1^{(t)}} \hat{\pi}(\delta_{j0} | \boldsymbol{\gamma}_F, \mathcal{D}) \hat{\pi}(\delta_{j\setminus 0} | \boldsymbol{\gamma}_F, \mathcal{D}).$$

3. Draw $\{\delta_j^{(t)} : j \in \mathcal{J}_2^{(t)}\}$ using Sampler 1, but with the current value $\boldsymbol{\gamma}^{(t-1)}$ instead of fixed $\boldsymbol{\gamma}_F$. (Unlike in Sampler 1, we do not update all the components of δ here.)
4. For $j = 1, \dots, p$, draw $[\gamma_j^x]^{(t)}$ from the conditional posterior $\pi(\gamma_j^x | [\boldsymbol{\gamma}^x]^{(t)}, [\boldsymbol{\gamma}^z]^{(t-1)}, \delta^{(t)}, \mathcal{D})$, which is a discrete distribution with probabilities,

$$\begin{aligned} \pi(\gamma_j^x = \ell | \boldsymbol{\gamma}^x, \boldsymbol{\gamma}^z, \delta, \mathcal{D}) &\propto \mathcal{L}(\mathbf{y} | \mathbf{W}, \gamma_j^x = \ell, \boldsymbol{\gamma}^x, \boldsymbol{\gamma}^z, \delta) \\ &\times \pi(\gamma_j^x = \ell, \boldsymbol{\gamma}^x, \boldsymbol{\gamma}^z) \pi(\delta_j | \gamma_j^x = \ell), \quad \ell = 0, 1, 2, \end{aligned}$$

where $[\boldsymbol{\gamma}^x]^{(t)} = ([\gamma_1^x]^{(t)}, \dots, [\gamma_{j-1}^x]^{(t)}, [\gamma_{j+1}^x]^{(t-1)}, \dots, [\gamma_p^x]^{(t-1)})$, $\boldsymbol{\gamma}_{-j}^x = (\gamma_j^x)_{j \neq j}$, and $\pi(\delta_j | \gamma_j^x = \ell)$ is the prior in Section 3.3.1 if $\ell = 2$ and the pseudo-prior in Section 3.3.2 otherwise.

5. For $k = 1, \dots, q$, draw $[\gamma_k^z]^{(t)}$ from the conditional posterior $\pi(\gamma_k^z | [\boldsymbol{\gamma}^x]^{(t)}, [\boldsymbol{\gamma}^z]^{(t)}, \delta^{(t)}, \mathcal{D})$, which is a Bernoulli distribu-

tion with probabilities,

$$\begin{aligned} \pi(\gamma_k^z = \ell | \boldsymbol{\gamma}^x, \boldsymbol{\gamma}^z, \delta, \mathcal{D}) &\propto \mathcal{L}(\mathbf{y} | \mathbf{W}, \boldsymbol{\gamma}^x, \gamma_k^z = \ell, \boldsymbol{\gamma}^z, \delta) \\ &\times \pi(\boldsymbol{\gamma}^x, \gamma_k^z = \ell, \boldsymbol{\gamma}^z), \quad \ell = 0, 1, \end{aligned}$$

where $[\boldsymbol{\gamma}^z]^{(t)} = ([\gamma_1^z]^{(t)}, \dots, [\gamma_{k-1}^z]^{(t)}, [\gamma_{k+1}^z]^{(t-1)}, \dots, [\gamma_p^z]^{(t-1)})$ and $\boldsymbol{\gamma}_{-k}^z = (\gamma_k^z)_{k \neq k}$.

An estimate of $\boldsymbol{\gamma}$, denoted by $\hat{\boldsymbol{\gamma}}$, can be obtained with the accumulated MCMC samples. Because some values of $\boldsymbol{\gamma}$ are ternary, ordinary Bayesian model selection criteria such as the median probability model (Barbieri and Berger 2004) are not directly applicable. To find the most likely model, we thus generalize the median probability model such that $\hat{\boldsymbol{\gamma}}$ is determined by the marginal maximum a posteriori estimates: $\hat{\gamma}_j^x = \arg \max_{\ell \in \{0,1,2\}} \pi(\gamma_j^x = \ell | \mathcal{D})$, $j = 1, \dots, p$, and $\hat{\gamma}_k^z = \arg \max_{\ell \in \{0,1\}} \pi(\gamma_k^z = \ell | \mathcal{D})$, $k = 1, \dots, q$, where the marginal posteriors are estimated by the collected MCMC draws.

4.3. Estimation After Model Selection

Once we have found the most likely model, as specified by $\hat{\boldsymbol{\gamma}}$, we can fit this model by obtaining a Markov chain with stationary distribution $\pi(\alpha, \boldsymbol{\theta}_{[\hat{\boldsymbol{\gamma}}, \delta]}, \sigma^2, g, \delta | \hat{\boldsymbol{\gamma}}, \mathcal{D})$.

Here, we present an efficient sampler for APLMs using the data-driven knot-selection technique described in Section 1. The convergence characteristics of the sampling scheme can be improved by collapsing the parameters $(\alpha, \boldsymbol{\theta}_{[\hat{\boldsymbol{\gamma}}, \delta]}, \sigma^2, g)$ out of the model. In the algorithm below, we update δ_j only for j such that $\hat{\gamma}_j^x = 2$, because the latent variables δ_j affect the model likelihood only for such j . We still use the full vector notation δ for simplicity, but the actual target distribution $\pi(\alpha, \boldsymbol{\theta}_{[\hat{\boldsymbol{\gamma}}, \delta]}, \sigma^2, g, \delta | \hat{\boldsymbol{\gamma}}, \mathcal{D})$ is obtained by drawing δ_j for j such that $\hat{\gamma}_j^x \neq 2$ together using the pseudo-prior distribution.

Sampler 3: Generating a Markov chain with stationary distribution $\pi(\alpha, \boldsymbol{\theta}_{[\hat{\boldsymbol{\gamma}}, \delta]}, \sigma^2, g, \delta | \hat{\boldsymbol{\gamma}}, \mathcal{D})$

Given starting values $\delta^{(0)} = (\delta_j^{(0)})_{j=1}^p$, iterate the following.

For iteration $t = 1, 2, \dots$:

1. For $\hat{\mathcal{J}}_2 = \{j : \hat{\gamma}_j^x = 2\}$, draw $\{\delta_j^{(t)} : j \in \hat{\mathcal{J}}_2\}$ using Sampler 1, but with the estimate $\hat{\boldsymbol{\gamma}}$ instead of $\boldsymbol{\gamma}_F$.
2. Draw $g_*^{(t)}$ from $\pi(g_* | \hat{\boldsymbol{\gamma}}, \delta^{(t)}, \mathcal{D})$ which is a beta distribution,

$$g_* | \hat{\boldsymbol{\gamma}}, \delta, \mathcal{D} \sim \text{Beta} \left(\frac{J_{[\hat{\boldsymbol{\gamma}}, \delta]}}{2} + \frac{1}{4}, \frac{n - J_{[\hat{\boldsymbol{\gamma}}, \delta]}}{2} - \frac{3}{4} \right),$$

and let $g^{(t)} = (1/g_*^{(t)} - 1) / (1 - R_{[\hat{\boldsymbol{\gamma}}, \delta^{(t)}]}^2)$.

3. Draw $[\sigma^2]^{(t)}$ from $\pi(\sigma^2 | \hat{\boldsymbol{\gamma}}, \delta^{(t)}, g^{(t)}, \mathcal{D})$ which is an inverse gamma distribution,

$$\sigma^2 | \hat{\boldsymbol{\gamma}}, \delta, g, \mathcal{D} \sim \text{IG} \left(\frac{n-1}{2}, \frac{\|\mathbf{y} - \bar{\mathbf{y}}\mathbf{1}_n\|^2 [1 + g(1 - R_{[\hat{\boldsymbol{\gamma}}, \delta]}^2)]}{2(1+g)} \right).$$

4. Draw $(\alpha^{(t)}, \boldsymbol{\theta}_{[\hat{\boldsymbol{\gamma}}, \delta^{(t)}]}^{(t)})$ from $\pi(\alpha, \boldsymbol{\theta}_{[\hat{\boldsymbol{\gamma}}, \delta^{(t)}]} | \hat{\boldsymbol{\gamma}}, \delta^{(t)}, [\sigma^2]^{(t)}, g^{(t)}, \mathcal{D})$ which is a product of two independent normal distributions,

$$\alpha | \hat{\boldsymbol{\gamma}}, \delta, \sigma^2, g, \mathcal{D} \sim N(\bar{\mathbf{y}}, n^{-1}\sigma^2),$$

$$\boldsymbol{\theta}_{[\hat{\boldsymbol{\gamma}}, \delta]} | \hat{\boldsymbol{\gamma}}, \delta, \sigma^2, g, \mathcal{D} \sim N \left(\frac{g}{1+g} \bar{\boldsymbol{\theta}}_{[\hat{\boldsymbol{\gamma}}, \delta]}, \frac{g\sigma^2}{1+g} (\mathbf{W}_{[\hat{\boldsymbol{\gamma}}, \delta]}^\top \mathbf{W}_{[\hat{\boldsymbol{\gamma}}, \delta]})^{-1} \right),$$

where $\bar{\boldsymbol{\theta}}_{[\hat{\boldsymbol{\gamma}}, \delta]} = (\mathbf{W}_{[\hat{\boldsymbol{\gamma}}, \delta]}^\top \mathbf{W}_{[\hat{\boldsymbol{\gamma}}, \delta]})^{-1} \mathbf{W}_{[\hat{\boldsymbol{\gamma}}, \delta]}^\top (\mathbf{y} - \bar{\mathbf{y}}\mathbf{1}_n)$.

The conditional distributions sampled by Sampler 3 are functionally incompatible (as defined in Section 2.2 of Hobert and Casella 1998). This means that there is no joint distribution for which the sampled conditionals are the full conditional distribution and changing the order of the draws may alter the stationary distribution of the resulting Markov chain (van Dyk and Park 2008; Park and Lee 2021). Formally, the sampled distributions are conditional distributions of a nested sequence of marginal distributions of the target posterior, that is, a sequence of partially collapsed distributions (van Dyk and Park 2008). In this case, it is straightforward to verify that the sampler's stationary distribution is indeed $\pi(\alpha, \theta_{[\hat{\gamma}, \delta]}, \sigma^2, g, \delta | \hat{\gamma}, \mathcal{D})$. Step 1 (combined with sampling all other entries of δ using the pseudo-prior) samples a set of complete conditional distributions of the marginal posterior distribution $\pi(\delta | \hat{\gamma}, \mathcal{D})$ and operates completely independently of Steps 2–4. Thus, its (Metropolis within blocked Gibbs) sampler maintains the marginal target distribution, $\pi(\delta | \hat{\gamma}, \mathcal{D})$, as its stationary distribution (e.g., van Dyk and Jiao 2015). The subsequent steps directly sample the conditional distributions $\pi(g | \delta, \hat{\gamma}, \mathcal{D})$, $\pi(\sigma^2 | g, \delta, \hat{\gamma}, \mathcal{D})$, and $\pi(\alpha, \theta_{[\hat{\gamma}, \delta]} | \sigma^2, g, \delta, \hat{\gamma}, \mathcal{D})$, each of which maintains the target posterior distribution. In general, care must be taken when implementing algorithms such as Sampler 3 that involve partially collapsed conditional distributions to be sure that the resulting stationary distribution coincides with the target (posterior) distribution; see Park and van Dyk (2009) for details.

Sampler 3 enables us to fit the APLM with a given set of covariates while averaging over the knot indicator variables, δ , which change over the Markov chain iterations (i.e., Bayesian model averaging, Raftery, Madigan, and Hoeting 1997). More specifically, model-averaged estimates of smooth functions are obtained by recovering f_j^* using (4) with posterior samples of δ_j and ϕ_{δ_j} in each iteration; that is, with T iterations of a Markov chain, we obtain

$$\hat{f}_j^*(\cdot) = \frac{1}{T} \sum_{t=1}^T \mathbf{b}_{\delta_j^{(t)}}(\cdot)^\top \boldsymbol{\phi}_{\delta_j^{(t)}}^{(t)}$$

for j such that $\hat{\gamma}_j^x = 2$. This estimate reduces the bias caused by fixing a single set of basis functions. We could also average over γ rather than conditioning on the estimate $\hat{\gamma}$. This approach is known to work well for prediction (Raftery, Madigan, and Hoeting 1997), but it is more difficult to interpret the results since this approach does not produce a set of included covariates. We do not consider this form of model averaged prediction in this article.

As a final remark, we recall that the procedure requires $\mathbf{Z}$ to be centered for identifiability. This can be done by the translation $\mathbf{Z} = \mathbf{Z}^* - n^{-1} \mathbf{1}_n \mathbf{1}_n^\top \mathbf{Z}^*$ with the original (uncentered) matrix $\mathbf{Z}^*$. Since this operation changes the intercept in the model, the intercept of the original location, denoted by α^* , can be restored by the translation $\alpha^* = \alpha - \mathbf{1}_n^\top \mathbf{Z}_{\hat{\gamma}^z}^* \boldsymbol{\beta}_{\hat{\gamma}^z}$ using the uncentered $\mathbf{Z}^*$ and samples of α and $\boldsymbol{\beta}_{\hat{\gamma}^z}$ in each MCMC iteration.

5. Numerical Illustration

We conduct a simulation study to evaluate the performance of MSBS. The test datasets are generated from the model

$$Y = \alpha + \sum_{j=1}^p f_j(X_j) + \sum_{k=1}^q \beta_k Z_k + \epsilon, \quad \epsilon \sim N(0, \sigma^2),$$

where parameters are set to $\alpha = 1$, $\beta_1 = 0.25$, $\beta_2 = 0.5$, $\beta_3 = 0.75$, $\beta_4 = \dots = \beta_q = 0$, and $\sigma^2 = 1$, and smooth functions are chosen as

$$\begin{aligned} f_1(x) &= 4x^2, \\ f_2(x) &= \sin(2\pi x), \\ f_3(x) &= 3 \exp\{-200(x - 0.2)^2\} + 0.5 \exp\{-50(x - 0.6)^2\}, \\ f_4(x) &= x, \\ f_5(x) &= 1.5x, \\ f_6(x) &= 2x, \\ f_j(x) &= 0, \quad j = 7, \dots, p, \end{aligned}$$

that is, three nonlinear functions and three linear functions are actually used for the smooth functions. (Note that in the data-generating process, we need not consider the identifiability restriction $E[f_j(X_j)] = 0$ since the functions are automatically adjusted by the intercept term.) Thus, we have a p -dimensional vector $\gamma^x = (2, 2, 2, 1, 1, 1, 0, \dots, 0)^\top$ and a q -dimensional vector $\gamma^z = (1, 1, 1, 0, \dots, 0)^\top$. The predictors $\{X_j, j = 1, \dots, p\}$ are generated by the Gaussian copula such that X_j are marginally distributed uniformly on $(0, 1)$ with pairwise correlation coefficient denoted by ρ . The predictors $\{Z_k, k = 1, \dots, q\}$ are generated from a q -dimensional multivariate normal distribution with zero mean, identity covariance matrix, and pairwise correlation coefficient denoted by ρ . We consider two different values for the pairwise correlation coefficient ($\rho = 0.5$ and $\rho = 0.95$), and three different cases for the dimension of covariates ($p = q = 10$, $p = q = 25$, and $p = q = 50$); note that a larger model with bigger p and q has more irrelevant variables because the number of nonzero coefficients are the same in all cases. We also consider three different values for the sample size ($n = 100$, $n = 200$, and $n = 500$), so that a total of 18 simulation settings are considered. For each of the simulation settings, 100 independent datasets are generated to account for a sampling error.

The effects of predictor variables are estimated using the method proposed in Section 4.2, and the misclassification rates for each possible value of the latent variables, γ^x and γ^z , are computed as

$$\begin{aligned} \text{MR}_\ell^x &= \frac{\sum_{j=1}^p \mathbb{1}(\hat{\gamma}_j^x \neq \ell, \gamma_j^x = \ell)}{\sum_{j=1}^p \mathbb{1}(\gamma_j^x = \ell)}, \quad \ell = 0, 1, 2, \\ \text{MR}_\ell^z &= \frac{\sum_{k=1}^q \mathbb{1}(\hat{\gamma}_k^z \neq \ell, \gamma_k^z = \ell)}{\sum_{k=1}^q \mathbb{1}(\gamma_k^z = \ell)}, \quad \ell = 0, 1. \end{aligned}$$

The overall misclassification rates for linear and nonlinear components are defined as

$$\text{MR}^x = \frac{\sum_{j=1}^p \mathbb{1}(\hat{\gamma}_j^x \neq \gamma_j^x)}{p}, \quad \text{MR}^z = \frac{\sum_{k=1}^q \mathbb{1}(\hat{\gamma}_k^z \neq \gamma_k^z)}{q},$$

and the total misclassification rate, MR_T , is defined as a weighted average of MR^x and MR^z . We compare our MSBS with the Bayesian methods of Sabanés Bové, Held, and Kauermann

Table 1. Simulation results with $p = q = 10$. The performance measures are averaged over 100 replications.

ρ	n	Method	Nonlinear components				Linear components			
			MR_2^x	MR_1^x	MR_0^x	MR^x	MR_1^z	MR_0^z	MR^z	MR_T
0.5	100	MSBS	0.50	0.47	0.03	0.30	0.27	0.05	0.12	0.21
		Sabanés Bové	0.81	0.38	0.09	0.39	0.32	0.07	0.14	0.27
		Scheipl	0.66	0.56	0.02	0.37	0.45	0.00	0.14	0.26
	200	MSBS	0.05	0.19	0.01	0.08	0.11	0.05	0.07	0.07
		Sabanés Bové	0.14	0.15	0.03	0.10	0.14	0.04	0.07	0.08
		Scheipl	0.19	0.25	0.01	0.13	0.28	0.00	0.08	0.11
	500	MSBS	0.00	0.02	0.02	0.01	0.01	0.02	0.02	0.02
		Sabanés Bové	0.00	0.01	0.02	0.02	0.02	0.02	0.02	0.02
		Scheipl	0.05	0.05	0.00	0.03	0.15	0.00	0.04	0.04
0.95	100	MSBS	0.78	0.98	0.01	0.53	0.83	0.03	0.27	0.40
		Sabanés Bové	0.79	0.97	0.01	0.53	0.81	0.04	0.27	0.40
		Scheipl	0.97	1.00	0.00	0.59	0.96	0.01	0.29	0.44
	200	MSBS	0.47	0.85	0.02	0.40	0.68	0.02	0.22	0.31
		Sabanés Bové	0.52	0.87	0.01	0.42	0.71	0.03	0.23	0.33
		Scheipl	0.65	0.96	0.01	0.48	0.81	0.00	0.25	0.37
	500	MSBS	0.22	0.56	0.01	0.24	0.39	0.02	0.14	0.18
		Sabanés Bové	0.24	0.59	0.00	0.25	0.40	0.02	0.13	0.19
		Scheipl	0.39	0.74	0.00	0.34	0.52	0.00	0.16	0.25

(2015) and Scheipl, Fahrmeir, and Kneib (2012), which we refer to as “Sabanés Bové” and “Scheipl.” These methods are implemented by using the R packages, *hypermgsplines* and *spikeSlabGAM*, respectively. (The *hypermgsplines* package is not currently maintained in the repositories; we installed it in an old version of R using the source files.) For all methods, 20 knots (knot-candidates for MSBS) are used to approximate smooth functions. Since all methods suitably penalize smooth functions in the procedure of model selection, the precise number of knots (or knot candidates) is not very important and it suffices to use a value that is not too small. Interestingly, we found that Scheipl is not invariant to the scale of a response variable, so observations are standardized before the analysis.

The results are presented in Tables 1–3, which show the performance measures averaged over the 100 replications. We observe that Sabanés Bové and Scheipl do not work for some setups with $p = q = 50$, and therefore those settings were excluded from Table 3. It is clear that MSBS always outperforms the other methods in terms of MR_2^x , the misclassification rate showing how well a method finds a true nonlinear component. This implies that MSBS can better identify a complex structure in a relationship between a predictor and a response. Tables 1 and 2 show that Sabanés Bové and Scheipl tend to be good at identifying a zero or linear component, especially with small sample size. This means that these methods generally prefer a structure that is simpler than the true structure for the predictor–response relationship. When it comes to the overall performance measured by MR^x , MR^z , and MR_T , MSBS clearly outperforms or is at least comparable to the other two methods. If our interest is primarily in understanding the true unknown complex structure for the effect of a predictor on a response, then MSBS performs very well.

Besides, the misclassification rates, MSBS is compared with two other Bayesian methods in terms of computational efficiency, via the multivariate effective sample size (mESS) of the Markov chain (Vats, Flegal, and Jones 2019) divided by the runtime required to obtain 10,000 iterations after 1000 burn-in

Table 2. Simulation results with $p = q = 25$. The performance measures are averaged over 100 replications.

ρ	n	Method	Nonlinear components				Linear components			
			MR_2^x	MR_1^x	MR_0^x	MR^x	MR_1^z	MR_0^z	MR^z	MR_T
0.5	100	MSBS	0.55	0.61	0.01	0.15	0.43	0.01	0.06	0.11
		Sabanés Bové	0.91	0.44	0.15	0.28	0.41	0.12	0.16	0.22
		Scheipl	0.78	0.70	0.00	0.18	0.61	0.00	0.07	0.13
	200	MSBS	0.04	0.24	0.00	0.04	0.21	0.01	0.04	0.04
		Sabanés Bové	0.32	0.23	0.01	0.08	0.25	0.01	0.04	0.06
		Scheipl	0.29	0.32	0.00	0.08	0.35	0.00	0.04	0.06
	500	MSBS	0.00	0.03	0.01	0.01	0.06	0.00	0.01	0.01
		Sabanés Bové	0.00	0.02	0.01	0.01	0.06	0.00	0.01	0.01
		Scheipl	0.19	0.15	0.00	0.04	0.20	0.00	0.02	0.03
0.95	100	MSBS	0.85	0.98	0.01	0.22	0.93	0.01	0.12	0.17
		Sabanés Bové	0.88	0.98	0.01	0.23	0.91	0.02	0.13	0.18
		Scheipl	0.99	1.00	0.00	0.24	1.00	0.00	0.12	0.18
	200	MSBS	0.55	0.95	0.00	0.18	0.80	0.01	0.11	0.14
		Sabanés Bové	0.62	0.92	0.00	0.19	0.77	0.02	0.11	0.15
		Scheipl	0.82	1.00	0.00	0.22	0.96	0.00	0.12	0.17
	500	MSBS	0.22	0.65	0.01	0.11	0.53	0.00	0.07	0.09
		Sabanés Bové	0.23	0.65	0.00	0.11	0.52	0.01	0.07	0.09
		Scheipl	0.46	0.88	0.00	0.16	0.73	0.00	0.09	0.12

Table 3. Simulation results with $p = q = 50$. The performance measures are averaged over 100 replications.

ρ	n	Method	Nonlinear components				Linear components				
			MR_2^X	MR_1^X	MR_0^X	MR^X	MR_1^Z	MR_0^Z	MR^Z	MR_T	
0.5	100	MSBS	0.65	0.72	0.00	0.09	0.52	0.01	0.04	0.06	
		Sabanés Bové	—	—	—	—	—	—	—	—	
		Scheipl	0.84	0.80	0.00	0.10	0.62	0.00	0.04	0.07	
	200	MSBS	0.06	0.28	0.00	0.02	0.24	0.00	0.02	0.02	
		Sabanés Bové	0.90	0.21	0.24	0.28	0.25	0.19	0.19	0.24	
		Scheipl	0.33	0.43	0.00	0.05	0.37	0.00	0.02	0.03	
	500	MSBS	0.00	0.05	0.00	0.00	0.07	0.00	0.01	0.00	
		Sabanés Bové	0.03	0.05	0.00	0.01	0.08	0.00	0.01	0.01	
		Scheipl	0.23	0.20	0.00	0.03	0.25	0.00	0.01	0.02	
	0.95	100	MSBS	0.87	1.00	0.00	0.11	0.92	0.01	0.06	0.09
			Sabanés Bové	—	—	—	—	—	—	—	—
			Scheipl	—	—	—	—	—	—	—	—
200		MSBS	0.60	0.95	0.00	0.10	0.85	0.00	0.06	0.08	
		Sabanés Bové	0.82	0.95	0.01	0.11	0.83	0.01	0.06	0.09	
		Scheipl	0.97	1.00	0.00	0.12	0.99	0.00	0.06	0.09	
500		MSBS	0.24	0.74	0.00	0.06	0.59	0.00	0.04	0.05	
		Sabanés Bové	0.46	0.82	0.00	0.08	0.56	0.01	0.04	0.06	
		Scheipl	0.54	0.96	0.00	0.09	0.94	0.00	0.06	0.07	

iterations. This measure represents the number of independent draws of a Markov chain Monte Carlo algorithm per unit time. The R package *mcmcse* gives an estimate of mESS. To compute mESS, the posterior variance of each γ_j^x and of each γ_k^z needs to be sufficiently far from zero. To achieve this, instead of using $\sigma^2 = 1$, here we use $\sigma^2 = n/50$ and $n/100$ for the $\rho = 0.5$ and 0.95 cases, respectively. The case with $p = q = 50$ is not considered for this simulation study, because too many MCMC iterations are required for the *mcmcse* package to work properly in this case. The performance measures are averaged over only 15 replications. The comparison of computational efficiency is summarized in Figure 1. In the small sample size case ($n = 100$), Sabanés Bové tends to be more efficient than the other two methods, including MSBS. As the sample size becomes larger ($n = 200$ and $n = 500$), however, the computational efficiency of MSBS significantly improves and outperforms the other two methods.

The aim of the article is to propose a new model selection strategy for additive partial linear models. Thus, we do not




Figure 1. Boxplots for multivariate effective sample size divided by runtime. The measures are averaged over 15 replications and show the computational advantage of MSBS, especially for large sample sizes.

detail the performance of the function estimates proposed in Section 4.3. The advantages of the knot-selection technique in nonparametric regression, however, have been illustrated elsewhere (Smith and Kohn 1996; Kohn, Smith, and Chan 2001; Jeong and Park 2016; Jeong, Park, and Park 2017; Park and Jeong 2018).

6. Application to Nutritional Epidemiology Study

In this section, we consider the nutritional epidemiology study discussed by Fletcher and Fairfield (2002) and Liu, Wang, and Liang (2011). There is a close relationship between beta-carotene and certain types of cancer, such as lung, colon, breast, and prostate. It is known that the antioxidant properties of beta-carotene help eliminate cancer-causing free radicals and thus can effectively reduce the risk of certain cancers. In addition, a sufficient beta-carotene supply can strengthen the body's autoimmune system to fight degenerative diseases such as cancer. Thus, it is of interest for clinicians and nutritionists to investigate the relationship between beta-carotene concentrations and other regulatory factors like age, gender, dietary intake, smoking status, and alcohol consumption. Some studies use simple linear models to explore the relationship, and the results are diverse and inconsistent (Nierenberg et al. 1989; Faure et al. 2006). Therefore, there is a need for a more advanced statistical analysis to closely investigate the relationship between beta-carotene and cancer risk factors.

MSBS is thus applied to identify significant linear/nonlinear effects of risk factors on the logarithm of beta-carotene concentration. The response variable is the plasma concentration of beta-carotene from 315 patients and the following risk factors are available: AGE (years), SEX (1=male, 2=female), SMOK-STAT (smoking status: 1=never, 2=former, 3=current Smoker), BMI (weight/height²), VITUSE (vitamin use: 1=yes, fairly often, 2=yes, not often, 3=no), CALORIES (number of calories consumed per day), FAT (grams of fat consumed per day), FIBER (grams of fiber consumed per day), ALCOHOL (number of alcoholic drinks consumed per week), CHOL (mg of cholesterol consumed per day), BETADIET (mcg of dietary beta-carotene consumed per day), and RETDIET (mcg of dietary retinol consumed per day).

Liu, Wang, and Liang (2011) used an additive partial linear model to examine the relationship between the logarithm of beta-carotene concentration and some personal characteristics by selecting significant linear predictors with their variable selection method after determining predictors having nonlinear effects via a preliminary analysis. Banerjee and Ghosal (2014) used a similar model with more predictors and chose significant nonlinear predictors as well as linear predictors. The previous studies treated both AGE and CHOL as predictors with nonlinear effects. As can be seen of (Liu, Wang, and Liang 2011, fig. 1), however, the significance of nonlinear effect of CHOL is uncertain. We thus apply MSBS to examine the true effects of the predictors on the logarithm of beta-carotene concentration.

We include the categorical variables SEX, SMOKSTAT, and VITUSE into the model as candidates of linear predictors, after centering in advance. The other covariates are continuous variables and are candidates for nonlinear predictors. Table 4 shows the posterior probabilities of the latent indicator variables for model selection and the resulting estimate $\hat{\gamma}$. Our results are that MSBS selects SEX, SMOKSTAT, BMI, VITUSE, and FIBER as predictors with linear effects, and selects AGE as a predictor with a nonlinear effect. Interestingly, CHOL is chosen to have no significant effect on the logarithm of beta-carotene concentration in our analysis, which is the main difference between our analysis and former studies analyzing the same data. The posterior summary statistics for predictors selected to have linear effects are provided in Table 5, and the estimated smooth function for AGE is given in Figure 2.

Table 4. Results of model selection in the beta-carotene study. The columns show the posterior probabilities of the latent indicator variables for model selection and the estimates of the latent indicator variables under the best model.

Predictor	Posterior probability			$\hat{\gamma}$
	Zero	Linear	Nonlinear	
AGE	0.034	0.017	0.949	2
SEX	0.316	0.684	–	1
SMOKSTAT	0.154	0.846	–	1
BMI	0.000	0.899	0.101	1
VITUSE	0.053	0.947	–	1
CALORIES	0.765	0.199	0.036	0
FAT	0.757	0.207	0.038	0
FIBER	0.104	0.732	0.164	1
ALCOHOL	0.836	0.119	0.045	0
CHOL	0.622	0.252	0.126	0
BETADIET	0.600	0.358	0.042	0
RETDIET	0.844	0.120	0.036	0

Table 5. Posterior summary statistics for predictors selected to have linear effects.

Parameter	Mean	Median	95% credible interval
Intercept	5.462	5.462	(4.433, 6.494)
SEX	0.290	0.289	(−0.054, 0.635)
SMOKSTAT	−0.141	−0.141	(−0.303, 0.019)
BMI	−0.033	−0.033	(−0.052, −0.014)
VITUSE	−0.129	−0.129	(−0.262, 0.003)
FIBER	0.024	0.024	(0.002, 0.045)
σ^2	0.411	0.410	(0.351, 0.482)




Figure 2. Posterior estimates of a smooth function for AGE in the beta-carotene study. The solid line represents the point-wise posterior median of the function and the gray region represents the point-wise 95% credible interval.

7. Discussion

This article proposes a new Bayesian model selection strategy and an efficient estimation method for additive partial linear models. Besides latent variables for model selection, we introduce additional latent variables that select basis functions to reduce a bias by data-adaptively controlling the smoothness of a nonlinear function. New prior distributions are carefully devised to overcome difficulties arising in the existing knot-selection methods. The resulting Bayesian model selection method, which we call MSBS, outperforms the existing methods in terms of overall model selection accuracy and computational efficiency.

The main focus of this article is on additive partial linear models with Gaussian errors, which can possibly provide a guide toward future research on the models with nonnormal errors. Thus, MSBS developed in this article can be generalized to analyze nonnormal response data within a generalized linear model framework. Such a generalization can be based on recent studies about mixtures of g -priors for generalized linear models (Sabanés Bové and Held 2011; Li and Clyde 2018).

It is also possible to naturally extend MSBS for a nonparametric additive model with multivariate additive components. The extension is straightforward if one uses a tensor product of splines (de Boor 1978) or radial basis functions in multiple dimensions (Kohn, Smith, and Chan 2001). As another extension, one may consider an adaptive sampler for our model selection problem instead of the given Metropolis–Hastings algorithm. This may be achieved by using the currently available sampling schemes (Nott and Kohn 2005; Ji and Schmidler 2013).

Acknowledgments

We thank to the associate editor and the two anonymous referees for their constructive advice.

Funding

S. Jeong's research was supported by the Yonsei University Research Fund of 2021-22-0032. T. Park's research was supported by Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education (2017R1D1A1B03033536) and by the National Research Foundation of Korea (NRF) grant funded by the Korea government(MSIT) (2020R1A2C1A01005949).

Appendices

Appendix A. R Package MSBS

Here we show how to use the R package for MSBS. Using the devtools package available at CRAN, our R package is installed and loaded by running

```
devtools::install_github("s-jeong/MSBS")
library(MSBS)
```

and then one can reproduce the results in Section 6 by running the examples in the help pages of the R functions `bmsaplm` and `estaplm`.

Appendix B. Proof of Proposition 1

By writing $|\mathbf{y}^x| = \sum_{j=1}^p \mathbb{1}(\gamma_j^x > 0)$ and $|\mathbf{y}^z| = \sum_{k=1}^q \gamma_k^z$, we have

$$\begin{aligned} \pi(|\mathbf{y}| = k) &= \sum_{\ell=0 \vee (k-q)}^{p \wedge k} \pi(|\mathbf{y}^x| = \ell, |\mathbf{y}^z| = k - \ell) \\ &= \sum_{\ell=0 \vee (k-q)}^{p \wedge k} \sum_{\mathbf{y}: |\mathbf{y}^x| = \ell, |\mathbf{y}^z| = k - \ell} \pi(\mathbf{y}). \end{aligned}$$

Note that the cardinality of the set $\{\mathbf{y} : |\mathbf{y}^x| = \ell, |\mathbf{y}^z| = k - \ell\}$ for $\ell \in \{0 \vee (k - q), \dots, p \wedge k\}$ is $\binom{p}{\ell} \binom{q}{k-\ell} 2^\ell$ and $\pi(\mathbf{y})$ is invariant on this set, where the factor 2^ℓ appears because each γ_j^x has two nonzero values. Hence, the rightmost side of the last equation equals

$$\begin{aligned} \sum_{\ell=0 \vee (k-q)}^{p \wedge k} \binom{p}{\ell} \binom{q}{k-\ell} B(k+1, p+q-k+1) \\ = \binom{p+q}{k} B(k+1, p+q-k+1) = \frac{1}{p+q+1}, \end{aligned}$$

which completes the proof.

Appendix C. Proof of Proposition 2

Define $\mathbf{y}_{-j}^x$ such that $\mathbf{y}^x = \{\gamma_j^x, \mathbf{y}_{-j}^x\}$ and let $|\mathbf{y}_{-j}^x| = |\mathbf{y}^x| - \mathbb{1}(\gamma_j^x > 0)$. Using the same argument in the proof of Proposition 1,

$$\begin{aligned} \pi(\gamma_j^x = 0) &= \sum_{k=0}^{p+q-1} \sum_{\ell=0 \vee (k-q)}^{(p-1) \wedge k} \pi(\gamma_j^x = 0, |\mathbf{y}_{-j}^x| = \ell, |\mathbf{y}^z| = k - \ell) \\ &= \sum_{k=0}^{p+q-1} \binom{p+q-1}{k} B(k+1, p+q-k+1) \\ &= \sum_{k=0}^{p+q-1} \frac{p+q-k}{(p+q)(p+q+1)} = \frac{1}{2}. \end{aligned}$$

Because $\pi(\gamma_j^x = 1)$ and $\pi(\gamma_j^x = 2)$ are equal by the definition in (9), it follows that $\pi(\gamma_j^x = 1) = \pi(\gamma_j^x = 2) = 1/4$. Similarly, $\pi(\gamma_j^z = 0) = 1/2$ can be verified.

References

- Baladandayuthapani, V., Mallick, B. K., and Carroll, R. J. (2005), “Spatially Adaptive Bayesian Penalized Regression Splines (P-splines),” *Journal of Computational and Graphical Statistics*, 14, 378–394. [325]
- Banerjee, S., and Ghosal, S. (2014), “Bayesian Variable Selection in Generalized Additive Partial Linear Models,” *Statistics*, 3, 363–378. [324,333]
- Barbieri, M. M., and Berger, J. O. (2004), “Optimal Predictive Model Selection,” *The Annals of Statistics*, 32, 870–897. [330]
- Bayarri, M. J., Berger, J. O., Forte, A., and García-Donato, G. (2012), “Criteria for Bayesian Model Choice With Application to Variable Selection,” *The Annals of Statistics*, 40, 1550–1577. [327]
- Belitser, E., and Serra, P. (2014), “Adaptive Priors Based on Splines With Random Knots,” *Bayesian Analysis*, 9, 859–882. [328]
- Berger, J. O., Pericchi, L. R., and Varshavsky, J. A. (1998), “Bayes Factors and Marginal Distributions in Invariant Situations,” *Sankhyā: The Indian Journal of Statistics*, 60, 307–321. [327]
- Carlin, B. P., and Chib, S. (1995), “Bayesian Model Choice Via Markov Chain Monte Carlo Methods,” *Journal of the Royal Statistical Society, Series B*, 57, 473–484. [325,328]
- Cripps, E., Carter, C., and Kohn, R. (2005), “Variable Selection and Covariance Selection in Multivariate Regression Models,” *Handbook of Statistics*, 25, 519–552. [327]
- Curtis, S. M., Banerjee, S., and Ghosal, S. (2014), “Fast Bayesian Model Assessment for Nonparametric Additive Regression,” *Computational Statistics & Data Analysis*, 71, 347–358. [324]
- de Boor, C. (1978), *A Practical Guide to Splines*, New York: Springer. [334]
- Dellaportas, P., Forster, J. J., and Ntzoufras, I. (2002), “On Bayesian Model and Variable Selection Using MCMC,” *Statistics and Computing*, 12, 27–36. [328]
- Denison, D. G., Mallick, B. K., and Smith, A. F. (1998), “Bayesian MARS,” *Statistics and Computing*, 8, 337–346. [325]
- DiMatteo, I., Genovese, C. R., and Kass, R. E. (2001), “Bayesian Curve-Fitting With Free-Knot Splines,” *Biometrika*, 88, 1055–1071. [325]
- Faure, H., Preziosi, P., Roussel, A., Bertrais, S., Galan, P., Hercberg, S., and Favier, A. (2006), “Factors Influencing Blood Concentration of Retinol, α -Tocopherol, Vitamin c, and β -Carotene in the French Participants of the Su. VI. Max Trial,” *European Journal of Clinical Nutrition*, 60, 706–717. [333]
- Fernandez, C., Ley, E., and Steel, M. F. (2001), “Benchmark Priors for Bayesian Model Averaging,” *Journal of Econometrics*, 100, 381–427. [327]
- Fletcher, R. H., and Fairfield, K. M. (2002), “Vitamins for Chronic Disease Prevention in Adults: Clinical Applications,” *Journal of the American Medical Association*, 287, 3127–3129. [333]
- George, E. I., and Foster, D. P. (2000), “Calibration and Empirical Bayes Variable Selection,” *Biometrika*, 87, 731–747. [327]
- Hobert, J. P., and Casella, G. (1998), “Functional Compatibility, Markov Chains, and Gibbs Sampling With Improper Posteriors,” *Journal of Computational and Graphical Statistics*, 7, 42–60. [331]
- Huang, J., Horowitz, J. L., and Wei, F. (2010), “Variable Selection in Nonparametric Additive Models,” *The Annals of Statistics*, 38, 2282–2313. [324]
- Jeong, S., Park, M., and Park, T. (2017), “Analysis of Binary Longitudinal Data With Time-Varying Effects,” *Computational Statistics and Data Analysis*, 112, 145–153. [333]
- Jeong, S., and Park, T. (2016), “Bayesian Semiparametric Inference on Functional Relationships in Linear Mixed Models,” *Bayesian Analysis*, 11, 1137–1163. [333]
- Ji, C., and Schmidler, S. C. (2013), “Adaptive Markov Chain Monte Carlo for Bayesian Variable Selection,” *Journal of Computational and Graphical Statistics*, 22, 708–728. [334]
- Jullion, A., and Lambert, P. (2007), “Robust Specification of the Roughness Penalty Prior Distribution in Spatially Adaptive Bayesian p-Splines Models,” *Computational Statistics & Data Analysis*, 51, 2542–2558. [325]
- Kohn, R., Smith, M., and Chan, D. (2001), “Nonparametric Regression Using Linear Combinations of Basis Functions,” *Statistics and Computing*, 11, 313–322. [325,326,327,330,333,334]
- Lang, S., and Brezger, A. (2004), “Bayesian P-Splines,” *Journal of Computational and Graphical Statistics*, 13, 183–212. [325]
- Li, Y., and Clyde, M. A. (2018), “Mixtures of g-Priors in Generalized Linear Models,” *Journal of the American Statistical Association*, 113, 1828–1845. [334]
- Lian, H., Liang, H., and Ruppert, D. (2015), “Separation of Covariates Into Nonparametric and Parametric Parts in High-Dimensional Partially Linear Additive Models,” *Statistica Sinica*, 25, 591–607. [324]
- Liang, F., Paulo, R., Molina, G., Clyde, M. A., and Berger, J. O. (2008), “Mixtures of g Priors for Bayesian Variable Selection,” *Journal of the American Statistical Association*, 103, 410–423. [327]
- Lin, Y., and Zhang, H. H. (2006), “Component Selection and Smoothing in Multivariate Nonparametric Regression,” *The Annals of Statistics*, 34, 2272–2297. [324]
- Liu, X., Wang, L., and Liang, H. (2011), “Estimation and Variable Selection for Semiparametric Additive Partial Linear Models,” *Statistica Sinica*, 21, 1225–1248. [333]
- Maruyama, Y., and George, E. I. (2011), “Fully Bayes Factors With a Generalized g-Prior,” *The Annals of Statistics*, 39, 2740–2765. [327]
- Moreno, E., Girón, J., and Casella, G. (2015), “Posterior Model Consistency in Variable Selection as the Model Dimension Grows,” *Statistical Science*, 30, 228–241. [327]

- Nierenberg, D. W., Stukel, T. A., Baron, J. A., Dain, B. J., and Greenberg, E. R. (1989), "Determinants of Plasma Levels of Beta-Carotene and Retinol," *American Journal of Epidemiology*, 130, 511–521. [333]
- Nott, D. J., and Kohn, R. (2005), "Adaptive Sampling for Bayesian Variable Selection," *Biometrika*, 92, 747–763. [334]
- Park, T., and Jeong, S. (2018), "Analysis of Poisson Varying-Coefficient Models With Autoregression," *Statistics*, 52, 34–49. [333]
- Park, T., and Lee, S. (2021), "Improving the Gibbs Sampler," *Wiley Interdisciplinary Reviews: Computational Statistics*, e1546. [331]
- Park, T., and van Dyk, D. A. (2009), "Partially Collapsed Gibbs Samplers: Illustrations and Applications," *Journal of Computational and Graphical Statistics*, 18, 283–305. [331]
- Raftery, A. E., Madigan, D., and Hoeting, J. A. (1997), "Bayesian Model Averaging for Linear Regression Models," *Journal of the American Statistical Association*, 92, 179–191. [331]
- Raskutti, G., Wainwright, M. J., and Yu, B. (2012), "Minimax-Optimal Rates for Sparse Additive Models Over Kernel Classes Via Convex Programming," *Journal of Machine Learning Research*, 13, 389–427. [324]
- Ravikumar, P., Liu, H., Lafferty, J., and Wasserman, L. (2009), "Sparse Additive Models," *Journal of the Royal Statistical Society, Series B*, 71, 1009–1030. [324]
- Rossell, D., and Rubio, F. J. (2019), "Additive Bayesian Variable Selection Under Censoring and Misspecification," arXiv:1907.13563. [324,325,326]
- Ruppert, D., Wand, M. P., and Carroll, R. J. (2003), *Semiparametric Regression*, Cambridge: Cambridge University Press. [325]
- Sabanés Bové, D., and Held, L. (2011), "Hyper-g Priors for Generalized Linear Models," *Bayesian Analysis*, 6, 387–410. [334]
- Sabanés Bové, D., Held, L., and Kauermann, G. (2015), "Objective Bayesian Model Selection in Generalized Additive Models With Penalized Splines," *Journal of Computational and Graphical Statistics* 24(2), 394–415. [325,332]
- Scheipl, F., Fahrmeir, L., and Kneib, T. (2012), "Spike-and-Slab Priors for Function Selection in Structured Additive Regression Models," *Journal of the American Statistical Association*, 107, 1518–1532. [325,332]
- Scheipl, F., and Kneib, T. (2009), "Locally Adaptive Bayesian p-Splines With a Normal-Exponential-Gamma Prior," *Computational Statistics & Data Analysis*, 53, 3533–3552. [325]
- Scott, J. G., and Berger, J. O. (2010), "Bayes and Empirical-Bayes Multiplicity Adjustment in the Variable-Selection Problem," *The Annals of Statistics*, 38, 2587–2619. [327]
- Shen, W., and Ghosal, S. (2015), "Adaptive Bayesian Procedures Using Random Series Priors," *Scandinavian Journal of Statistics*, 42, 1194–1213. [328]
- Smith, M., and Kohn, R. (1996), "Nonparametric Regression Using Bayesian Variable Selection," *Journal of Econometrics*, 75, 317–343. [325,333]
- van Dyk, D. A., DeGennaro, S., Stein, N., Jefferys, W. H., and von Hippel, T. (2009), "Statistical Analysis of Stellar Evolution," *The Annals of Applied Statistics*, 3, 117–143. [328]
- van Dyk, D. A., and Jiao, X. (2015), "Metropolis–Hastings Within Partially Collapsed Gibbs Samplers," *Journal of Computational and Graphical Statistics*, 24, 301–327. [331]
- van Dyk, D. A., and Park, T. (2008), "Partially Collapsed Gibbs Samplers: Theory and Methods," *Journal of the American Statistical Association*, 193, 790–796. [325,331]
- Vats, D., Flegal, J. M., and Jones, G. L. (2019), "Multivariate Output Analysis for Markov Chain Monte Carlo," *Biometrika*, 106, 321–337. [332]
- Wang, M., and Maruyama, Y. (2018), "Posterior Consistency of g-Prior for Variable Selection With a Growing Number of Parameters," *Journal of Statistical Planning and Inference*, 196, 19–29. [327]
- Wu, Y., and Stefanski, L. A. (2015), "Automatic Structure Recovery for Additive Models," *Biometrika*, 102, 381–395. [324]
- Xue, L. (2009), "Consistent Variable Selection in Additive Models," *Statistica Sinica*, 19, 1281–1296. [324]
- Zellner, A. (1986), "On Assessing Prior Distributions and Bayesian Regression Analysis With g-Prior Distributions," *Bayesian Inference and Decision Techniques: Essays in Honor of Bruno De Finetti*, 6, 233–243. [325,327]