

Adaptive Bayesian Nonparametric Regression Using a Kernel Mixture of Polynomials with Application to Partial Linear Models

Fangzheng Xie^{*} and Yanxun Xu[†]

Abstract. We propose a kernel mixture of polynomials prior for Bayesian nonparametric regression. The regression function is modeled by local averages of polynomials with kernel mixture weights. We obtain the minimax-optimal contraction rate of the full posterior distribution up to a logarithmic factor by estimating metric entropies of certain function classes. Under the assumption that the degree of the polynomials is larger than the unknown smoothness level of the true function, the posterior contraction behavior can adapt to this smoothness level provided an upper bound is known. We also provide a frequentist sieve maximum likelihood estimator with a near-optimal convergence rate. We further investigate the application of the kernel mixture of polynomials to partial linear models and obtain both the near-optimal rate of contraction for the nonparametric component and the Bernstein-von Mises limit (i.e., asymptotic normality) of the parametric component. The proposed method is illustrated with numerical examples and shows superior performance in terms of computational efficiency, accuracy, and uncertainty quantification compared to the local polynomial regression, DiceKriging, and the robust Gaussian stochastic process.

Keywords: Bayesian nonparametric regression, Bernstein-von Mises limit, metric entropies, partial linear model, rate of contraction.

1 Introduction

The standard nonparametric regression model is of the form $y_i = f(\mathbf{x}_i) + e_i$, where y_i 's are observations at given design points, \mathbf{x}_i 's are in the design space $\mathcal{X} \subset \mathbb{R}^p$, and e_i 's are independently $N(0, \sigma^2)$ distributed residuals, $i = 1, \dots, n$. The inference task is to estimate the unknown function $f : \mathcal{X} \rightarrow \mathbb{R}$. Nonparametric regression methods have been widely used in a variety of applications, such as pattern recognition (Györfi et al., 2006; Devroye et al., 2013), image processing and reconstruction (Takeda et al., 2007), electronic healthcare records (Xu et al., 2016b), and semiparametric econometrics (Robinson, 1988; Klein and Spady, 1993).

Frequentist methods for nonparametric regression typically compute a fixed estimated function through the given data $(\mathbf{x}_i, y_i)_{i=1}^n$. In contrast, Bayesian nonparametric techniques first impose a carefully-selected prior distribution on the unknown function f and then find the posterior distribution of f given the observed data $(\mathbf{x}_i, y_i)_{i=1}^n$, providing a natural way for uncertainty quantification through the full posterior distributions

^{*}Department of Applied Mathematics and Statistics, Johns Hopkins University, fxie5@jhu.edu

[†]Department of Applied Mathematics and Statistics, Johns Hopkins University, yanxun.xu@jhu.edu

instead of a point estimate given by frequentist approaches. One of the most popular Bayesian nonparametric regression methods is the Gaussian process (Rasmussen and Williams, 2006) due to its tractability. Nevertheless the computational burden of the Gaussian process in likelihood function evaluation resulting from the inversion of the covariance matrix prevents its scalability to big data.

In this paper, we propose a novel prior model for nonparametric regression, called *the kernel mixture of polynomials*, that features attractive theoretical properties, efficient computation, and flexibility for extension. Theoretically, we show that by using the kernel mixture of polynomials for Bayesian nonparametric regression, the rate of contraction with respect to the L_2 -topology is minimax-optimal (Stone, 1982; Györfi et al., 2006) (up to a logarithmic factor) and, under the assumption that the degree of the polynomials is large enough, is adaptive to the smoothness level in the sense that the prior does not depend on the smoothness level of the true function. It is worth mentioning that most papers concerning posterior convergence for Bayesian nonparametric regression only discuss the rate of contraction with respect to the weaker empirical L_2 -norm (van der Vaart and van Zanten, 2008, 2009; De Jonge et al., 2010; Bhattacharya et al., 2014), i.e., the convergence of the function at the given design points. There is little discussion about the rate of contraction with respect to the exact L_2 -norm for general Bayesian nonparametric regression methods. van der Vaart and Zanten (2011), Yoo et al. (2016), and Yang et al. (2017) address this issue only in the context of Gaussian process regression. In particular, the rate of contraction for Gaussian processes with respect to the exact L_2 -norm requires prior knowledge of the smoothness level of the true regression function. We also obtain a sieve maximum likelihood estimator with the near-optimal convergence rate as a frequentist point estimator, which could potentially be useful for designing scalable optimization algorithms.

From the computational perspective, the proposed kernel mixture of polynomials model avoids the cumbersome $O(n^3)$ inversion of large covariance matrices, greatly facilitating computational efficiency for posterior inference compared to Gaussian process priors, while maintaining the same level of accuracy (see Section 5.2 for detailed comparisons). Such a nice computational advantage makes it attractive to the big-data regime. The code for implementation is publicly available at <https://github.com/fangzhengxie/Kernel-Mixture-of-Polynomials.git>.

In addition, the kernel mixture of polynomials is flexible for extension due to its nonparametric nature. As a specific example, we study the application of this prior model to the partial linear model. The partial linear model is a classical semiparametric regression model of the form $y_i = \mathbf{z}_i^T \boldsymbol{\beta} + \eta(\mathbf{x}_i) + e_i$, where $\mathbf{z}_i, \mathbf{x}_i$'s are design points, $\boldsymbol{\beta}$ is the linear coefficient, η is some unknown function, and e_i 's are independent $N(0, 1)$ residuals, $i = 1, \dots, n$. The literatures of partial linear models from both the frequentist perspective (Engle et al., 1986; Chen et al., 1988; Speckman, 1988; Hastie and Tibshirani, 1990; Fan and Li, 1999) and Bayesian approaches (Lenk, 1999; Bickel et al., 2012; Tang et al., 2015; Yang et al., 2015) are rich. However, there is little discussion regarding the theoretical properties of the Bayesian partial linear model. To the best of our knowledge, only Bickel et al. (2012) and Yang et al. (2015) discuss the asymptotic behavior of the marginal posterior distribution of $\boldsymbol{\beta}$ with Gaussian process priors on η .

We impose the kernel mixture of polynomials prior on η and obtain both a near-optimal rate of contraction for η and the Bernstein-von Mises limit (i.e., asymptotic normality) of the marginal posterior of β .

The layout of this paper is as follows. Section 2 presents the setup for the kernel mixture of polynomials prior for nonparametric regression. Section 3 elaborates on the convergence properties of the kernel mixture of polynomials for nonparametric regression. Section 4 presents the application of the kernel mixture of polynomials to the partial linear model. Section 5 illustrates the proposed methodology using numerical examples. We conclude the paper with several discussions in Section 6.

2 Preliminaries

2.1 Setup

Recall that the Gaussian nonparametric regression is of the form

$$\begin{aligned} y_i &= f(\mathbf{x}_i) + e_i, \\ \mathbf{x}_i &\sim p_{\mathbf{x}} \text{ independently,} \\ e_i &\sim N(0, \sigma^2) \text{ independently, } i = 1, \dots, n, \end{aligned}$$

where $p_{\mathbf{x}}$ is the marginal density of the design points $(\mathbf{x}_i)_{i=1}^n$ supported on the p -dimensional unit hypercube $\mathcal{X} = [0, 1]^p$. We assume that $p_{\mathbf{x}}$ is bounded away from 0 and ∞ , and is known and fixed. The true but unknown regression function f_0 is assumed to be in the α -Hölder function class $\mathfrak{C}^{\alpha, B}(\mathcal{X})$ with envelope B , the class of functions f that are $\lceil \alpha - 1 \rceil$ -times continuously differentiable with

$$\max_{|\mathbf{s}| \leq \lceil \alpha - 1 \rceil} \|D^{\mathbf{s}} f\|_{\infty} + \max_{|\mathbf{s}| = \lceil \alpha - 1 \rceil} \sup_{\mathbf{x}_1 \neq \mathbf{x}_2} \frac{|D^{\mathbf{s}} f(\mathbf{x}_1) - D^{\mathbf{s}} f(\mathbf{x}_2)|}{\|\mathbf{x}_1 - \mathbf{x}_2\|^{\alpha - |\mathbf{s}|}} \leq B$$

for all $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{X}$, where $\lceil \alpha - 1 \rceil$ denotes the minimum integer no less than $\alpha - 1$, $D^{\mathbf{s}} = \partial^{|\mathbf{s}|} / \partial x_1^{s_1} \dots \partial x_p^{s_p}$ is the mixed partial derivative operator, and $|\mathbf{s}| = \sum_{j=1}^p s_j$. When $|\mathbf{s}| = 0$, by convention we define $D^{\mathbf{0}} f(\mathbf{x}) = f(\mathbf{x})$. We use \mathbb{P}_0 and \mathbb{E}_0 to denote the probability and expected value under p_0 , respectively. For readers' convenience, descriptions of additional notations are provided in Section A in the Supplementary Material (Xie and Xu, 2019).

The goal is to estimate the unknown function f_0 . Leaving the Bayesian framework for a moment, let us consider the frequentist Nadaraya-Watson estimator (Nadaraya, 1964; Watson, 1964) of the form

$$\hat{f}(\mathbf{x}) = \sum_{i=1}^n \left[\frac{\varphi_h(\mathbf{x} - \mathbf{x}_i)}{\sum_{i=1}^n \varphi_h(\mathbf{x} - \mathbf{x}_i)} \right] y_i, \quad (1)$$

where $\varphi_h : \mathbb{R}^p \rightarrow [0, +\infty)$ is the kernel function parametrized by the bandwidth parameter $h \in (0, +\infty)$ and is assumed to decrease when $\|\mathbf{x}\|$ increases. It is a local averaging

estimator (Györfi et al., 2006), since the summand $\varphi_h(\mathbf{x} - \mathbf{x}_i)/\sum_i \varphi_h(\mathbf{x} - \mathbf{x}_i)$ can be treated as the weight received from y_i . As such a simple local averaging estimator does not yield an optimal rate of convergence when the true regression function is α -Hölder for $\alpha \geq 2$ (Devroye et al., 2013), Fan and Gijbels (1996) considers a more general local polynomial regression to capture higher-order curvature information of the unknown regression function and gain an optimal rate of convergence. Inspired by these two classical approaches for nonparametric regression, we develop the kernel mixture of polynomials model.

Firstly, for a given integer K , we partition the design space \mathcal{X} into K^p disjoint hypercubes: $\mathcal{X} = \bigcup_{\mathbf{k} \in [K]^p} \mathcal{X}_K(\mathbf{k})$, where $\mathcal{X}_K(\mathbf{k}) = \prod_{j=1}^p ((k_j - 1)/K, k_j/K]$, $\boldsymbol{\mu}_{\mathbf{k}}^* = [(2k_1 - 1)/(2K), \dots, (2k_p - 1)/(2K)]^T$, $\mathbf{k} = [k_1, \dots, k_p]^T \in [K]^p$, and $\boldsymbol{\mu}_{\mathbf{k}}^*$ is the center of $\mathcal{X}_K(\mathbf{k})$. The idea of the partition is to breakdown the problem of estimating the regression function over the entire domain \mathcal{X} into sub-problems of estimating the regression function within each block.

Next, we introduce the notion of (boxed) kernel functions in order to detect the local behavior of the underlying regression function f_0 around the block $\mathcal{X}_K(\mathbf{k})$. Formally, a continuous function $\varphi : \mathbb{R}^p \rightarrow [0, 1]$ is a boxed kernel function if it is supported on $\{\mathbf{x} : \|\mathbf{x}\|_\infty \leq 1\}$, does not increase when $\|\mathbf{x}\|_\infty$ increases, and $\varphi(\mathbf{x}) \leq \mathbf{1}(\|\mathbf{x}\|_\infty \leq 1)$. We consider the univariate bump kernel $\varphi(x) = \exp[-(1 - x^2)^{-1}] \mathbf{1}(|x| < 1)$ in this paper. Other examples of boxed kernels include the triangle kernel, the Epanechnikov kernel, etc. For convenience we denote $\varphi_h(\mathbf{x}) = \varphi(\mathbf{x}/h)$, where $h > 0$ is the bandwidth parameter. For each $K \in \mathbb{N}_+$ and each $\mathbf{k} \in [K]^p$, define the kernel mixture weight as

$$w_{\mathbf{k}}(\mathbf{x}) = \frac{\varphi_h(\mathbf{x} - \boldsymbol{\mu}_{\mathbf{k}})}{\sum_{\mathbf{l} \in [K]^p} \varphi_h(\mathbf{x} - \boldsymbol{\mu}_{\mathbf{l}})}, \quad (2)$$

where $\boldsymbol{\mu}_{\mathbf{k}} \in \mathcal{X}_K(\mathbf{k})$, and $h > 0$. The kernel mixture weight $w_{\mathbf{k}}(\mathbf{x})$ is motivated by the form of Nadaraya-Watson estimator (1) and is designed to average the signal from f_0 locally around $\boldsymbol{\mu}_{\mathbf{k}}$. It can also be shown that the denominator $D(\mathbf{x}) := \sum_{\mathbf{l}} \varphi_h(\mathbf{x} - \boldsymbol{\mu}_{\mathbf{l}})$ is strictly non-zero by considering the structure of the disjoint blocks $\mathcal{X}_K(\mathbf{k})$'s. See Figure 1 for some examples of the kernel mixture weights $w_k(x)$ in the univariate case.

Thirdly, in order to enhance the estimation accuracy, we need additional tools to capture the higher-order curvature information (*e.g.*, gradient, Hessian, etc.) of the underlying regression function surface, motivated by the local polynomial regression. Given K and the kernel mixture weights (2), we define the kernel mixture of polynomials system to be the set of functions of the form

$$\psi_{\mathbf{k}\mathbf{s}}(\mathbf{x}) = w_{\mathbf{k}}(\mathbf{x})(\mathbf{x} - \boldsymbol{\mu}_{\mathbf{k}}^*)^{\mathbf{s}}, \quad \mathbf{k} \in [K]^p, \quad \mathbf{s} \in \{\mathbf{s} \in [m]^p : |\mathbf{s}| = 0, 1, \dots, m\}, \quad (3)$$

where $(\mathbf{x} - \boldsymbol{\mu}_{\mathbf{k}}^*)^{\mathbf{s}}$ denotes the monomial $\prod_{j=1}^p [x_j - (2k_j - 1)/2K]^{s_j}$, $\mathbf{x} = [x_1, \dots, x_p]^T$, and $\mathbf{s} = [s_1, \dots, s_p]^T$. In this work we require that the degree m of the kernel mixture of polynomials system to be no less than α , i.e., $m \geq \alpha$. A similar assumption was also adopted in De Jonge et al. (2012) in the context of conditional Gaussian tensor-product spline models, where the degree of polynomials in the B-splines was required to be no less than the smoothness level of the underlying true function. Note that for

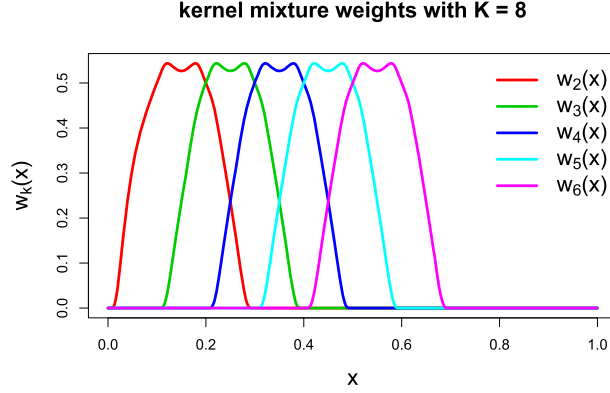


Figure 1: Examples of $w_k(x)$ for $k = 2, 3, 4, 5, 6$ with $K = 8$ and equidistant $\mu_k = (2k - 1)/(2K)$, $k \in [K]$.

any $m < \alpha$, f_0 being an α -Hölder function implies that f_0 is also an m -Hölder function. Therefore, the analysis in this work can be easily adapted to the case where the relation between m and α is unknown by simply replacing α with $\min\{m, \alpha\}$. When $\mathbf{s} = \mathbf{0}$, $\psi_{\mathbf{ks}}$ closely resembles the term $\varphi_h(\mathbf{x} - \mathbf{x}_i)/\sum_i \varphi_h(\mathbf{x} - \mathbf{x}_i)$ appearing in the summation of the Nadaraya-Watson estimator (1). When $\mathbf{s} \neq \mathbf{0}$, $\psi_{\mathbf{ks}}$ is able to capture higher-order curvature information from the underlying regression function due to its polynomial structure.

Finally, with the above ingredients, we define the kernel mixture of polynomials model by $\mathcal{M} = \{p_{f,\sigma}(\mathbf{x}, y) : f \in \bigcup_{K=1}^{\infty} \mathcal{F}_K, \sigma \in [\underline{\sigma}, \bar{\sigma}]\}$, a class of distributions indexed by the regression function f and the standard deviation of the noise σ , where $p_{f,\sigma}(\mathbf{x}, y) = \phi_{\sigma}(y - f(\mathbf{x}))p_{\mathbf{x}}(\mathbf{x})$, f lies in the union of the function classes $\mathcal{F}_K = \bigcup_{h \in [\underline{h}, \bar{h}]} \mathcal{F}_K(h)$ for some constants \underline{h}, \bar{h} with $1 < \underline{h} < \bar{h} < \infty$, and

$$\mathcal{F}_K(h) = \left\{ \sum_{\mathbf{k} \in [K]^p} \sum_{\mathbf{s}: |\mathbf{s}|=0}^m \xi_{\mathbf{ks}} \psi_{\mathbf{ks}}(\mathbf{x}) : \mu_{\mathbf{k}} \in \mathcal{X}_K(\mathbf{k}), \max_{|\mathbf{s}|=0, \dots, m} |\xi_{\mathbf{ks}}| \leq B, \mathbf{k} \in [K]^p \right\}. \quad (4)$$

The parameters σ , K , $(\mu_{\mathbf{k}} : \mathbf{k} \in [K]^p)$, and $(\xi_{\mathbf{ks}} : \mathbf{k} \in [K]^p, |\mathbf{s}| = 0, 1, \dots, m)$ are to be assigned a hierarchical prior in Section 2.2. In other words, the kernel mixture of polynomials model parametrizes the regression function f through

$$f(\mathbf{x}) = \sum_{\mathbf{k} \in [K]^p} \sum_{\mathbf{s}: |\mathbf{s}|=0}^m \xi_{\mathbf{ks}} \psi_{\mathbf{ks}}(\mathbf{x}), \quad (5)$$

where $\{\psi_{\mathbf{ks}}(\mathbf{x})\}_{\mathbf{ks}}$ serves as certain basis functions that mimic the behavior of f locally around $\mu_{\mathbf{k}} \in \mathcal{X}_K(\mathbf{k})$, and $\xi_{\mathbf{ks}}$ is the coefficient or the amplitude of the corresponding basis function $\psi_{\mathbf{ks}}(\mathbf{x})$.

Remark 1. In order that the centers $\mu_{\mathbf{k}}$'s in the kernel mixture weights are helpful to detecting both the local and global behavior of f , they need to be relatively spread. This also explains why in the first step above, we require that μ_k 's lie in the disjoint blocks $\in \mathcal{X}_K(\mathbf{k})$'s. In addition, such a partition restriction avoids the label-switching phenomenon in the Markov chain Monte Carlo sampler for posterior inference, which could potentially affect the mixing of the Markov chain (Celeux et al., 2000; Jasra et al., 2005). Alternatively, repulsive priors (Affandi et al., 2013; Xu et al., 2016a; Xie and Xu, 2017) can be incorporated to gain well-spread kernel centers, but this could still cause the label-switching issue.

Remark 2. The kernel bandwidth parameter h , which is typically endowed with a prior in the literature of Bayesian kernel methods (Ghosal et al., 2007a; Shen et al., 2013), also plays a key role in establishing the convergence properties of the kernel mixture of polynomials. In the current context, h can be arbitrarily close to zero asymptotically, since we require that K ranges over all positive integers and that Kh stays bounded as $K \rightarrow \infty$. Therefore, $\bigcup_{K=1}^{\infty} \mathcal{F}_K$ is rich enough to provide good approximation to arbitrary $f_0 \in \mathfrak{C}^{\alpha, B}(\mathcal{X})$.

Remark 3. In a related work of Rockova and van der Pas (2017), the authors consider the posterior contraction of Bayesian regression trees and their ensembles, which are essentially piece-wise constant functions supported on hypercube partitions. Due to the non-smoothness nature of these classes of functions, their contraction results are restricted to functions with smoothness level $\alpha \leq 1$. In contrast, the proposed kernel mixture of polynomials model can be regarded as a smoother version of the Bayesian regression trees, and, as will be seen in Section 3, the contraction result holds for smoother functions as well ($\alpha \leq m$).

2.2 Prior specification

We define a prior distribution Π for (f, σ) through (5) by imposing the hierarchical priors on the parameters $(\mu_{\mathbf{k}} : \mathbf{k} \in [K]^p)$, $(\xi_{\mathbf{k}\mathbf{s}} : \mathbf{k} \in [K]^p, |\mathbf{s}| = 0, 1, \dots, m)$, σ , h , and K as follows:

- The standard deviation σ of the residuals $(e_i)_{i=1}^n$ follows π_{σ} that is continuous and non-vanishing on $[\underline{\sigma}, \bar{\sigma}]$, independent of the remaining parameters.
- The prior for K satisfies the following condition:

$$\exp[-b_0 x^p (\log x^p)^{r_0}] \lesssim \Pi(K \geq x) \lesssim \exp[-b_1 x^p (\log x^p)^{r_0}] \quad (6)$$

for some constants $b_0, b_1 > 0$, $b_1 \leq b_0$, and $r_0 \geq 0$.

Given K , the conditional prior for the rest of the parameters are given as follows:

- The kernel centers $\mu_{\mathbf{k}}$'s are sampled as $\mu_{\mathbf{k}} = \mu_{\mathbf{k}}^* + \tilde{\mu}_{\mathbf{k}}/(2K)$, where $\tilde{\mu}_{\mathbf{k}}$ independently follows π_{μ} for each $\mathbf{k} \in [K]^p$ for some continuous non-vanishing density π_{μ} on the hypercube $[-1, 1]^p$. For example, π_{μ} can be taken as the uniform distribution on $[-1, 1]^p$.

- The coefficient $\xi_{\mathbf{ks}}$ for $\psi_{\mathbf{ks}}(\mathbf{x})$ follow π_ξ independently for each $\mathbf{s} \in [m]^p$ with $|\mathbf{s}| = 0, 1, \dots, m$ given K , where π_ξ is a continuous non-vanishing density on $[-B, B]$. For example, one can take π_ξ to be the normal distribution truncated on $[-B, B]$, or the uniform distribution on $[-B, B]$.
- The bandwidth h follows a non-vanishing density $\pi(h \mid K)$ supported on $[\underline{h}/K, \bar{h}/K]$ given K , where $1 < \underline{h} < \bar{h}$. The uniform distribution on $[\underline{h}/K, \bar{h}/K]$ satisfies the condition.

Remark 4. Kruijer et al. (2010) adopts the same tail probability condition (6) for the number of support points in the context of nonparametric density estimation. Special cases of (6) include the geometric distribution when $r_0 = 0$, and the Poisson distribution when $r_0 = 1$. In Section 3, we show that both $r_0 = 0$ and $r_0 = 1$ yield the same rate of contraction, but any $r_0 > 1$ (i.e., thinner tail) leads to a slower rate of contraction.

3 Convergence properties of the kernel mixture of polynomials regression

In this section, we establish the convergence results of the kernel mixture of polynomials for nonparametric regression, and obtain a frequentist sieve maximum likelihood estimator with the corresponding convergence rate. For nonparametric regression problems, when the true regression function f_0 is in $\mathfrak{C}^{\alpha, B}(\mathcal{X})$, $p_{\mathbf{x}}(\mathbf{x}) = 1$, and $e_i \sim N(0, 1)$, $i = 1, \dots, n$, it has been shown that the minimax rate of convergence for any estimator with respect to the L_2 -norm is $n^{-\alpha/(2\alpha+p)}$ (Stone, 1982; Györfi et al., 2006). The optimal rate of contraction cannot be faster than the minimax rate of convergence. Theorem 1 below, which is one of the main results of this section, asserts that the rate of contraction with respect to the $L_2(\mathbb{P}_{\mathbf{x}})$ -topology is minimax-optimal up to a logarithmic factor. Furthermore, under the assumption that the degree m of the kernel mixture of polynomials system is no less than the smoothness level α , rate of contraction is adaptive to the smoothness level α of the underlying true f_0 .

Theorem 1 (Rate of contraction). *Assume that f_0 is in the α -Hölder function class $\mathfrak{C}^{\alpha, B}(\mathcal{X})$ with envelope B . Suppose Π is the prior constructed in Section 2.2 with $m \geq \alpha$. Then for some large constant $M > 0$, it holds that*

$$\Pi(\|f - f_0\|_{L_2(\mathbb{P}_{\mathbf{x}})}^2 > M\epsilon_n^2 \mid \mathcal{D}_n) \rightarrow 0 \quad (7)$$

in \mathbb{P}_0 -probability, where \mathcal{D}_n denotes the data $(\mathbf{x}_i, y_i)_{i=1}^n$, $\epsilon_n = n^{-\alpha/(2\alpha+p)}(\log n)^{t/2}$, and $t > 2\alpha \max(r_0, 1)/(2\alpha + p) + \max(0, 1 - r_0)$.

We sketch the proof below and defer the details to the Supplementary Material (Xie and Xu, 2019). Write the posterior distribution $\Pi(\cdot \mid \mathcal{D}_n)$ as follows:

$$\Pi(\mathcal{A} \mid \mathcal{D}_n) = \frac{\int_{\mathcal{A}} \exp[\ell_n(f, \sigma) - \ell_n(f_0, \sigma_0)] \Pi(df d\sigma)}{\int \exp[\ell_n(f, \sigma) - \ell_n(f_0, \sigma_0)] \Pi(df d\sigma)} := \frac{\mathfrak{N}_n(\mathcal{A})}{\mathfrak{D}_n}$$

for any measurable set \mathcal{A} , where $\ell_n(f, \sigma)$ is the log-likelihood function $\ell_n(f, \sigma) = \sum_{i=1}^n \log p_{f, \sigma}(\mathbf{x}_i, y_i)$. To prove the rates of contraction results in Theorem 1, it suffices to verify a set of sufficient conditions, one variation of the prior-concentration-and-testing framework originally presented by Ghosal et al. (2000). This framework, along with other variations, have been widely applied not only to nonparametric regression, but also to density estimation (Ghosal and van der Vaart, 2001; Ghosal et al., 2007b; Kruijer et al., 2010; Shen et al., 2013) and high-dimensional statistics (Castillo and van der Vaart, 2012; Bhattacharya et al., 2015; Ročková, 2018; Pati et al., 2014; Gao and Zhou, 2015).

For the purpose of illustrating the proof framework, we introduce the notions of covering number and metric entropy. For a metric space (\mathcal{F}, d) , for any $\epsilon > 0$, the ϵ -covering number of (\mathcal{F}, d) , denoted by $\mathcal{N}(\epsilon, \mathcal{F}, d)$, is defined to be the minimum number of ϵ -balls of the form $\{g \in \mathcal{F} : d(f, g) < \epsilon\}$ that are needed to cover \mathcal{F} . The ϵ -bracketing number of (\mathcal{F}, d) , denoted by $\mathcal{N}_{[\cdot]}(\epsilon, \mathcal{F}, d)$, is defined to be the minimum number of brackets of the form $[l_i, u_i]$ that are needed to cover \mathcal{F} such that $l_i, u_i \in \mathcal{F}$ and $d(l_i, u_i) < \epsilon$. We refer to $\log \mathcal{N}(\epsilon, \mathcal{F}, d)$ as the metric entropy, and $\log \mathcal{N}_{[\cdot]}(\epsilon, \mathcal{F}, d)$ as the bracketing (metric) entropy. The bracketing integral $\int_0^\epsilon \sqrt{\log \mathcal{N}_{[\cdot]}(u, \mathcal{F}, d)} du$ is denoted by $J_{[\cdot]}(\epsilon, \mathcal{F}, d)$.

We now specify the prior-concentration-and-testing framework, which are two sufficient conditions for (7). The first one is a prior concentration condition. Namely, there exists another sequence $(\underline{\epsilon}_n)_{n=1}^\infty$ with $\underline{\epsilon}_n \leq \epsilon_n$, such that

$$\Pi(p \in B_{\text{KL}}(p_0, \underline{\epsilon}_n)) \geq \exp(-n\underline{\epsilon}_n^2),$$

where $B_{\text{KL}}(p_0, \epsilon)$ is the Kullback-Leibler ball of radius ϵ centered at p_0 defined as follows:

$$B_{\text{KL}}(p_0, \epsilon) = \left\{ p_{f, \sigma} : D_{\text{KL}}(p_0 \| p_{f, \sigma}) < \epsilon^2, \mathbb{E}_0 \left[\log \frac{p_0(\mathbf{x}, y)}{p_{f, \sigma}(\mathbf{x}, y)} \right]^2 < \epsilon^2 \right\}.$$

Secondly, we require a summability condition: For the sequence $(\underline{\epsilon}_n)_{n=1}^\infty$ above, there exists a sequence of sub-classes of densities $(\mathcal{M}_n)_{n=1}^\infty$, $\mathcal{M}_n \subset \mathcal{M}$ (recall the definition of \mathcal{M} in Section 2.1), and for each \mathcal{M}_n a partition $(\mathcal{M}_{nm})_{m=1}^\infty$ with $\mathcal{M}_n = \bigcup_{m=1}^\infty \mathcal{M}_{nm}$, such that $\Pi(p \in \mathcal{M}_n^c) \leq \exp(-4n\underline{\epsilon}_n^2)$ and

$$\exp(-n\underline{\epsilon}_n^2) \sum_{m=1}^\infty \sqrt{\mathcal{N}(\epsilon_n, \mathcal{M}_{nm}, H)} \sqrt{\Pi(p \in \mathcal{M}_{nm})} \rightarrow 0.$$

The sequence of sub-classes of densities $(\mathcal{M}_n)_{n=1}^\infty$ is referred to as sieves in the literature (Shen and Wong, 1994).

The prior concentration condition plays a fundamental role in proving convergence of posterior distributions, as it guarantees that with probability tending to one, the denominator \mathfrak{D}_n appearing in the posterior distribution $\Pi(\cdot | \mathcal{D}_n)$ does not decay super-exponentially (see, for example, Lemma 8.10 in Ghosal and van der Vaart, 2017):

$$\mathbb{P}_0(\mathfrak{D}_n \geq \Pi(p \in B_{\text{KL}}(p_0, \underline{\epsilon}_n)) \exp[-(1+c)n\underline{\epsilon}_n^2]) \rightarrow 1$$

for any positive $c > 0$. Under the current setup, the prior concentration condition largely depends on how well the kernel mixture of polynomials is able to approximate any α -Hölder function. To this end, we introduce the following lemma to demonstrate the approximation power of the kernel mixture of polynomials.

Lemma 1 (Approximation lemma). *Assume that f_0 is in the α -Hölder function class $\mathfrak{C}^{\alpha,B}(\mathcal{X})$ with envelope B . Let f be of the form (5) and $\mu_{\mathbf{k}} \in \mathcal{X}_K(\mathbf{k})$. Then under the assumption $m \geq \alpha$, there exists some constant C_1 such that for sufficiently small ϵ the following holds whenever $K \geq \epsilon^{-1/\alpha}$,*

$$\begin{aligned} B_K^* &:= \left\{ f : \max_{\mathbf{k} \in [K]^p, |\mathbf{s}|=0,1,\dots,\lceil\alpha-1\rceil} \left| \xi_{\mathbf{k}\mathbf{s}} - \frac{D^{\mathbf{s}} f_0(\mu_{\mathbf{k}}^*)}{s_1! \dots s_p!} \right| \leq \epsilon, \mathbf{k} \in [K]^p \right\} \\ &\subset \left\{ f : \|f - f_0\|_{L_2(\mathbb{P}_{\mathbf{x}})}^2 < C_1 \epsilon^2 \right\}. \end{aligned}$$

The intuition of Lemma 1 is that for each $\mu_{\mathbf{k}}^*$, the function f_0 can be well approximated by the Taylor polynomials of degree m locally around $\mu_{\mathbf{k}}^*$. In fact, the coefficients $\xi_{\mathbf{k}\mathbf{s}}$ are selected to be sufficiently close to the Taylor polynomial coefficients. Then the local approximation effect around $\mu_{\mathbf{k}}^*$ for each $\mathbf{k} \in [K]^p$ is accumulated through the kernel mixture weights $(w_{\mathbf{k}}(\mathbf{x}) : \mathbf{k} \in [K]^p)$. Lemma 1 may be of independent interest for numerical function approximation as well.

Moving forward to the summability condition, it states that there exists a subset \mathcal{M}_n of the entire model \mathcal{M} that occupies most of the prior probability, and at the same time can be covered by the union of a collection of blocks $(\mathcal{M}_{nm})_{m \geq 1}$ with low model complexity (metric entropies). Therefore, it is desired that useful metric entropy bounds can be obtained for the kernel mixture of polynomials model. The following proposition, which is one of the major technical contributions of this paper, directly tackles this issue.

Proposition 1 (Metric entropy bound). *Under the assumption that $m \geq \alpha$, there exists some constant $c_2 > 0$, such that for sufficiently small $\epsilon > 0$ and any $r \in [1, \infty)$,*

$$\log \mathcal{N}_{[\cdot]}(2\epsilon, \mathcal{F}_K, \|\cdot\|_{L_r(\mathbb{P}_{\mathbf{x}})}) \leq \log \mathcal{N}(\epsilon, \mathcal{F}_K, \|\cdot\|_{\infty}) \leq c_2 K^p \left(\log \frac{1}{\epsilon} \right).$$

Besides the rate of contraction, which is a frequentist large sample evaluation of the full posterior distribution, we also obtain a frequentist sieve maximum likelihood estimator with a convergence rate as a result of the metric entropy bounds. This convergence rate is also minimax optimal up to a logarithmic factor. Interestingly, this rate is tighter than the rate of contraction of the full posterior, but the price we pay for the rate improvement is that the construction of the sieve depends on the smoothness level α and the rate is non-adaptive.

Theorem 2. *Assume that f_0 is in the α -Hölder function class $\mathfrak{C}^{\alpha,B}(\mathcal{X})$ with envelope B . Consider the sieve maximum likelihood estimator $\hat{f}_K(x)$ defined by*

$$\hat{f}_K(\mathbf{x}) = \arg \max_{f \in \mathcal{G}_K} \sum_{i=1}^n \log \phi_{\sigma_0}(y_i - f(\mathbf{x}_i)),$$

where

$$\mathcal{G}_K = \left\{ \sum_{\mathbf{k} \in [K]^p} \sum_{\mathbf{s}; |\mathbf{s}|=0}^m \xi_{\mathbf{k}\mathbf{s}} \psi_{\mathbf{k}\mathbf{s}}(\mathbf{x}) : Kh \in [\underline{h}, \bar{h}], \boldsymbol{\mu}_{\mathbf{k}} \in \overline{\mathcal{X}_K(\mathbf{k})}, \max_{\mathbf{k} \in [K]^p, |\mathbf{s}|=0, \dots, m} |\xi_{\mathbf{k}\mathbf{s}}| \leq B \right\},$$

and $m \geq \alpha$. If $K_n = \lceil (n/\log n)^{1/(2\alpha+p)} \rceil$, then

$$\lim_{n \rightarrow \infty} \mathbb{P}_0 \left(\left\| f_0 - \hat{f}_{K_n} \right\|_{L_2(\mathbb{P}_{\mathbf{x}})} \geq M(\log n/n)^{\alpha/(2\alpha+p)} \right) = 0$$

for some large constant $M > 0$.

4 Application to partial linear models

The kernel mixture of polynomials model can be extended to other contexts thanks to its flexible nonparametric nature. In this section we present a semiparametric application: we use the kernel mixture of polynomials to model the nonparametric component in partial linear models. The partial linear model is of the form $y_i = \mathbf{z}_i^T \boldsymbol{\beta} + \eta(\mathbf{x}_i) + e_i$, where $\mathbf{z}_i, \mathbf{x}_i$'s are design points, $\boldsymbol{\beta}$ is the linear coefficient, $\eta : \mathcal{X} \rightarrow \mathbb{R}$ is an unknown nonparametric function, and e_i 's are independent $N(0, 1)$ residuals. In many applications, the estimation of the nonparametric component η is of great interest. For example, in Xu et al. (2016b), the parametric term $\mathbf{z}_i^T \boldsymbol{\beta}$ models the baseline disease progression and the nonparametric term $\eta(\mathbf{x}_i)$ models the individual-specific treatment effect deviations over time. When the regression coefficient $\boldsymbol{\beta}$ is of more interest, the estimation of η can still be critical since it could affect the inference of $\boldsymbol{\beta}$. As will be seen later, we prove the convergence results for both the nonparametric component η (Theorem 3) and the parametric component $\boldsymbol{\beta}$ (Theorem 4). Furthermore, as a consequence of the metric entropy result (Proposition 1), we obtain the Bernstein-von Mises limit of the marginal posterior distribution of $\boldsymbol{\beta}$.

4.1 Setup and prior specification

Let $\mathcal{X} = [0, 1]^p \subset \mathbb{R}^p$ be the design space of the nonparametric component, $\mathcal{Z} \subset \mathbb{R}^q$ be the design space of the parametric component, and $p(\mathbf{x}, \mathbf{z}) : \mathcal{X} \times \mathcal{Z} \rightarrow (0, \infty)$ be a continuous density function supported on $\mathcal{X} \times \mathcal{Z}$. We incorporate the partial linear model with the kernel mixture of polynomials prior for the nonparametric component η through $\mathcal{P} = \{p_{\boldsymbol{\beta}, \eta}(\mathbf{x}, \mathbf{z}, y) : \boldsymbol{\beta} \in \mathbb{R}^q, \eta \in \bigcup_{K=1}^{\infty} \mathcal{F}_K\}$, a class of distributions indexed by the linear coefficient $\boldsymbol{\beta}$ and the nonparametric component η , where $p_{\boldsymbol{\beta}, \eta}(\mathbf{x}, \mathbf{z}, y) = \phi(y - \mathbf{z}^T \boldsymbol{\beta} - \eta(\mathbf{x})) p_{\mathbf{x}, \mathbf{z}}(\mathbf{x}, \mathbf{z})$ and $\mathcal{F}_K = \bigcup_{Kh \in [\underline{h}, \bar{h}]} \mathcal{F}_K(h)$ with $\mathcal{F}_K(h)$ given by (4). We assume that the data $\mathcal{D}_n = (\mathbf{x}_i, \mathbf{z}_i, y_i)_{i=1}^n$ are independently sampled from $p_0(\mathbf{x}, \mathbf{z}, y) = \phi(y - \mathbf{z}^T \boldsymbol{\beta}_0 - \eta_0(\mathbf{x})) p_{\mathbf{x}, \mathbf{z}}(\mathbf{x}, \mathbf{z})$ for some $\boldsymbol{\beta}_0 \in \mathbb{R}^q$ and some function $\eta_0 \in \mathfrak{C}^{\alpha, L}(\mathcal{X})$. Several additional assumptions regarding the parametric component $\mathbf{z}^T \boldsymbol{\beta}$ are needed: The design space $\mathcal{Z} \subset \mathbb{R}^q$ for \mathbf{z} is compact with $\sup_{\mathbf{z} \in \mathcal{Z}} \|\mathbf{z}\|_1 \leq \bar{B}$ for some $\bar{B} > 0$; The sampling distribution for \mathbf{z} satisfies $\mathbb{E}\mathbf{z} = \mathbf{0}$ and $\mathbb{E}\mathbf{z}\mathbf{z}^T$ being non-singular; The density

of the design points $(\mathbf{x}_i, \mathbf{z}_i)_{i=1}^n$ factorizes as $p_{(\mathbf{x}, \mathbf{z})}(\mathbf{x}, \mathbf{z}) = p_{\mathbf{x}}(\mathbf{x})p_{\mathbf{z}}(\mathbf{z})$, i.e., \mathbf{x} and \mathbf{z} are independent.

For the prior specification, we assume η follows the kernel mixture of polynomials prior Π_η constructed in Section 2.2 with $\sigma = 1$. For the parametric component β , we impose a standard Gaussian prior $\Pi_\beta = N(0, \mathbf{I}_q)$, independent of Π_η . The joint prior is denoted by $\Pi = \Pi_\eta \times \Pi_\beta$.

4.2 Convergence results

We first tackle the convergence of the nonparametric component η . The following theorem not only addresses the rate of contraction of the marginal posterior of η , but also serves as one of the building blocks for proving the Bernstein-von Mises limit of the marginal posterior of β . The proofs of Theorem 3 and Theorem 4 are deferred to the Supplementary Material (Xie and Xu, 2019).

Theorem 3 (Nonparametric rate). *Assume that η_0 is in the α -Hölder function class $\mathfrak{C}^{\alpha, B}(\mathcal{X})$ with envelope B . Under the setup and prior specification in Section 4.1, with $m \geq \alpha$,*

$$\Pi(\|\eta - \eta_0\|_{L_2(\mathbb{P}_{\mathbf{x}})} > M\epsilon_n \mid \mathcal{D}_n) \rightarrow 0$$

in \mathbb{P}_0 -probability for some large constant $M > 0$, where $\epsilon_n = n^{-\alpha/(2\alpha+p)}(\log n)^{t/2}$, and $t > 2\alpha \max(r_0, 1)/(2\alpha + p) + \max(0, 1 - r_0)$.

Now we turn to the convergence results for the parametric component. The focus is the asymptotic normality of the marginal posterior distribution of β , i.e., the Bernstein-von Mises limit (Doob, 1949). To achieve this, we need the notion of the *least favorable submodel* for semiparametric models (Bickel et al., 1998). For each fixed $\beta \in \mathbb{R}^q$, the least favorable curve η_β^* is defined by the minimizer of the Kullback-Leibler divergence over all η : $\eta_\beta^*(x) = \arg \inf_{\eta \in \mathcal{F}} D_{\text{KL}}(p_0 \| p_{\beta, \eta})$. Under the assumptions $\mathbb{E}\mathbf{z} = \mathbf{0}$ and $p_{(\mathbf{x}, \mathbf{z})}(\mathbf{x}, \mathbf{z}) = p_{\mathbf{x}}(\mathbf{x})p_{\mathbf{z}}(\mathbf{z})$, for each β , it can be shown that $\eta_\beta^*(\mathbf{x})$ coincides with $\eta_0(\mathbf{x})$. The least favorable submodel is defined to be $\{p_{\beta, \eta_\beta^*} : \beta \in \mathbb{R}^q\}$, which in turn coincides with $\{p_{\beta, \eta_0} : \beta \in \mathbb{R}^q\}$ in our context.

Theorem 4. *Assume that η_0 is in the α -Hölder function class $\mathfrak{C}^{\alpha, B}(\mathcal{X})$ with envelope B . Under the setup and prior specification in Section 4.1, if $\alpha > p/2$ and $m \geq \alpha$, then*

$$\sup_F |\Pi(\sqrt{n}(\beta - \beta_0) \in F \mid \mathcal{D}_n) - \Phi(F \mid \Delta_n, (E\mathbf{z}\mathbf{z}^T)^{-1})| \rightarrow 0$$

in \mathbb{P}_0 -probability, where $\Phi(\cdot \mid \Delta_n, (E\mathbf{z}\mathbf{z}^T)^{-1})$ is the $N(\Delta_n, (E\mathbf{z}\mathbf{z}^T)^{-1})$ probability measure and

$$\Delta_n = \frac{1}{n} \sum_{i=1}^n (E\mathbf{z}\mathbf{z}^T)^{-1} \mathbf{z} [y_i - \eta_0(\mathbf{x}_i) - \mathbf{z}_i^T \beta_0].$$

The proof of Theorem 4 is based on verifying a set of sufficient conditions in Yang et al. (2015), which are provided in the Supplementary Material (Xie and Xu, 2019). However, we remark that the metric entropy results (Proposition 1) in Section 3 and the rate of contraction for η (Theorem 3) are also of fundamental interest in the verification process.

5 Numerical studies

We perform numerical studies to evaluate the kernel mixture of polynomials for non-parametric regression and the partial linear model. The posterior inference for all examples is carried out by a Markov chain Monte Carlo (MCMC) sampler with the number of burn-in iterations being 1000. Then we collect 1000 post-burn-in MCMC samples for posterior analysis. To determine K , we collect posterior samples for each fixed $K \in \{K_{\min}, \dots, K_{\max}\}$ and find the optimal K by minimizing the deviance information criterion (Gelman et al., 2014) over K . Numerical evidence shows that all Markov chains converge within 1000 iterations. The kernel we use throughout is the bump kernel $\varphi(x) = \exp[-1/(1 - x^2)]\mathbf{1}(|x| < 1)$.

5.1 A synthetic example for nonparametric regression

We first consider a synthetic example for nonparametric regression. Following Knapik et al. (2011) and Yoo et al. (2016), we consider the true function to be $f_0(x) = \sqrt{2} \sum_{s=1}^{\infty} s^{-3/2} \sin(s) \cos[(s - 1/2)\pi x]$, which has smoothness level $\alpha = 1$. We generate $n = 1000$ observations $(x_i, y_i)_{i=1}^n$ given by $y_i = f_0(x_i) + e_i$, where the design points $(x_i)_{i=1}^n$ are independently and uniformly sampled over $\mathcal{X} = [0, 1]$, and $(e_i)_{i=1}^n$ are independent $N(0, 0.2^2)$ residuals. For the prior specification, we let $\pi_\mu = \text{Unif}(-1, 1)$, $\pi_\beta = N(0, 10^2)$, $\pi_\xi = N(0, 10^2)\mathbf{1}(|\xi| \leq 50)$, and $Kh \sim \text{Unif}(\underline{h}, \bar{h})$. The hyperparameters are set as $\underline{h} = 1.2$, $\bar{h} = 2$, $B = 50$, $K = 15$, and $m = 2$. The range of K is set to be $\{6, 7, \dots, 15\}$.

For comparison we consider 3 competitors for estimating f_0 : the local polynomial regression (Fan and Gijbels, 1996), implemented in the `locpol` package (Cabrera, 2012), DiceKriging method (Roustant et al., 2012), and the robust Gaussian stochastic process emulation (RobustGaSP, Gu et al., 2017), implemented in the `RobustGaSP` package (Gu et al., 2016). The point-wise posterior means and 95%-credible/confidence intervals for $f(x)$ using the 4 nonparametric regression methods are plotted in Figure 2, respectively. We also compute the mean-squared errors of the posterior means of the four methods, where the ground true f_0 is evaluated at 1000 equidistant design points. In terms of accuracy measured by the mean-squared errors (marked in the bottom-left corner of each panel), the kernel mixture of polynomials performs better than DiceKriging and similarly to the local polynomials and RobustGaSP.

We further investigate the uncertainty quantification of the kernel mixture of polynomials and the 3 competitors by performing 1000 repeated experiments. Namely, we simulate 1000 synthetic datasets using the same setup as above, perform the corresponding inference for each dataset, and then compute the average coverage and average lengths of the point-wise credible/confidence band. We are particularly interested in the behavior of coverage of the point-wise credible/confidence band for $x \in [0.30, 0.35]$, where the true regression function f_0 exhibits a bump. As shown in the left panel of Figure 3, DiceKriging and the local polynomial regression provide the best coverage, but the widths of the point-wise confidence band are much larger (see the right panel of

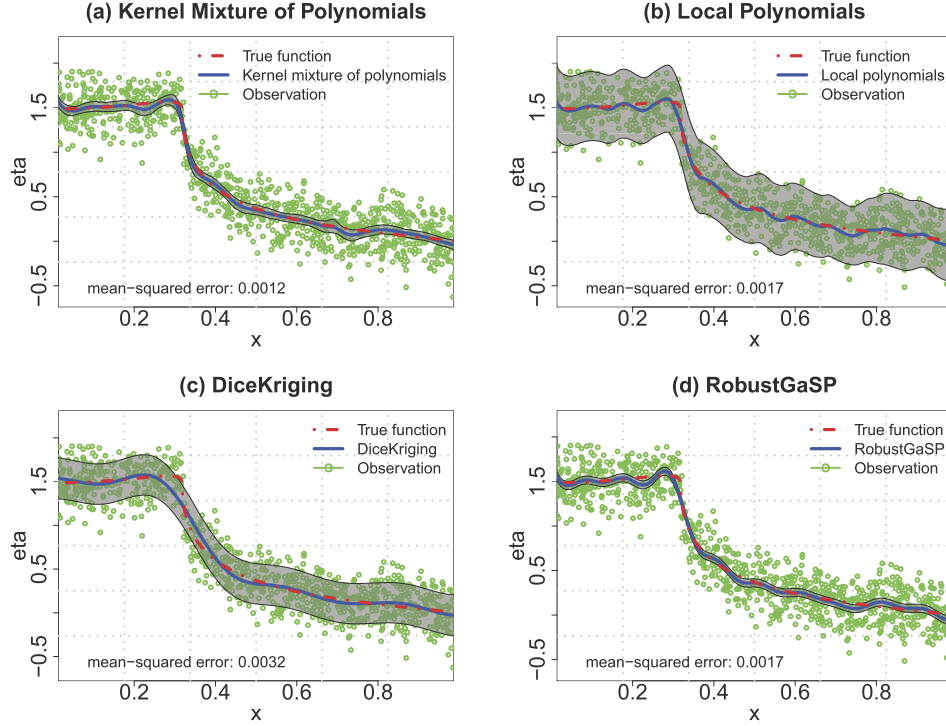


Figure 2: Synthetic example for nonparametric regression. Shaded regions are the point-wise 95%-credible/confidence intervals. The solid lines are point-wise posterior means/point estimators of f and the dot-dashed lines are the ground true f_0 . Scatter points are the observations.

Figure 3). RobustGaSP provides the worst coverage among the 4 approaches, and the corresponding confidence band is still wider than that of the kernel mixture of polynomials. In contrast, the proposed kernel mixture of polynomial regression provides better credible band coverage for $x \in [0.30, 0.35]$ than RobustGaSP. In addition, the average length of the point-wise credible band of the kernel mixture of polynomials is narrower than those of the other 3 competitors. However, we also notice that this feature is at the cost of worse coverage than the local polynomial regression and DiceKriging. Namely, the narrower point-wise confidence band is over-confident near the bump of f_0 than DiceKriging and the local polynomial regression. One potential reason for this phenomenon could be that the model over-smooths the true regression function f_0 near the bump based on a random sample of limited size, but f_0 is much smoother elsewhere (Yoo et al., 2016).

As pointed out by one of the referees, for the kernel mixture of polynomials, the bad coverage behavior of the point-wise credible band near the bump of f_0 may be alleviated by using the credible set for the entire function. Following Szabó et al. (2015),

we consider the following L_2 -credible set for f :

$$\hat{C}_n(\gamma_n) = \{f : \|f - f_0\|_{L_2(\mathbb{P}_x)} \leq \gamma_n\},$$

where $\gamma_n > 0$ is the credible radius such that $\Pi(f \in \hat{C}_n(\gamma_n) \mid \mathcal{D}_n) = 95\%$. Numerically, the credible set $C_n(\gamma_n)$ can be computed as follows: let $(f_t^{\text{mc}})_{t=1}^{n_{\text{mc}}}$ be n_{mc} posterior samples of f drawn from the MCMC. Let γ_n be the 95%-quantile of $(\|f_t^{\text{mc}} - \hat{f}\|_{L_2(\mathbb{P}_x)})_{t=1}^{n_{\text{mc}}}$, where \hat{f} is the point-wise posterior mean function. Then the L_2 -credible set $\hat{C}_n(\gamma_n)$ can be constructed by

$$\left\{ f : \min_{t: \|f_t^{\text{mc}} - \hat{f}\|_{L_2(\mathbb{P}_x)} \leq \gamma_n} f_t^{\text{mc}}(\mathbf{x}) \leq f(\mathbf{x}) \leq \max_{t: \|f_t^{\text{mc}} - \hat{f}\|_{L_2(\mathbb{P}_x)} \leq \gamma_n} f_t^{\text{mc}}(\mathbf{x}) \text{ for all } \mathbf{x} \in \mathcal{X} \right\}.$$

Using the aforementioned 1000 replicates of synthetic datasets, we examine the corresponding point-wise coverage behavior of the L_2 -credible sets of the entire function locally around the bump ($x \in [0.25, 0.4]$), compute the average lengths of these credible sets, and plot them in Figure 3. The L_2 -credible sets exhibit better point-wise coverage behavior than the point-wise credible bands at the cost of slightly wider lengths. This is also in accordance with the observations in Szabó et al. (2015): the overall coverage is satisfactory, but the coverage near the bump is slightly lower. We also observe that the lengths of these L_2 -credible sets are close to those obtained from RobustGaSP, but the coverage is better.

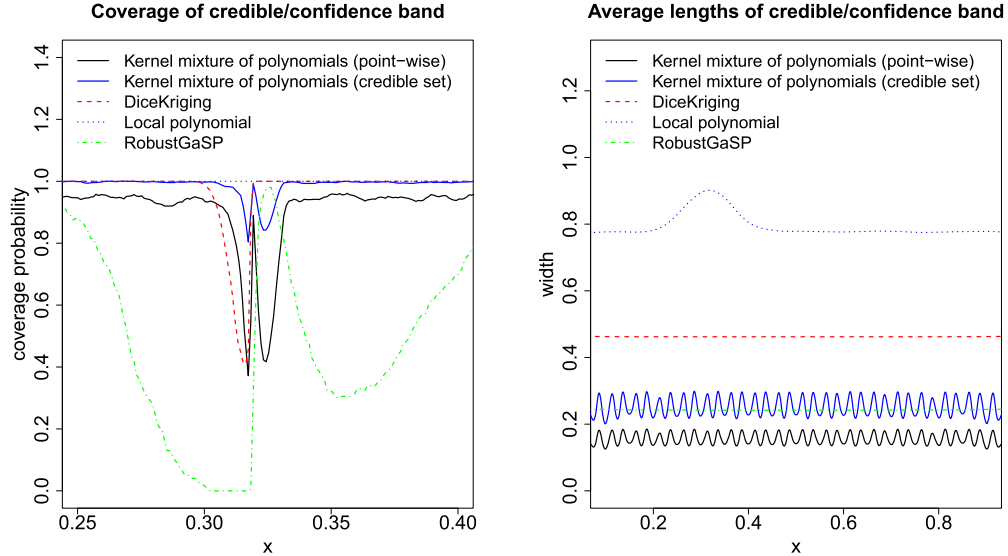


Figure 3: Coverage of the credible band/confidence band/credible set and the average lengths of the credible band/confidence band/credible set for nonparametric regression under 1000 repeated experiments.

5.2 Computation comparison with Gaussian processes

One motivation of developing the kernel mixture of polynomials for Bayesian non-parametric regression is to tackle the computational bottleneck of classical Gaussian processes. In this section we compare the computational costs of the kernel mixture of polynomials against Gaussian processes with 3 different covariance functions: the squared-exponential covariance function:

$$K(x, x'; \psi) = \exp \left[-(x - x')^2 / \psi^2 \right],$$

the Matérn covariance function with roughness parameter $3/2$:

$$K(x, x'; \psi) = \left(1 + \sqrt{3}|x - x'|/\psi \right) \exp \left(-\sqrt{3}|x - x'|/\psi \right),$$

and the Matérn covariance function with roughness parameter $5/2$:

$$K(x, x'; \psi) = \left(1 + \sqrt{5}|x - x'|/\psi + 5|x - x'|^2/(3\psi^2) \right) \exp \left(-\sqrt{5}|x - x'|/\psi \right).$$

The parameter ψ appearing in these covariance functions is referred to as the range parameter. Here we follow the suggestion by van der Vaart and van Zanten (2009) and impose an inverse-Gamma hyperprior on ψ , *i.e.*, $\pi(\psi) \propto \psi^{-a_\psi-1} \exp(-b_\psi/\psi)$ for some $a_\psi, b_\psi \geq 2$. Such a hierarchical formulation for Gaussian processes with a hyperprior on the range parameter is also referred to as rescaled Gaussian processes in the literature (van der Vaart and van Zanten, 2007).

We adopt the same setup as in Subsection 5.1 except that $e_i \sim N(0, 0.1^2)$, and for each rescaled Gaussian process described above, posterior inference is carried out using classical Markov chain Monte Carlo, with 1000 posterior samples collected after burn-in. The hyperparameter for π_ψ is set to be $a_\psi = b_\psi = 2$. The posterior means and point-wise 95% credible intervals for the kernel mixture of polynomials and 3 rescaled Gaussian processes are visualized in Figure 4 (a), (b), (c), and (d), respectively. Furthermore, we examine the mean-squared errors of the posterior means as well as the computational costs of the 4 approaches. As shown in Table 1, the proposed kernel mixture of polynomials not only yields a point estimate with smaller mean-squared error, but is also much more efficient than the rescaled Gaussian processes in terms of runtime expenses. The main reason is that when updating the range parameter ψ in a single iteration of MCMC using Gaussian processes, the inverse of the covariance matrix needs to be re-computed, whereas the kernel mixture of polynomials does not suffer from such a computational bottleneck.

5.3 A synthetic example for the partial linear model

We consider a synthetic example to evaluate the performance of the kernel mixture of polynomials for the partial linear model. We simulate $n = 500$ observations according to the model $y_i = \mathbf{z}_i^T \beta_0 + \eta_0(x_i) + e_i$, where β_0 is provided in Table 2, $(e_i)_{i=1}^n$ are independent $N(0, 1)$ residuals that are independent of $(x_i, \mathbf{z}_i)_{i=1}^n$, and $\eta_0(x) =$

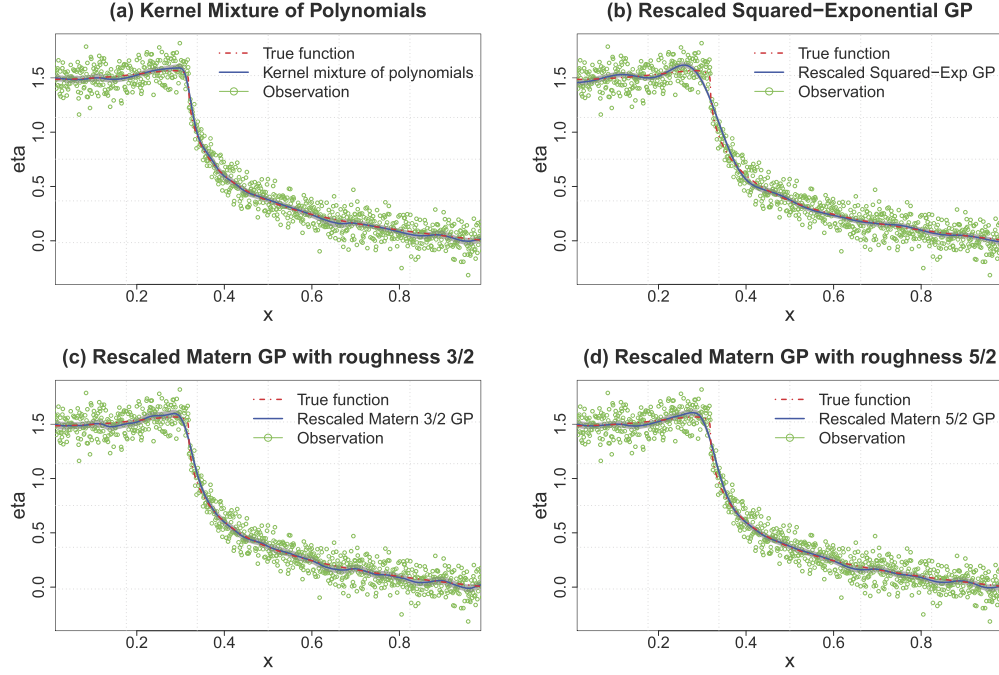


Figure 4: Comparison with rescaled Gaussian processes: Shaded regions are point-wise 95%-credible/confidence intervals. The solid lines are point-wise posterior means/point estimators of f and the dot-dashed lines are the ground true f_0 . Scatter points are the observations.

Table 1: Comparison with Gaussian processes: Mean-squared errors and computational costs.

Method	mean-squared error	computational cost
Kernel mixture of polynomials	3.447×10^{-4}	562s
Squared-exponential Gaussian process	1.266×10^{-3}	198806s
Matérn 3/2 Gaussian process	5.736×10^{-4}	24219s
Matérn 5/2 Gaussian process	6.918×10^{-4}	20300s

$2.5 \exp(-x) \sin(10\pi x)$. The nonparametric function η_0 is highly nonlinear and hence brings natural challenge to estimation. The design points $(\mathbf{z}_i)_{i=1}^n$ for the linear component follow $\text{Unif}([-1, 1]^8)$ independently, and the design points $(x_i)_{i=1}^n$ for η are independently sampled from $\text{Unif}(0, 1)$. The hyperparameters for the kernel mixture of polynomials prior are set as follows: $\underline{h} = 1.2$, $\bar{h} = 2$, $B = 100$, and $m = 3$. For the prior specification, we assume $\pi_\mu = \text{Unif}(-1, 1)$, $\pi_\beta = \mathcal{N}(\mathbf{0}, 10^2 \mathbf{I}_8)$, $\pi_\xi = \mathcal{N}(0, 10^2) \mathbf{1}(|\xi| \leq 100)$, and $Kh \sim \text{Unif}(\underline{h}, \bar{h})$. We set the range of K to be $\{6, 7, \dots, 15\}$.

For the parametric component, we compute the posterior means and the posterior 95% credible intervals for β . For comparison, we calculate the least-squared estimate of

β : $\hat{\beta}_{LS} = \arg \min_{\beta \in \mathbb{R}^q} \sum_{i=1}^n (y_i - \mathbf{z}_i^T \beta)^2$. The comparison is provided in Table 2, where the column “KMP” stands for the posterior means of β under the kernel mixture of polynomials prior. From the posterior summary of β we see that the underlying true β_0 lie in the posterior 95%-credible intervals of the kernel mixture of polynomials prior, and the corresponding posterior means outperform the least-squared estimate in terms of accuracy. For the nonparametric component, we plot the point-wise posterior means

Table 2: Simulation example: Inference of β with $\eta_0(x) = 2.4 \exp(-x) \sin(10\pi x)$.

β	β_0	KMP	95%-credible intervals	$\hat{\beta}_{LS}$
β_1	1.0338	1.0579	(0.8949, 1.2150)	1.1976
β_2	0.1346	0.1003	(-0.0560, 0.2516)	0.1733
β_3	0.2854	0.3481	(0.1832, 0.5090)	0.4427
β_4	0.6675	0.6449	(0.5007, 0.7887)	0.8386
β_5	0.6732	0.7212	(0.5630, 0.8838)	0.7427
β_6	0.5293	0.5433	(0.3971, 0.6866)	0.6274
β_7	-0.5073	-0.4759	(-0.6337, -0.3158)	-0.3464
β_8	-3.3942	-3.3031	(-3.4492, -3.1450)	-3.5253

of η along with the point-wise 95% credible intervals in Figure 5 (a). The mean-squared error of the posterior means at 500 equidistant points on $(0, 1)$ is 0.0351. Based on the least-squared estimate of β , we also consider the local polynomial regression, DiceKriging method, and RobustGaSP as three alternatives to estimate the nonparametric component. The numerical comparisons are illustrated in Panels (b), (c), and (d) in Figure 5, respectively. The kernel mixture of polynomials outperforms the local polynomial in terms of both the mean-squared error. Both DiceKriging and RobustGaSP fail to detect the nonlinearity of η_0 , giving rise to significantly larger mean-squared error of the point-wise posterior means.

To examine the uncertainty quantification of the 4 approaches, we perform 1000 repeated experiments. The coverage and average lengths of the point-wise credible/confidence band are presented in Figure 6. The averages of the regression function estimates and their point-wise credible/confidence bands based on 1000 repeated experiments are computed and visualized in Figure 7. It can be seen that the kernel mixture of polynomials provides decent coverage to η_0 and the corresponding average lengths of the point-wise credible band are smaller than those using the other three approaches. In contrast, all the three competitors fail to detect the signal of η_0 after averaging from repeated experiments, producing inappropriate point-wise confidence band (either too narrow or too wide). In particular, DiceKriging provides poor coverage of the point-wise confidence band, and the average widths of the point-wise confidence bands using the local polynomials and RobustGaSP are much wider than the magnitude of η_0 .

5.4 The partial linear model for the wage data

We further analyze the cross-sectional data on wages (Wooldridge, 2015), a benchmark dataset for the partial linear model. This dataset is also available in the `np` package

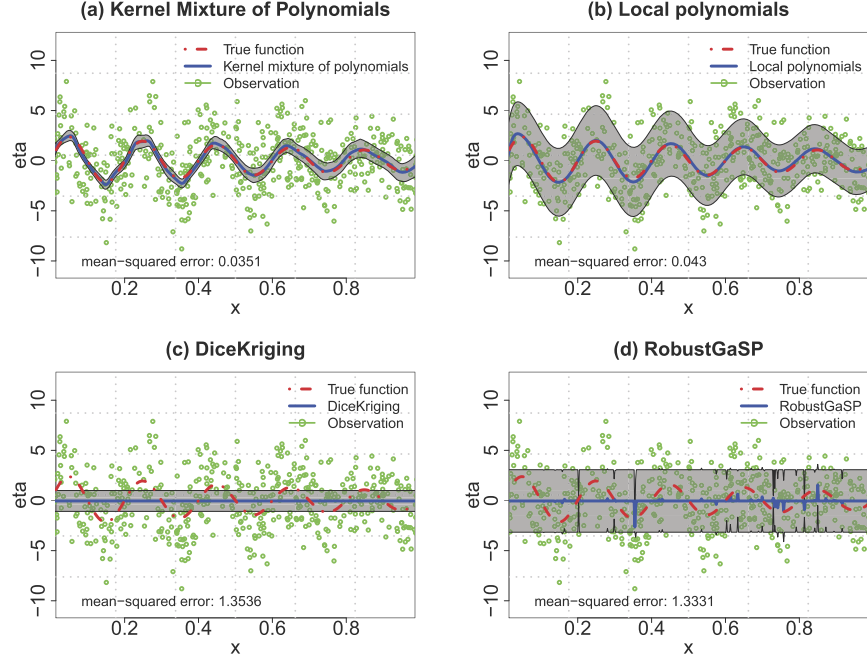


Figure 5: Synthetic example for the partial linear model. Shaded regions are point-wise 95% credible/confidence intervals. The solid lines are point-wise posterior means/point estimators of η and the dot-dashed lines are η_0 . Scatter points are the observations. Panels (b), (c), and (d) are computed using the discrepancy data $(x_i, y_i - \mathbf{z}_i^T \hat{\beta}_{LS})_{i=1}^n$, where $\hat{\beta}_{LS}$ is the least-squared estimate of β .

(Hayfield et al., 2008). It consists of 526 observations with 24 variables and are taken from U.S. Current Population Survey for the year 1976. In particular, we are interested in modeling the hourly wage on the logarithm scale as the response with respect to 5 variables: years of education (educ), years of potential experience (exper), years with current employer, gender, and marital status. Choi and Woo (2015) and Hayfield et al. (2008) suggest the following form of the model:

$$y_i = \beta_1 z_i^{\text{female}} + \beta_2 z_i^{\text{married}} + \beta_3 z_i^{\text{educ}} + \beta_4 z_i^{\text{tenure}} + \eta(x_i^{\text{exper}}) + e_i,$$

where e_i 's are independent $N(0, \sigma^2)$ residuals. The z_i^{female} are ± 1 -valued, where $z_i^{\text{female}} = 1$ indicates that the i th observation is a female, and -1 otherwise. We set $z_i^{\text{married}} = 1$ if the i th observation is married, and -1 otherwise. We centralize z_i^{educ} and z_i^{tenure} before applying the partial linear model, i.e., $\sum_{i=1}^n z_i^{\text{educ}} = \sum_{i=1}^n z_i^{\text{tenure}} = 0$. The x_i^{exper} 's are re-scaled so that they lie in $(0, 1)$. To evaluate the performance of the proposed method, we use 300 observations as the training data, with the rest of the 226 observations left as the testing data to compute the prediction mean-squared error. The prior specification and hyperparameters for the MCMC sampler are set as follows: $\underline{h} = 1.2$, $\bar{h} = 2$, $m = 3$,

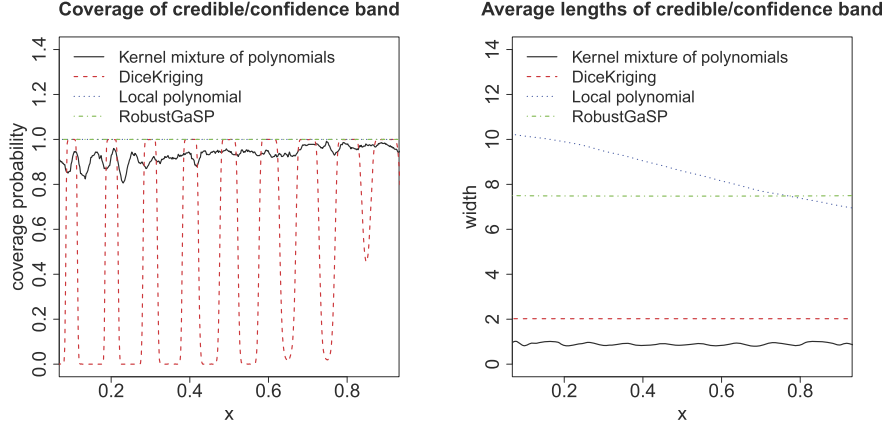


Figure 6: Coverage of point-wise credible/confidence band and average lengths of point-wise credible/confidence band for partial linear model under 1000 repeated experiments.

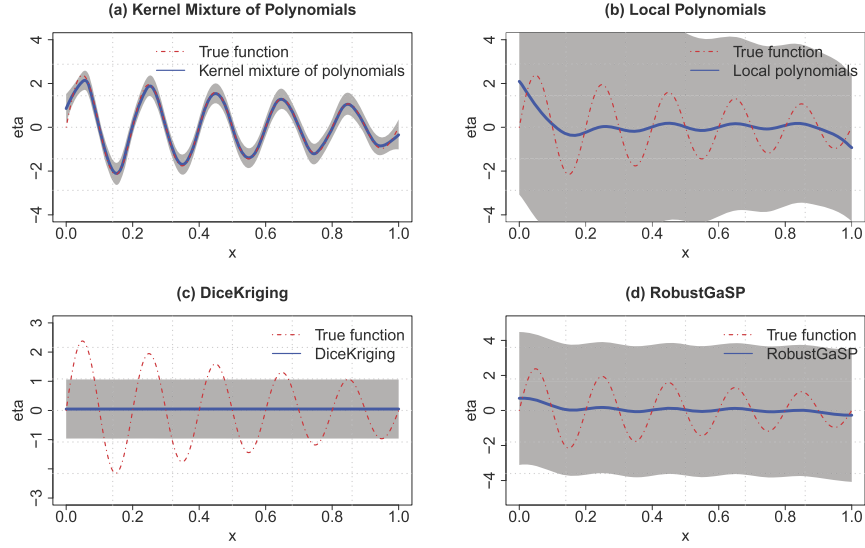


Figure 7: Averages of regression function estimates (blue solid lines) and the corresponding point-wise credible/confidence band (shaded regions) for partial linear model using 1000 repeated experiments.

$\pi(\sigma^2) \propto [\sigma^2]^{-2} \exp(-1/\sigma^2)$ (i.e., the inverse-Gamma density), $\pi_{\beta} = N(\mathbf{0}, 10^2 \mathbf{I}_4)$, and $\pi_{\xi} = N(0, 10^2) \mathbf{1}(|\xi| \leq 100)$. The range of K is set to be $\{11, 12, \dots, 20\}$.

We calculate the posterior means and the posterior 95%-credible intervals for β . For comparison, we also provide the least-squared estimate of β and the estimate computed

by the **np** package (Hayfield et al., 2008). The results are summarized in Table 3 (the “KMP” column represents the posterior means of β), showing that the kernel mixture of polynomials estimate is closer to the **np** package estimate than the least-squared estimate, and all 3 point estimates of β lie in the posterior 95%-credible intervals. For

Table 3: Wage data example: Inference of β .

	KMP	95% credible intervals	np package	$\hat{\beta}_{LS}$
female	-0.1214	(-0.2534, -0.0071)	-0.1287	-0.0921
married	0.0249	(-0.0943, 0.1629)	0.0279	0.3209
educ	0.0903	(0.0373, 0.1404)	0.0891	0.1257
tenure	0.0175	(-0.0032, 0.0372)	0.0167	0.0152

the nonparametric component, we compute the kernel mixture of polynomials prediction on the testing dataset. The comparison with the true testing responses is demonstrated in Panel (a) of Figure 8. The 3 alternatives based on the **np** estimate of β for estimating the nonparametric component η are: the local polynomial regression, DiceKriging, and RobustGaSP. The performance of these 3 competitors are visualized in Panels (b), (c), and (d) in Figure 8, respectively. The local polynomial regression estimate does not outperform the kernel mixture of polynomials in terms of the prediction mean-squared error, and the prediction curve is highly non-smooth. DiceKriging does not work in this scenario: the prediction mean-squared error is large, the prediction curve is highly non-smooth, and the point-wise confidence intervals show singularity in estimating the covariance matrix. RobustGaSP, though gives similar prediction mean-squared error compared to the kernel mixture of polynomials, does not capture the local nonlinearity of the nonparametric component. In addition, the point-wise 95% confidence/credible intervals for the local polynomial regression and RobustGaSP are wider than those given by the kernel mixture of polynomials when $x \in (0, 0.6)$. Since the data are dense in the region $(0, 0.6)$, the point-wise credible intervals estimated by the kernel mixture of polynomials are thinner in this region; In contrast, the design points are sparser in the region $x \geq 0.6$, and correspondingly there exists larger uncertainty in estimating η . Namely, the uncertainty of the kernel mixture of polynomials is adaptive to the distribution of the design points.

6 Discussion

In this work we assume that an upper bound B for the supremum norm of the derivatives of the true regression function f_0 is known and is used to construct the prior for the coefficients $(\xi_{\mathbf{k}s} : \mathbf{k} \in [K]^p, |s| = 0, \dots, m)$. In addition, the prior for the standard deviation σ of the residuals is supported on a compact interval $[\underline{\sigma}, \bar{\sigma}]$ containing σ_0 . These two restrictions can be potentially inconvenient, as B , $\underline{\sigma}$, and $\bar{\sigma}$ are typically unknown in practice. Here we provide the following non-adaptive alternative solution. The prior specification is slightly modified as follows: First K is fixed at $K_n = \lceil (n/\log n)^{1/(2\alpha+p)} \rceil$, h is fixed at $2/K$, and $\mu_{\mathbf{k}}$ is fixed at $\mu_{\mathbf{k}}^*$ for all $\mathbf{k} \in [K_n]^p$. Namely, we consider a simplified

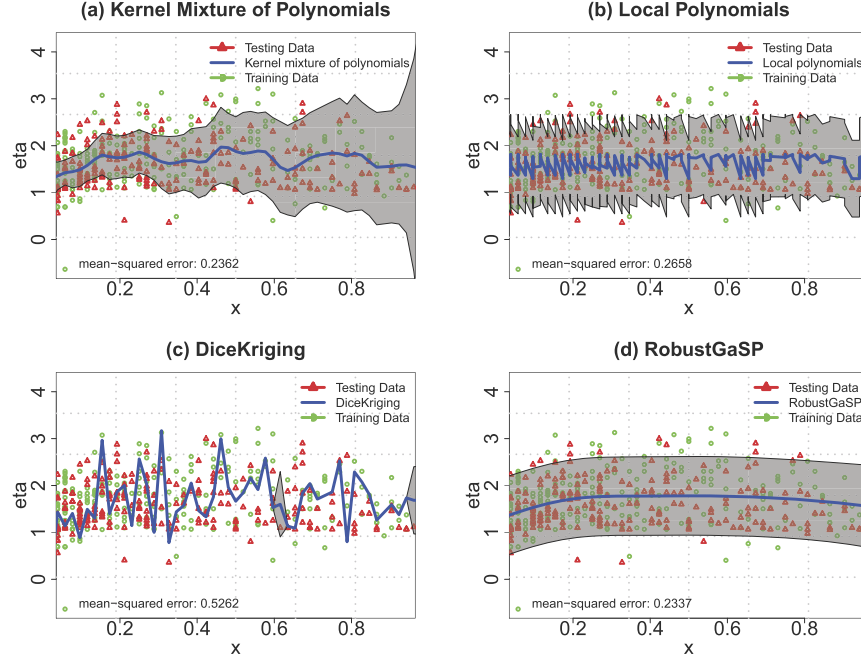


Figure 8: Wage data example. Shaded regions are point-wise 95%-credible/confidence intervals. The solid lines are point-wise posterior means/point estimators of η . Circle scatter points are the training observations, and triangle scatter points are testing responses. Panels (b), (c), and (d) are computed using the discrepancy data $(x_i, y_i - \mathbf{z}_i^T \hat{\boldsymbol{\beta}}_{np})_{i=1}^n$, where $\hat{\boldsymbol{\beta}}_{np}$ is the estimate of $\boldsymbol{\beta}$ computed using the np package.

kernel mixture of polynomials:

$$f(\mathbf{x}) = \sum_{\mathbf{k} \in [K_n]^p} \sum_{\mathbf{s}: |\mathbf{s}|=0}^m \xi_{\mathbf{k}\mathbf{s}} \psi_{\mathbf{k}\mathbf{s}}(\mathbf{x}), \quad \text{with} \quad \psi_{\mathbf{k}\mathbf{s}}(\mathbf{x}) = \frac{\varphi_h(\mathbf{x} - \boldsymbol{\mu}_{\mathbf{k}}^*)(\mathbf{x} - \boldsymbol{\mu}_{\mathbf{k}}^*)^{\mathbf{s}}}{\sum_{\mathbf{l} \in [K_n]^p} \varphi_h(\mathbf{x} - \boldsymbol{\mu}_{\mathbf{l}}^*)}. \quad (8)$$

Then the prior Π^n on (f, σ) is imposed as follows: $\xi_{\mathbf{k}\mathbf{s}} \mid \sigma \sim N(0, n^2 \sigma^2)$ independently for all $\mathbf{k} \in [K_n]^p$ and $0 \leq |\mathbf{s}| \leq m$ given σ . Note here we relax the assumption that $\xi_{\mathbf{k}\mathbf{s}}$ is supported on $[-B, B]$. Denote the corresponding induced conditional prior on f given σ by $\Pi_f^n(\cdot \mid \sigma)$. We impose σ^2 with an inverse-Gamma prior $\Pi_\sigma = \text{IG}(a_\sigma, b_\sigma)$, with $\pi_{\sigma^2}(\sigma^2) \propto (\sigma^2)^{-a_\sigma/2-1} \exp[-b_\sigma/(2\sigma^2)]$ for some $a_\sigma, b_\sigma \geq 2$. In particular, we impose the prior on $\xi_{\mathbf{k}\mathbf{s}}$ with a scaling factor n so that for large n the prior is weakly informative. Then under the above setup and prior specification $\Pi^n(\cdot, \cdot) = \Pi_f^n(\cdot \mid \sigma) \times \Pi_\sigma(\cdot)$, the posterior contraction with respect to $L_2(\mathbb{P}_n)$ can be shown to be minimax-optimal up to a logarithmic factor (see Section F in the Supplementary Material, Xie and Xu, 2019).

In general, in developing contraction rates for nonparametric regression with respect to the $\|\cdot\|_{L_2(\mathbb{P}_x)}$ distance, it requires significant work to drop the boundedness assump-

tion on the space of regression functions. The underlying reason is that when the space of functions are not uniformly bounded, existence of certain test functions satisfying the so-called Ghosal & van der Vaart conditions originally presented in Ghosal et al. (2000) may not exist. To the best of our knowledge, Xie et al. (2017) is the only literature addressing the posterior contraction of Bayesian nonparametric regression with respect to $\|\cdot\|_{L_2(\mathbb{P}_\mathbf{x})}$ in a systematic framework without imposing uniform boundedness on function spaces. It will be interesting to extend the techniques there to study the posterior contraction of the kernel mixture of polynomials without requiring the boundedness assumption.

Another feasible extension is variable selection in high dimensions. To be more specific, suppose that the design space $\mathcal{X} = [0, 1]^p$ is with high dimensionality in the sense that $p \gg n$, and the true regression function f_0 only depends on a subset of the covariates $\{x_{j_1}, \dots, x_{j_q}\} \subset \{x_1, \dots, x_p\}$, where $q < n$ is the intrinsic dimension. To tackle such a high-dimensional variable selection problem in nonparametric regression, we modify the kernel mixture of polynomials as follows: Let the kernel function $\varphi(x_1, \dots, x_p) : \mathbb{R}^p \rightarrow [0, 1]$ be of the product form $\varphi(\mathbf{x}) = \varphi(x_1, \dots, x_p) = \prod_{j=1}^p \varphi^1(x_j)$, where $\varphi^1 : \mathbb{R} \rightarrow [0, 1]$ is a univariate kernel function. We introduce the auxiliary binary variables $(z_j)_{j=1}^p$ to indicate whether the j th covariate is active or not, and modify the kernel mixture of polynomial system as follows:

$$\psi_{\mathbf{ks}}(x_1, \dots, x_p \mid z_1, \dots, z_p) = \frac{\prod_{j=1}^p [\varphi_h^1(x_j - \mu_{\mathbf{k}j})]^{z_j}}{\sum_{\mathbf{l} \in [K]^p} \prod_{j=1}^p [\varphi_h^1(x_j - \mu_{\mathbf{l}j})]^{z_j}} \prod_{j=1}^p (x_j - \mu_{\mathbf{k}j}^*)^{s_j z_j}.$$

Then by letting

$$f(x_1, \dots, x_p \mid z_1, \dots, z_p) = \sum_{\mathbf{k} \in [K]^p} \sum_{\mathbf{s}: |\mathbf{s}|=0}^m \xi_{\mathbf{ks}} \psi_{\mathbf{ks}}(x_1, \dots, x_p \mid z_1, \dots, z_p),$$

we can see that if $z_j = 0$, f does not depend on x_j . In fact, the basis function $\psi_{\mathbf{ks}}$ only depends on covariates that are active, *i.e.*, $(x_j : z_j = 1, j = 1, \dots, p)$. One can further impose independent Bernoulli(p) prior on (z_1, \dots, z_p) and complete the hierarchical Bayesian model by adopting the prior specification in Subsection 2.2 for the rest of the parameters. We believe that it will be interesting to investigate the convergence properties of such a modification of the kernel mixture of polynomials.

There are also several other potential extensions of the current work. Firstly, we develop the theoretical results under the assumption that the residuals $(e_i)_{i=1}^n$ are Gaussian. In cases where the residuals are only assumed to be sub-Gaussian, further exploration of the convergence properties can be investigated. Secondly, the design points are assumed to be random in the present paper. In cases where the design points are fixed, which is also a common phenomenon in many physical experiments (Tuo and Wu, 2015), theoretical results for the kernel mixture of polynomials can be further extended using the techniques developed for non-independent nor-identically distributed observations by Ghosal et al. (2007a) or fixed-design nonparametric regression by Xie et al. (2017). Secondly, we have only considered the case where the true regression function f_0 is in the α -Hölder function class. It is also interesting to extend the current framework to the

case where f_0 is in the α -Sobolev function space. Roughly speaking, when α is an integer, a function f is called α -Sobolev if the corresponding $(\alpha - 1)$ th derivatives are squared integrable. The almost-sure existence of derivatives of α -Sobolev functions guarantees that the approximation lemma (Lemma 1) is still applicable, and hence it is reasonable to expect that our theory also applies to an α -Sobolev f_0 . In addition, when applying the kernel mixture of polynomials to the partial linear model, we only consider the case where $\mathbb{E}\mathbf{z} = 0$ and \mathbf{x} is independent of \mathbf{z} , indicating that the linear component and the nonparametric component are orthogonal. On one hand, the idea of orthogonality has been explored in the literature of calibration of inexact computer models (Plumlee and Joseph, 2016; Plumlee, 2017), and therefore exploring the application of the kernel mixture of polynomials to calibration of orthogonal computer models is a promising extension. On the other hand, it is also interesting to investigate the convergence theory when the two components are not orthogonal to each other. We also expect that the kernel mixture of polynomials can be applied to other semiparametric models besides the partial linear model, such as the single-index model (Ichimura, 1993), the projection pursuit regression (Friedman and Stuetzle, 1981), etc. Finally, we have developed a theoretical support for a sieve maximum likelihood estimator with compact restrictions on the parameter spaces. In particular, the loss function is of the least-squared form. From the computational perspective, an efficient optimization technique that is also scalable to big data can be designed to obtain the frequentist estimator in light of the rich literature of solving nonlinear least-squared problems (Nocedal and Wright, 2006).

Supplementary Material

Supplementary Material for “Adaptive Bayesian Nonparametric Regression Using a Kernel Mixture of Polynomials with Application to Partial Linear Models” (DOI: [10.1214/19-BA1148SUPP](https://doi.org/10.1214/19-BA1148SUPP); .pdf). The supplementary material contains additional notations, proofs for Section 3, Section 4, posterior contraction for unknown σ^2 discussed in Section 6 and its proof, and cited theorems and results.

References

- Affandi, R. H., Fox, E., and Taskar, B. (2013). “Approximate inference in continuous determinantal processes.” In *Advances in Neural Information Processing Systems*, 1430–1438. 164
- Bhattacharya, A., Pati, D., and Dunson, D. (2014). “Anisotropic function estimation using multi-bandwidth Gaussian processes.” *Annals of Statistics*, 42(1): 352. MR3189489. doi: <https://doi.org/10.1214/13-AOS1192>. 160
- Bhattacharya, A., Pati, D., Pillai, N. S., and Dunson, D. B. (2015). “Dirichlet-Laplace priors for optimal shrinkage.” *Journal of the American Statistical Association*, 110(512): 1479–1490. MR3449048. doi: <https://doi.org/10.1080/01621459.2014.960967>. 166
- Bickel, P., Kleijn, B., et al. (2012). “The semiparametric Bernstein-von Mises theorem.”

- The Annals of Statistics*, 40(1): 206–237. MR3013185. doi: <https://doi.org/10.1214/11-AOS921>. 160
- Bickel, P. J., Klaassen, C. A., Ritov, Y., Wellner, J. A., et al. (1998). “Efficient and adaptive estimation for semiparametric models.” MR1623559. 169
- Cabrera, J. (2012). “locpol: Kernel local polynomial regression. R package version 0.4-0.” 170
- Castillo, I. and van der Vaart, A. (2012). “Needles and straw in a haystack: Posterior concentration for possibly sparse sequences.” *The Annals of Statistics*, 40(4): 2069–2101. MR3059077. doi: <https://doi.org/10.1214/12-AOS1029>. 166
- Celeux, G., Hurn, M., and Robert, C. P. (2000). “Computational and inferential difficulties with mixture posterior distributions.” *Journal of the American Statistical Association*, 95(451): 957–970. MR1804450. doi: <https://doi.org/10.2307/2669477>. 164
- Chen, H. et al. (1988). “Convergence rates for parametric components in a partly linear model.” *The Annals of Statistics*, 16(1): 136–146. MR0924861. doi: <https://doi.org/10.1214/aos/1176350695>. 160
- Choi, T. and Woo, Y. (2015). “A partially linear model using a Gaussian process prior.” *Communications in Statistics-Simulation and Computation*, 44(7): 1770–1786. MR3334556. doi: <https://doi.org/10.1080/03610918.2013.833226>. 176
- De Jonge, R., Van Zanten, J., et al. (2010). “Adaptive nonparametric Bayesian inference using location-scale mixture priors.” *The Annals of Statistics*, 38(6): 3300–3320. MR2766853. doi: <https://doi.org/10.1214/10-AOS811>. 160
- De Jonge, R., Van Zanten, J., et al. (2012). “Adaptive estimation of multivariate functions using conditionally Gaussian tensor-product spline priors.” *Electronic Journal of Statistics*, 6: 1984–2001. MR3020254. doi: <https://doi.org/10.1214/12-EJS735>. 162
- Devroye, L., Györfi, L., and Lugosi, G. (2013). *A probabilistic theory of pattern recognition*, volume 31. Springer Science & Business Media. MR1383093. doi: <https://doi.org/10.1007/978-1-4612-0711-5>. 159, 162
- Doob, J. L. (1949). “Application of the theory of martingales.” *Le calcul des probabilités et ses applications*, 23–27. MR0033460. 169
- Engle, R. F., Granger, C. W., Rice, J., and Weiss, A. (1986). “Semiparametric estimates of the relation between weather and electricity sales.” *Journal of the American statistical Association*, 81(394): 310–320. 160
- Fan, J. and Gijbels, I. (1996). *Local polynomial modelling and its applications: monographs on statistics and applied probability 66*, volume 66. CRC Press. MR1383587. 162, 170
- Fan, Y. and Li, Q. (1999). “Root-n-consistent estimation of partially linear time series models.” *Journal of Nonparametric Statistics*, 11(1-3): 251–269. MR1719473. doi: <https://doi.org/10.1080/10485259908832783>. 160

- Friedman, J. H. and Stuetzle, W. (1981). “Projection pursuit regression.” *Journal of the American statistical Association*, 76(376): 817–823. [MR0650892](#). 181
- Gao, C. and Zhou, H. H. (2015). “Rate-optimal posterior contraction for sparse PCA.” *The Annals of Statistics*, 43(2): 785–818. [MR3325710](#). doi: <https://doi.org/10.1214/14-AOS1268>. 166
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2014). *Bayesian data analysis*, volume 2. CRC press Boca Raton, FL. [MR3235677](#). 170
- Ghosal, S., Ghosh, J. K., and van der Vaart, A. W. (2000). “Convergence rates of posterior distributions.” *Annals of Statistics*, 28(2): 500–531. [MR1790007](#). doi: <https://doi.org/10.1214/aos/1016218228>. 166, 180
- Ghosal, S. and van der Vaart, A. (2017). *Fundamentals of nonparametric Bayesian inference*, volume 44. Cambridge University Press. [MR3587782](#). doi: <https://doi.org/10.1017/9781139029834>. 166
- Ghosal, S., van der Vaart, A., et al. (2007a). “Convergence rates of posterior distributions for non-i.i.d observations.” *The Annals of Statistics*, 35(1): 192–223. [MR2332274](#). doi: <https://doi.org/10.1214/009053606000001172>. 164, 180
- Ghosal, S., van der Vaart, A., et al. (2007b). “Posterior convergence rates of Dirichlet mixtures at smooth densities.” *The Annals of Statistics*, 35(2): 697–723. 166
- Ghosal, S. and van der Vaart, A. W. (2001). “Entropies and rates of convergence for maximum likelihood and Bayes estimation for mixtures of normal densities.” *Annals of Statistics*, 29(5): 1233–1263. [MR1873329](#). doi: <https://doi.org/10.1214/aos/1013203453>. 166
- Gu, M., Palomo, J., and Berger, J. O. (2016). “Robust GaSP: an R Package for objective Bayesian emulation of complex computer model codes.” *Technical Report*. 170
- Gu, M., Wang, X., and Berger, J. O. (2017). “Robust Gaussian Stochastic Process Emulation.” *arXiv preprint arXiv:1708.04738*. [MR3851764](#). doi: <https://doi.org/10.1214/17-AOS1648>. 170
- Györfi, L., Kohler, M., Krzyzak, A., and Walk, H. (2006). *A distribution-free theory of nonparametric regression*. Springer Science & Business Media. [MR1920390](#). doi: <https://doi.org/10.1007/b97848>. 159, 160, 162, 165
- Hastie, T. and Tibshirani, R. (1990). *Generalized additive models*. Wiley Online Library. [MR1082147](#). 160
- Hayfield, T., Racine, J. S., et al. (2008). “Nonparametric econometrics: The np package.” *Journal of Statistical Software*, 27(5): 1–32. [MR3410191](#). doi: [https://doi.org/10.1108/S0731-9053\(2009\)25](https://doi.org/10.1108/S0731-9053(2009)25). 176, 178
- Ichimura, H. (1993). “Semiparametric least squares (SLS) and weighted SLS estimation of single-index models.” *Journal of Econometrics*, 58(1-2): 71–120. [MR1230981](#). doi: [https://doi.org/10.1016/0304-4076\(93\)90114-K](https://doi.org/10.1016/0304-4076(93)90114-K). 181

- Jasra, A., Holmes, C. C., and Stephens, D. A. (2005). “Markov chain Monte Carlo methods and the label switching problem in Bayesian mixture modeling.” *Statistical Science*, 50–67. MR2182987. doi: <https://doi.org/10.1214/088342305000000016>. 164
- Klein, R. W. and Spady, R. H. (1993). “An efficient semiparametric estimator for binary response models.” *Econometrica: Journal of the Econometric Society*, 387–421. MR1209737. doi: <https://doi.org/10.2307/2951556>. 159
- Knapik, B. T., van der Vaart, A. W., van Zanten, J. H., et al. (2011). “Bayesian inverse problems with Gaussian priors.” *The Annals of Statistics*, 39(5): 2626–2657. MR2906881. doi: <https://doi.org/10.1214/11-AOS920>. 170
- Kruijer, W., Rousseau, J., van der Vaart, A., et al. (2010). “Adaptive Bayesian density estimation with location-scale mixtures.” *Electronic Journal of Statistics*, 4: 1225–1257. MR2735885. doi: <https://doi.org/10.1214/10-EJS584>. 165, 166
- Lenk, P. J. (1999). “Bayesian inference for semiparametric regression using a Fourier representation.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(4): 863–879. MR1722244. doi: <https://doi.org/10.1111/1467-9868.00207>. 160
- Nadaraya, E. A. (1964). “On estimating regression.” *Theory of Probability & Its Applications*, 9(1): 141–142. 161
- Nocedal, J. and Wright, S. (2006). *Numerical optimization*. Springer Science & Business Media. MR2244940. 181
- Pati, D., Bhattacharya, A., Pillai, N. S., and Dunson, D. (2014). “Posterior contraction in sparse Bayesian factor models for massive covariance matrices.” *The Annals of Statistics*, 42(3): 1102–1130. MR3210997. doi: <https://doi.org/10.1214/14-AOS1215>. 166
- Plumlee, M. (2017). “Bayesian calibration of inexact computer models.” *Journal of the American Statistical Association*, 1–12. MR3735376. doi: <https://doi.org/10.1080/01621459.2016.1211016>. 181
- Plumlee, M. and Joseph, V. R. (2016). “Orthogonal Gaussian process models.” *arXiv preprint arXiv:1611.00203*. MR3791080. 181
- Rasmussen, C. E. and Williams, C. K. (2006). *Gaussian processes for machine learning*, volume 1. MIT press Cambridge. MR2514435. 160
- Robinson, P. M. (1988). “Root-N-consistent semiparametric regression.” *Econometrica: Journal of the Econometric Society*, 931–954. MR0951762. doi: <https://doi.org/10.2307/1912705>. 159
- Ročková, V. (2018). “Bayesian estimation of sparse signals with a continuous spike-and-slab prior.” *The Annals of Statistics*, 46(1): 401–437. MR3766957. doi: <https://doi.org/10.1214/17-AOS1554>. 166
- Rockova, V. and van der Pas, S. (2017). “Posterior concentration for Bayesian regression trees and their ensembles.” *arXiv preprint arXiv:1708.08734*. 164

- Roustant, O., Ginsbourger, D., and Deville, Y. (2012). “DiceKriging, DiceOptim: Two R packages for the analysis of computer experiments by kriging-based metamodeling and optimization.” *Journal of Statistical Software*, 51(1): 54p. 170
- Shen, W., Tokdar, S. T., and Ghosal, S. (2013). “Adaptive Bayesian multivariate density estimation with Dirichlet mixtures.” *Biometrika*, 100(3): 623–640. MR3094441. doi: <https://doi.org/10.1093/biomet/ast015>. 164, 166
- Shen, X. and Wong, W. H. (1994). “Convergence rate of sieve estimates.” *The Annals of Statistics*, 580–615. MR1292531. doi: <https://doi.org/10.1214/aos/1176325486>. 166
- Speckman, P. (1988). “Kernel smoothing in partial linear models.” *Journal of the Royal Statistical Society. Series B (Methodological)*, 413–436. MR0970977. 160
- Stone, C. J. (1982). “Optimal global rates of convergence for nonparametric regression.” *The annals of statistics*, 1040–1053. MR0673642. 160, 165
- Szabó, B., van der Vaart, A. W., and van Zanten, J. H. (2015). “Frequentist coverage of adaptive nonparametric Bayesian credible sets.” *Ann. Statist.*, 43(4): 1391–1428. MR3357861. doi: <https://doi.org/10.1214/14-AOS1270>. 171, 172
- Takeda, H., Farsiu, S., and Milanfar, P. (2007). “Kernel regression for image processing and reconstruction.” *IEEE Transactions on image processing*, 16(2): 349–366. MR2462728. doi: <https://doi.org/10.1109/TIP.2006.888330>. 159
- Tang, Y., Sinha, D., Pati, D., Lipsitz, S., and Lipshultz, S. (2015). “Bayesian partial linear model for skewed longitudinal data.” *Biostatistics*, 16(3): 441–453. MR3365438. doi: <https://doi.org/10.1093/biostatistics/kxv005>. 160
- Tuo, R. and Wu, C. J. (2015). “Efficient calibration for imperfect computer models.” *The Annals of Statistics*, 43(6): 2331–2352. MR3405596. doi: <https://doi.org/10.1214/15-AOS1314>. 180
- van der Vaart, A. and van Zanten, H. (2007). “Bayesian inference with rescaled Gaussian process priors.” *Electronic Journal of Statistics*, 1: 433–448. MR2357712. doi: <https://doi.org/10.1214/07-EJS098>. 173
- van der Vaart, A. and Zanten, H. v. (2011). “Information rates of nonparametric Gaussian process methods.” *Journal of Machine Learning Research*, 12(Jun): 2095–2119. MR2819028. 160
- van der Vaart, A. W. and van Zanten, J. H. (2008). “Rates of contraction of posterior distributions based on Gaussian process priors.” *The Annals of Statistics*, 1435–1463. MR2418663. doi: <https://doi.org/10.1214/009053607000000613>. 160
- van der Vaart, A. W. and van Zanten, J. H. (2009). “Adaptive Bayesian estimation using a Gaussian random field with inverse Gamma bandwidth.” *The Annals of Statistics*, 2655–2675. MR2541442. doi: <https://doi.org/10.1214/08-AOS678>. 160, 173
- Watson, G. S. (1964). “Smooth regression analysis.” *Sankhyā: The Indian Journal of Statistics, Series A*, 359–372. MR0185765. 161

- Wooldridge, J. M. (2015). *Introductory econometrics: A modern approach*. Nelson Education. 175
- Xie, F., Jin, W., and Xu, Y. (2017). “A Theoretical Framework for Bayesian Nonparametric Regression: Orthonormal Random Series and Rates of Contraction.” *arXiv preprint arXiv:1712.05731*. 180
- Xie, F. and Xu, Y. (2017). “Bayesian Repulsive Gaussian Mixture Model.” *arXiv preprint arXiv:1703.09061*. 164
- Xie, F. and Xu, Y. (2019). “Supplementary Material for “Adaptive Bayesian Nonparametric Regression Using a Kernel Mixture of Polynomials with Application to Partial Linear Models”.” *Bayesian Analysis*. doi: <https://doi.org/10.1214/19-BA1148>. 161, 165, 169, 179
- Xu, Y., Mueller, P., and Telesca, D. (2016a). “Bayesian inference for latent biologic structure with determinantal point processes (DPP).” *Biometrics*, 72(3): 955–964. MR3545688. doi: <https://doi.org/10.1111/biom.12482>. 164
- Xu, Y., Xu, Y., and Saria, S. (2016b). “A Bayesian Nonparametric Approach for Estimating Individualized Treatment-Response Curves.” In *Machine Learning for Healthcare Conference*, 282–300. 159, 168
- Yang, Y., Bhattacharya, A., and Pati, D. (2017). “Frequentist coverage and sup-norm convergence rate in Gaussian process regression.” *arXiv preprint arXiv:1708.04753*. 160
- Yang, Y., Cheng, G., and Dunson, D. B. (2015). “Semiparametric Bernstein-von Mises Theorem: Second Order Studies.” *arXiv preprint arXiv:1503.04493*. MR3020419. doi: <https://doi.org/10.1214/13-EJS768>. 160, 169
- Yoo, W. W., Ghosal, S., et al. (2016). “Supremum norm posterior contraction and credible sets for nonparametric multivariate regression.” *The Annals of Statistics*, 44(3): 1069–1102. MR3485954. doi: <https://doi.org/10.1214/15-AOS1398>. 160, 170, 171