

Optimal Bayesian Estimation in Random Covariate Design with a Rescaled Gaussian Process Prior

Debdeep Pati

*Department of Statistics
Florida State University
Tallahassee, FL 32306-4330, USA*

DEBDEEP@STAT.FSU.EDU

Anirban Bhattacharya

*Department of Statistics
Texas A&M University
College Station, TX 77843-3143, USA*

ANIRBANB@STAT.TAMU.EDU

Guang Cheng

*Department of Statistics
Purdue University
West Lafayette, IN 47907-2066, USA*

CHENGG@PURDUE.EDU

Editor: Zhihua Zhang

Abstract

In Bayesian nonparametric models, Gaussian processes provide a popular prior choice for regression function estimation. Existing literature on the theoretical investigation of the resulting posterior distribution almost exclusively assume a fixed design for covariates. The only random design result we are aware of (van der Vaart and van Zanten, 2011) assumes the assigned Gaussian process to be supported on the smoothness class specified by the true function with probability one. This is a fairly restrictive assumption as it essentially rules out the Gaussian process prior with a squared exponential kernel when modeling rougher functions. In this article, we show that an appropriate rescaling of the above Gaussian process leads to a rate-optimal posterior distribution even when the covariates are independently realized from a known density on a compact set. The proofs are based on deriving sharp concentration inequalities for frequentist kernel estimators; the results might be of independent interest.

Keywords: Bayesian, convergence rate, Gaussian process, nonparametric regression, random design, rate-optimal

1. Introduction

Gaussian processes (Rasmussen, 2004; Seeger, 2004; Rasmussen and Williams, 2006) are widely used in the machine learning community as a principled probabilistic approach to function estimation. A mean-zero Gaussian process is completely specified by its covariance kernel; popular choices include the squared-exponential and Matérn families. Recently, there has been significant interest in frequentist convergence properties of Bayesian posteriors in Gaussian process models. Ghosal and Roy (2006); Choi and Schervish (2007); Tokdar and Ghosh (2007) established posterior consistency in a variety of settings including nonparametric regression, classification and density estimation. Seeger et al. (2008) used an

information criterion to evaluate closeness of the posterior distribution to the truth; see also van der Vaart and van Zanten (2011). A major focus in the recent literature (van der Vaart and van Zanten, 2007, 2008a, 2009, 2011; Bhattacharya et al., 2014; Shang and Cheng, 2014) has been on deriving the posterior convergence rate (Ghosal et al., 2000), which is defined as the minimum possible sequence $\epsilon_n \rightarrow 0$ such that for some constant $M > 0$,

$$E_{\theta_0} \Pi(\|\theta - \theta_0\| < M\epsilon_n \mid \mathcal{D}_n) \rightarrow 1, \quad (1)$$

where \mathcal{D}_n denotes the data, θ is the parameter of interest with some known transformation $\Psi(\theta)$ assigned a Gaussian process prior, θ_0 is the true data generating parameter and $\|\cdot\|$ is a distance measure relevant to the statistical problem. In the context of nonparametric regression, classification and density estimation, it has been established that the posterior convergence rate based on appropriate Gaussian process priors coincides with the minimax optimal rate $n^{-\alpha/(2\alpha+d)}$ for d -variate α -smooth functions up to a logarithmic factor (van der Vaart and van Zanten, 2007, 2008a), with rate-adaptivity to the unknown smoothness achieved in van der Vaart and van Zanten (2009); Bhattacharya et al. (2014).

In this paper, we focus on a non-parametric regression problem,

$$Y_i = f(X_i) + \epsilon_i, \quad \epsilon_i \sim N(0, 1), \quad (2)$$

with f assigned a mean-zero Gaussian process prior. The above-mentioned literature on posterior convergence rates under (2) typically assume that the covariates X_i 's are fixed by design, in which the empirical L_2 norm $\|f - f_0\|_{2,n} = (1/n \sum_{i=1}^n |f(x_i) - f_0(x_i)|^2)^{1/2}$ is used as a discrepancy measure in (1). The empirical L_2 norm evaluates the discrepancy of the estimated function from the true function only at the observed data-points and is not suitable to assess out-of-sample predictive performance. In this paper, we consider a *random design* setup where the covariates X_i 's are drawn independently from a known distribution q , and derive the posterior convergence rates under an integrated $L_1(q)$ metric:

$$\|f - f_0\|_{1,q} = \int |f(x) - f_0(x)|q(x)dx.$$

The above integrated $L_1(q)$ metric is more relevant for studying statistical efficiency in a random design setting. From a technical standpoint, dealing with the integrated metric is challenging since one cannot directly leverage on properties of multivariate Gaussian distributions as in the case of the empirical L_2 norm to construct “test functions”; a key ingredient in Bayesian asymptotics.

In the frequentist literature, existing results (Baraud, 2002; Brown et al., 2002; Birgé, 2004) on the convergence rates (with respect to an integrated metric) in random design regression require an appropriate lower bound on the smoothness of the underlying true function. For example, Brown et al. (2002); Birgé (2004) assumed that the univariate function f_0 belongs to a Lipschitz class with smoothness index $\alpha > 1/2$. Moreover, Birgé (2004) demonstrated the necessity of the $\alpha > 1/2$ condition by establishing a lower bound for the asymptotic risk for $\alpha \leq 1/2$. Similar lower bound condition will be assumed in our main Bayesian Theorem as well.

As far as we are aware, the only Bayesian literature considering the random design setting in (2) is van der Vaart and van Zanten (2011) who assigned Gaussian processes

with Matérn or squared exponential kernels. Specifically, they obtained an optimal rate $n^{-\alpha/(2\alpha+d)}$ (up to a logarithmic factor, with respect to $L_2(q)$ norm) under a particularly strong assumption that the Gaussian process prior assigns probability one to the smoothness class containing the true function. Since the squared-exponential kernel has infinitely smooth sample paths, their result only delivers the optimal rate for analytic functions, but provides a highly suboptimal $(\log n)^{-t}$ rate for α -smooth functions. This significantly limits the applicability of their result in the sense that it rules out the use of a squared-exponential kernel for less smooth (but more commonly used) functions. An influential idea developed in van der Vaart and van Zanten (2007, 2009) is to scale the sample paths of a Gaussian process with a squared-exponential kernel to enable better approximation of α -smooth functions. The scaling is typically dependent on the smoothness of the true function and the sample size. However, Theorem 2 of van der Vaart and van Zanten (2011) is applicable only to priors without scaling. This is not evident from the statement of their theorem, but a closer inspection of their proof (ref. Page 2113) reveals that they have assumed $\tau^2 := \int \|f\|_{\alpha|\infty}^2 d\Pi(f)$ to be a global constant for every f in the support of the prior. This may not hold for a rescaled Gaussian process.

In this article, we show that an appropriately rescaled Gaussian process prior with a squared-exponential covariance kernel leads to a rate-optimal posterior distribution (with respect to $L_1(q)$ norm) for any α -smooth function d -variate f_0 in a random design setting if $\alpha > d/2$. While van der Vaart and van Zanten (2011) conjectured (see pp. 2103 after Theorem 2) that their smoothness assumption on the prior is unavoidable for the $L_2(q)$ norm, our result shows that this situation turns out to be different under the $L_1(q)$ norm. Specifically, we develop exponentially consistent test functions under the $L_1(q)$ norm using concentration inequalities for the Nadaraya–Watson kernel estimator. Existence of such test functions plays a key role in Bayesian asymptotic theory (Ghosal et al., 2000). For example, the classical Birgé – Le Cam testing theory (Birgé, 1984; Le Cam, 1986) for the Hellinger metric provides appropriate tests in a wide variety of settings. Giné and Nickl (2011) proposed an alternative framework for constructing tests based on concentration inequalities of frequentist estimators which is particularly useful for stronger norms; see also Ray (2013); Pati et al. (2014); Shang and Cheng (2014) for similar ideas in different contexts.

2. Posterior Convergence in Random Design Regression

2.1 Notations

Let $C[0, 1]^d$ and $C^\alpha[0, 1]^d$ denote the space of continuous functions and the Hölder space of α -smooth functions $f : [0, 1]^d \rightarrow \mathbb{R}$, respectively, endowed with the supremum norm $\|f\|_\infty = \sup_{t \in [0, 1]^d} |f(t)|$. For $\alpha > 0$, the Hölder space $C^\alpha[0, 1]^d$ consists of functions $f \in C[0, 1]^d$ that have bounded mixed partial derivatives up to order $[\alpha]$, with the partial derivatives of order $[\alpha]$ being Lipschitz continuous of order $\alpha - [\alpha]$. Let $\|\cdot\|_1$ and $\|\cdot\|_2$ respectively denote the L_1 and L_2 norm on $[0, 1]^d$ with respect to the Lebesgue measure (i.e., the uniform distribution). To distinguish the L_2 norm with respect to the Lebesgue measure on \mathbb{R}^d , we use the notation $\|\cdot\|_{2,d}$.

We write “ \lesssim ” for inequality up to a constant multiple. Let $\phi(t) = (2\pi)^{-1/2} \exp(-t^2/2)$ denote the standard normal density, and let $\phi_n(x) = \prod_{i=1}^n \phi(x_i)$ for $x \in \mathbb{R}^n$. Let a star denote a convolution, i.e., $f_1 \star f_2(y) = \int f_1(y-t)f_2(t)dt$. We denote the Fourier transform of f , whenever defined, by \hat{f} , with $\hat{f}(\lambda) = (2\pi)^{-d} \int \exp(i\langle \lambda, t \rangle) f(t) dt$, where $\langle \lambda, t \rangle$ denotes the complex inner product. Under this convention, the inverse Fourier transform $f(t) = \int \exp(-i\langle \lambda, t \rangle) \hat{f}(\lambda) d\lambda$ and $\hat{h} = (2\pi)^d \hat{f} \hat{g}$ when $h = f \star g$.

Throughout C, C', C_1, C_2, \dots are generically used to denote positive constants whose values might change from one line to another, but are independent from everything else. $Z_{1:n}$ is used as a shorthand for Z_1, \dots, Z_n .

In the sequel, we consider a Gaussian process prior Π on the regression function f with $\mathbb{E}f(x) = 0$ and covariance kernel $c(x, x') = \text{cov}(f(x), f(x'))$. In particular, we focus on the squared-exponential kernel $c_a(x, x') = \exp(-a^2 \|x - x'\|^2)$ indexed by an “inverse-bandwidth” parameter a . We next recall some important facts relevant to our setting from van der Vaart and van Zanten (2009) regarding the spectral measure and reproducing kernel Hilbert space of Gaussian process priors. For the squared-exponential kernel c_a , the spectral measure μ_a admits a density ω_a with respect to Lebesgue measure, where $\omega_a(\lambda) = a^{-d} \omega(\lambda/a)$, with $\omega(\lambda) = \exp(-\|\lambda\|^2/4)/(2^d \pi^{d/2})$. The reproducing kernel Hilbert space \mathbb{H}^a associated with a Gaussian process prior Π consists of (real parts of) functions $h(t) = \int \exp(i\langle \lambda, t \rangle) \xi(\lambda) d\mu_a(\lambda)$, where μ_a is the spectral measure of Π and $\xi \in L_2(\mu_a)$. The squared Hilbert space norm of h above is given by $\|h\|_{\mathbb{H}^a}^2 = \left\| \xi \omega_a^{1/2} \right\|_{2,d}^2 = \int \xi^2(\lambda) \omega_a(\lambda) d\lambda$; let \mathbb{H}_1^a denote the unit ball of the reproducing kernel Hilbert space $\{h \in \mathbb{H}^a : \|h\|_{\mathbb{H}^a} \leq 1\}$. Finally, let \mathbb{B}_1 denote the unit ball of $C[0, 1]^d$ with respect to the supremum norm. For a detailed review of reproducing kernel Hilbert space of Gaussian process priors and connections with posterior contraction rates, kindly refer to van der Vaart and van Zanten (2008b).

2.2 Main Result

Consider the nonparametric regression model (2). We assume a random design setup, where given the regression function $f : [0, 1]^d \rightarrow \mathbb{R}$, the data $(X_1, Y_1), \dots, (X_n, Y_n)$ are independently generated, with X_i having a density q on $[0, 1]^d$ that is bounded away from zero and infinity. Let $q(y, x) = q(y | x)q(x)$ denote the joint density of (Y, X) given f , where $q(y | x) = \phi\{y - f(x)\}$. The joint data likelihood given f is therefore

$$q^{(n)}(Y_{1:n}, X_{1:n} | f) = \prod_{i=1}^n q(Y_i, X_i) = \prod_{i=1}^n \phi\{Y_i - f(X_i)\} q(X_i).$$

Similarly, we define $q^{(n)}(Y_{1:n} | X_{1:n}, f)$ and $q^{(n)}(X_{1:n})$ as the density of $(Y_{1:n} | X_{1:n}, f)$ and $X_{1:n}$ respectively. Let $\mathbb{E}_{Y,X}^f(\mathbb{P}_{Y,X}^f)$ denote an expectation (probability) with respect to $q^{(n)}(Y_{1:n}, X_{1:n} | f)$. Similarly define $\mathbb{E}_{Y|X}^f(\mathbb{P}_{Y|X}^f)$ and $\mathbb{E}_X^f(\mathbb{P}_X^f)$. When f is clear from the context, we shall drop it from the superscript.

We assume a mean zero Gaussian process prior Π on f with a squared exponential kernel $\exp(-a_n^2 \|x - x'\|^2)$ and denote the corresponding posterior measure by $\Pi(\cdot | Y_{1:n}, X_{1:n})$, so that

$$\Pi(f | Y_{1:n}, X_{1:n}) \propto q^{(n)}(Y_{1:n} | X_{1:n}, f) \Pi(f).$$

Assuming the true regression function is f_0 , we study concentration of the posterior $\Pi(\cdot | Y_{1:n}, X_{1:n})$ in an $L_1(q)$ neighborhood of f_0 .

Theorem 1 *Assume that $f_0 \in C^\alpha[0, 1]^d$ with $\alpha > d/2$ and Π is a mean-zero Gaussian process prior with a squared exponential covariance kernel $c(x, x') = \exp(-a_n^2 \|x - x'\|^2)$. Set $a_n = n^{1/(2\alpha+d)}$. Then with $\epsilon_n = n^{-\alpha/(2\alpha+d)} \log^{3t_1/2} n$ for $t_1 \geq (d + 1)/2$, and some fixed sufficiently large constant $M > 0$,*

$$\mathbb{E}_{Y, X}^{f_0} \Pi(\|f - f_0\|_{1,q} > M\epsilon_n | Y_{1:n}, X_{1:n}) \rightarrow 0. \tag{3}$$

As stated previously, the condition $\alpha > d/2$ is necessary to obtain the optimal rate. van der Vaart and van Zanten (2009) showed that the squared-exponential covariance kernel without rescaling leads to a very slow $(\log n)^{-l}$ contraction rate for α -smooth functions both in the fixed and random design settings. This is not surprising as the sample paths of such a GP are analytic. The effect of scaling the prior using the “inverse bandwidth” a to yield the optimal posterior concentration was first noted by van der Vaart and van Zanten (2007) in a fixed design context, who showed (for $d = 1$) that a deterministic scaling $a_n = n^{1/(2\alpha+1)}$ produces priors that are suitable for modeling α -smooth functions. Theorem 1 assures that the same rescaling idea continues to work in the *random* design setting for an integrated L_1 norm.

The optimal rescaling in Theorem 1 requires knowledge of the true smoothness α . If there is a mismatch between the prior regularity and the function smoothness, one would typically expect a sub-optimal rate. Corollary 2 quantifies this fact; while we only derive an upper bound to the posterior convergence rate, such bounds are usually tight (van der Vaart and van Zanten, 2009). In absence of any prior knowledge regarding the smoothness, one may resort to an empirical or fully Bayes approach as in van der Vaart and van Zanten (2009); Szabó et al. (2013). The related theoretical investigation will be considerably harder in such cases.

Corollary 2 *Under the conditions of Theorem 1, if $a_n = n^{1/(2\beta+d)}$ for $\beta > d/2, \beta \neq \alpha$, the conclusion of Theorem 1 holds with $\epsilon_n = n^{-\alpha \wedge \beta / (2\beta+d)} \log^{3t_1/2} n$ for $t_1 \geq d/(4 - 2\kappa)$ for any $0 < \kappa < 2$.*

Observe that the optimal rate $n^{-\alpha/(2\alpha+d)}$ is attained (upto logarithmic terms) if and only if $\alpha = \beta$. A scaling $n^{1/(2\beta+d)}$ for $\beta < \alpha$ makes the prior rougher compared to the true function while $\beta > \alpha$ renders smoother prior realizations. In both cases, sub-optimal rates are obtained. This is in accordance with the findings for GP priors with Matérn covariance kernel; refer to Theorem 5 in van der Vaart and van Zanten (2011).

Note that by taking κ very close to 0, we can improve on the power of the $\log n$ term in Corollary 2 from that in Theorem 1. The difference in the power of $\log n$ stems from the fact that the corollary only targets sub-optimal rates as opposed to Theorem 1. Hence a portion of the power of $\log n$ can be eliminated in Corollary 2.

2.3 Contributions Beyond Literature

The proof of Theorem 1 follows from a general set of sufficient conditions for posterior concentration in model (2); kindly refer to Theorem 3 stated in the next Section. In

particular, we exploit concentration inequalities for suitable kernel estimators to construct the aforementioned exponentially consistent sequence of test functions. Such techniques have been used previously to show convergence rates in density estimation (Giné and Nickl, 2011) and in linear inverse problems (Ray, 2013). Their techniques do not directly apply to our case partly due to the lack of concentration bounds for kernel based estimators. Giné and Nickl (2011); Ray (2013) construct estimators based on truncated spectral representations which are well suited to sieve priors. However, to deal with a Gaussian process prior with a squared-exponential covariance kernel, we need to construct test functions based on the Nadaraya–Watson kernel estimator and derive sharp concentration bounds for this class of estimators in Lemma 4 and 5.

The choice of the norm dictating the neighborhood around the true parameter plays a critical role in Bayesian asymptotics. A fundamental tool for relating the likelihood ratio with the neighborhood in consideration is a sequence of exponentially consistent test functions (Ghosal et al., 2000). In a regression setting, such test functions are guaranteed to exist for the empirical L_2 norm by exploiting a direct relation between the empirical L_2 norm and the likelihood ratio of the multivariate Gaussian densities involved; refer to Section 4 of van der Vaart and van Zanten (2011). However, the integrated norm involves covariate points different from the observations, which makes the problem more challenging. van der Vaart and van Zanten (2011) applied Bernstein’s inequality to extrapolate to the $L_2(q)$ norm from the empirical L_2 norm. However, as stated in the Introduction, their approach only works for priors that are supported on the true smoothness class with probability one.

Among other related work, Section 4 of Kleijn and van der Vaart (2006) considers random design regression, where a correspondence between the Kullback–Leibler and $L_2(q)$ neighborhood is established to derive the test function, assuming the prior support consists of uniformly bounded functions. However, this assumption does not hold for the rescaled Gaussian process prior in Theorem 1. In particular, the *sieves* constructed in van der Vaart and van Zanten (2007) of the form $M_n \mathbb{H}_1^{q_n} + \epsilon_n \mathbb{B}_1$ with $M_n \rightarrow \infty$ do not correspond to sup-norm bounded subsets of $C[0, 1]^d$.

We comment here that convergence in the integrated metric has been settled in the binary regression setting. Using a logistic link function, a direct agreement can be established between the integrated L_1 metric on the function space and the Hellinger distance between the resulting densities arising from the Bernoulli likelihood; see for example, Section 3.2 of van der Vaart and van Zanten (2008a). Second, in this paper we implicitly refer to Gaussian processes which are specified by a kernel function; specifically, kernel functions which do not admit a finite series representation, such as the squared-exponential kernel. If a Gaussian process is specified via a truncated orthogonal series representation with independent Gaussian priors on the coefficients, the integrated metric can be related to the L_2 norm of the coefficient vector (Bontemps, 2011).

3. Auxiliary Results

We now state a general theorem which presents a set of sufficient conditions for proving Theorem 1. From now onwards, we shall assume the covariate distribution q to be a uniform distribution on $[0, 1]^d$ for notational simplicity; modifying our construction to a general q , which is bounded from above and below, is straightforward. The $L_1(q)$ norm $\|\cdot\|_{1,q}$ with q

the uniform distribution on $[0, 1]^d$ shall be simply denoted by $\|\cdot\|_1$ following our convention in Section 2.1. A proof of Theorem 3 can be found in the Appendix.

Theorem 3 *Let ϵ'_n, δ_n be sequences such that $\epsilon'_n, \delta_n \rightarrow 0$ and $n\epsilon_n^2, n\delta_n^2 \rightarrow \infty$. Let $U_n = \{f : \|f - f_0\|_1 > M\epsilon'_n\}$ for some fixed $M > 0$. Suppose that there exists a sequence of estimators \tilde{f}_n for f based on $(Y_{1:n}, X_{1:n})$ and a sequence of subsets/sieves \mathcal{P}_n of $C[0, 1]^d$ such that*

$$\Pi(\mathcal{P}_n^c) \leq \exp\{-(C + 4)n\delta_n^2\}, \tag{PCS}$$

$$\left\| \mathbb{E}_{Y,X}^{f_0} \tilde{f}_n - f_0 \right\|_1 < \epsilon'_n, \tag{BT}$$

$$\mathbb{P}_{Y,X}^{f_0} \left(\left\| \tilde{f}_n - \mathbb{E}_{Y,X}^{f_0} \tilde{f}_n \right\|_1 > \epsilon'_n \right) \leq \exp\{-(C + 4)n\delta_n^2\}, \tag{DT}$$

$$\sup_{f \in \mathcal{P}_n \cap U_n} \left\| \mathbb{E}_{Y,X}^f \tilde{f}_n - f \right\|_1 < \epsilon'_n, \tag{BS}$$

$$\sup_{f \in \mathcal{P}_n \cap U_n} \mathbb{P}_{Y,X}^f \left(\left\| \tilde{f}_n - \mathbb{E}_{Y,X}^f \tilde{f}_n \right\|_1 > \epsilon'_n \right) \leq \exp\{-(C + 4)n\delta_n^2\}, \tag{DS}$$

$$\Pi \left(\|f - f_0\|_\infty \leq \delta_n \right) \geq \exp\{-n\delta_n^2\}. \tag{PCN}$$

Then, $\mathbb{E}_{Y,X}^{f_0} \Pi(U_n | Y_{1:n}, X_{1:n}) \rightarrow 0$.

Condition (PCS) implies that the prior probability of the complement of the sieve \mathcal{P}_n is exponentially small. Condition (BT) assumes a sufficiently accurate estimator \tilde{f}_n with bias smaller than ϵ_n at f_0 while (DT) assumes an exponential concentration bound of \tilde{f}_n from its expectation under $q^{(n)}(\cdot | f_0)$. (BS) and (DS) assume similar conditions as (BT) and (DT) under $q^{(n)}(\cdot | f)$ for any $f \in \mathcal{P}_n \cap U_n$. The conditions (BT), (DT); (BS), (DS) jointly guarantee the existence of exponentially consistent test functions; see Lemma 9 in the Appendix. Condition (PCN) assumes that the prior Π places “enough” mass in an ϵ_n -neighborhood of the truth f_0 in terms of the sup-norm.

3.1 Verifying the Conditions of Theorem 3 to Prove Theorem 1

Letting $\delta_n = \epsilon'_n = \epsilon_n$ with ϵ'_n and ϵ_n as in the statement of Theorem 3 and Theorem 1 respectively, we now proceed to construct \mathcal{P}_n and \tilde{f}_n that satisfy the conditions of Theorem 3. While we choose the same sieve as in van der Vaart and van Zanten (2007), part of the technical challenge lies in the fact that the concentration bounds need to be derived not just for the truth, but rather for every function in the sieve. This requires precise control on the *size* of the functions in the sieve \mathcal{P}_n . We show in Proposition 7 below that the functions in the chosen sieve are uniformly bounded in L_2 norm, although they are unbounded in the supremum norm.

Let $\psi : \mathbb{R}^d \rightarrow \mathbb{C}$ be a function such that $\int \psi(t) dt = 1$, $\int t^k \psi(t) dt = 0$ for any non-zero multi-index $k = (k_1, \dots, k_d)$ of non-zero integers, $\int |t|^{\max\{\alpha, 2\}} |\psi(t)| dt < \infty$, and the functions $|\hat{\psi}|/\omega$ and $|\hat{\psi}|^2/\omega$ are uniformly bounded; see proof of Lemma 4.3 in van der Vaart and van Zanten (2009). Define,

$$\tilde{f}_n(x) = \frac{1}{n} \sum_{i=1}^n \psi_{\sigma_n}(x - X_i) Y_i, \tag{4}$$

where $\psi_\sigma(t) = \sigma^{-d}\psi(t/\sigma)$ for $\sigma > 0$ and set $\sigma_n = n^{-1/(2\alpha+d)} \log^{-t_2} n$, $t_2 = 1/(2 - \kappa)$ for some $0 < \kappa < 1$. Next, with $M_n = a_n^{d/2}$, set

$$\mathcal{P}_n = M_n \mathbb{H}_1^{a_n} + \epsilon_n \mathbb{B}_1. \quad (5)$$

Assume $f_0 \in C^\alpha[0, 1]^d$. Let \tilde{f}_n and \mathcal{P}_n be as in (4) and (5) respectively. We show below that the conditions of Theorem 3 are satisfied with $\epsilon_n = n^{-\alpha/(2\alpha+d)} \log^{t_1} n$ for $t_1 \geq \max\{t_2 d/2, (d+1)/2\} = (d+1)/2$, provided $\alpha > d/2$. (PCS) follow from the proof of Theorem 3.1 in van der Vaart and van Zanten (2009). To verify (PCN), observe from the proof of Theorem 3.1 in van der Vaart and van Zanten (2009) that with $a_n = n^{1/(2\alpha+d)}$,

$$\mathbb{P}(\|f - f_0\|_\infty \leq \delta_n) \geq \exp\{-n^{d/(2\alpha+d)}(\log n)^{d+1}\}$$

for $\delta_n \geq n^{-\alpha/(2\alpha+d)}$. Hence (PCN) is satisfied with $\delta_n = \epsilon_n$.

We verify (BS) and (DS); the verifications of (BT) and (DT) follow along similar lines.

We first show that (DS) holds. Fix $f \in \mathcal{P}_n \cap U_n$. We drop the superscript f from \mathbb{E}^f in the sequel. Let $f_n(x) = \psi_{\sigma_n} \star f(x) = \int \psi_{\sigma_n}(x-t)f(t)dt$ and $f_n^X(x) = n^{-1} \sum_{i=1}^n \psi_{\sigma_n}(x - X_i)f(X_i)$. Observe that $\mathbb{E}_{Y,X} f_n = f_n$ and $\mathbb{E}_{Y|X} f_n = f_n^X$. Then,

$$\begin{aligned} \mathbb{P}_{Y,X} \left(\left\| \tilde{f}_n - \mathbb{E}_{Y,X} \tilde{f}_n \right\|_1 > \epsilon_n \right) &= \mathbb{P}_{Y,X} \left(\left\| \tilde{f}_n - f_n \right\|_1 > \epsilon_n \right) \\ &\leq \mathbb{P}_{Y,X} \left(\left\| \tilde{f}_n - f_n \right\|_1 > \epsilon_n, \|f_n^X - f_n\|_1 \leq \epsilon_n/2 \right) + \mathbb{P}_X \left(\|f_n^X - f_n\|_1 \geq \epsilon_n/2 \right) \\ &\leq \mathbb{E}_X \mathbb{P}_{Y|X} \left(\left\| \tilde{f}_n - f_n^X \right\|_1 \geq \epsilon_n/2 \mid X_{1:n} \right) + \mathbb{P}_X \left(\|f_n^X - f_n\|_1 \geq \epsilon_n/2 \right). \end{aligned} \quad (6)$$

Lemmata 4 and 5 below deliver the desired bounds for the two terms appearing in (6).

Lemma 4 *Under conditions of Theorem 1,*

$$\mathbb{P}_{Y|X} \left(\left\| \tilde{f}_n - f_n^X \right\|_1 \geq \epsilon_n/2 \mid X_{1:n} \right) \leq \exp(-Cn\epsilon_n^2) \quad \text{a.s.}$$

for some constant $C > 0$.

Proof For simplicity of notation, we suppress the term ‘‘a.s.’’ in the displays that follow. Let $T(x) = n(\tilde{f}_n - f_n^X)(x) = \sum_{i=1}^n \psi_{\sigma_n}(x - X_i)Z_i$, where $Z_i = Y_i - f(X_i)$ with $Z_{1:n} \mid X_{1:n}, f \sim N_n(0, \mathbf{I})$. Given $X_{1:n}$, T is a random element of $L_1[0, 1]^d$ and $\|T\|_1$ is a non-negative random variable. By the Hahn–Banach theorem, there exists a bounded linear functional G on $L_\infty[0, 1]^d$ such that $G(h) = \int T(x)h(x)dx$ for all $h \in L_\infty[0, 1]^d$ and $\|T\|_1 = \|G\|_{\mathcal{F}}$, where $\|G\|_{\mathcal{F}} = \sup_{h \in \mathcal{F}} |G(h)|$ and \mathcal{F} is a countable dense subset of $\{h \in L_\infty[0, 1]^d : \|h\|_\infty \leq 1\}$.

By definition, $G(h) = \sum_{i=1}^n a_i Z_i$, where $a_i = \int \psi_{\sigma_n}(x - X_i)h(x)dx$. Thus, given $X_{1:n}$, $\{G(h) : h \in \mathcal{F}\}$ is a Gaussian process and

$$\mathbb{P}_{Y|X} \left(\left\| \tilde{f}_n - f_n^X \right\|_1 \geq \epsilon_n/2 \mid X_{1:n} \right) = \mathbb{P}_{Y|X} \left(\|G\|_{\mathcal{F}} \geq n\epsilon_n/2 \mid X_{1:n} \right). \quad (7)$$

By Borell's inequality (Adler, 1990),

$$\mathbb{P}_{Y|X}(\|G\|_{\mathcal{F}} - \mathbb{E}_{Y|X} \|G\|_{\mathcal{F}} \geq t \mid X_{1:n}) \leq 2 \exp\{-t^2/(2\sigma_{\mathcal{F}}^2)\}, \tag{8}$$

where $\sigma_{\mathcal{F}}^2 = \sup_{h \in \mathcal{F}} \mathbb{E}_{Y|X} G(h)^2$. We now proceed to estimate $\sigma_{\mathcal{F}}^2$ and $\mathbb{E}_{Y|X} \|G\|_{\mathcal{F}}$. For any $h \in \mathcal{F}$,

$$\begin{aligned} \mathbb{E}_{Y|X} G(h)^2 &= \sum_{i=1}^n \left\{ \int_{[0,1]^d} \psi_{\sigma_n}(x - X_i) h(x) dx \right\}^2 \\ &\leq \|h\|_{\infty}^2 \sum_{i=1}^n \left\{ \int_{[0,1]^d} |\psi_{\sigma_n}(x - X_i)| dx \right\}^2 \leq C_1 n, \end{aligned} \tag{9}$$

where $C_1 = \int |\psi(t)| dt$. Hence, $\sigma_{\mathcal{F}}^2 \leq C_1 n$.

We next bound $\mathbb{E}_{Y|X} \|G\|_{\mathcal{F}} = \mathbb{E}_{Y|X} \|T\|_1$. By Jensen's inequality, $\mathbb{E}_{Y|X} \|T\|_1 \leq (\mathbb{E}_{Y|X} \|T\|_1^2)^{1/2}$. Further,

$$(\mathbb{E}_{Y|X} \|T\|_1^2)^{1/2} = \left[\int \left\{ \int |T(x)| dx \right\}^2 \phi_n(z) dz \right]^{1/2} \leq \int \left\{ \int T(x)^2 \phi_n(z) dz \right\}^{1/2} dx,$$

where the above inequality follows from an integral version of Minkowski's inequality. Recalling $T(x) = \sum_{i=1}^n \psi_{\sigma_n}(x - X_i) Z_i$, $\int T(x)^2 \phi_n(z) dz = \mathbb{E}[T(x)^2 \mid X_{1:n}] = \sum_{i=1}^n \psi_{\sigma_n}(x - X_i)^2$. Substituting this in the above display and using Jensen's inequality one more time, we get

$$\begin{aligned} \mathbb{E}_{Y|X} \|T\|_1 &\leq \int \left\{ \sum_{i=1}^n \psi_{\sigma_n}(x - X_i)^2 \right\}^{1/2} dx \\ &\leq \left\{ \sum_{i=1}^n \int \psi_{\sigma_n}(x - X_i)^2 dx \right\}^{1/2} \leq C_2 (n/\sigma_n^d)^{1/2}, \end{aligned} \tag{10}$$

where $C_2 = \{\int \psi(t)^2 dt\}^{1/2}$. In (8), set $t = n\epsilon_n/4$. From the above calculations, $\mathbb{E}_{Y|X} \|G\|_{\mathcal{F}} \leq C_2 (n/\sigma_n^d)^{1/2} \leq n\epsilon_n/4$ since $t_1 \geq t_2 d/2$. Using $\sigma_{\mathcal{F}}^2 \leq C_1 n$, we finally obtain $\mathbb{P}_{Y|X} \left(\|G\|_{\mathcal{F}} \geq n\epsilon_n/2 \mid X_{1:n} \right) \leq 2 \exp(-Cn\epsilon_n^2)$. ■

Lemma 5 *Under conditions of Theorem 1,*

$$\mathbb{P}_X \left(\|f_n^X - f_n\|_1 \geq \epsilon_n/2 \right) \leq \exp(-n\epsilon_n^2).$$

Proof As in Lemma 4, we express the desired probability in terms of a tail bound for the supremum of a stochastic process. However, the stochastic process in this case is no longer a Gaussian process and we cannot use Borell's inequality here. We instead use Bosquet's version of Talagrand's inequality for the supremum of a centered empirical process. The following Proposition 6 is adapted from Bousquet (2003) which also appears in Section 3.1 of Giné and Nickl (2011).

Proposition 6 *Assume X_1, \dots, X_n are independent and identically distributed as P . Let \mathcal{G} be a countable set of real valued functions and assume all functions $g \in \mathcal{G}$ are P -measurable, square integrable and satisfy $\mathbb{E}_P[g] = 0$. Assume $K_1 = \sup_{g \in \mathcal{G}} \|g\|_\infty < \infty$ and let $W = \sup_{g \in \mathcal{G}} |\sum_{i=1}^n g(X_i)|$. Further, let $\sigma_{\mathcal{G}}^2 = \sup_{g \in \mathcal{G}} \mathbb{E}_P[g(X_1)^2]$ and $K_2 = n\sigma_{\mathcal{G}}^2 + K_1\mathbb{E}_P[W]$. Then, for any $t > 0$,*

$$\mathbb{P}\left\{W \geq \mathbb{E}_P W + (2K_2 t)^{1/2} + \frac{K_1 t}{3}\right\} \leq \exp(-t).$$

Let $L_x(t) = \psi_{\sigma_n}(x-t)f(t) - \psi_{\sigma_n} \star f(x)$ for $x, t \in [0, 1]^d$ and $W = \int_{[0,1]^d} |\sum_{i=1}^n L_x(X_i)| dx$. Clearly, $\mathbb{P}_X(\|f_n^X - f_n\|_1 > \epsilon_n/2) = \mathbb{P}_X(W > n\epsilon_n/2)$. By an application of Hahn–Banach theorem as in the proof of Lemma 4, $W = \|G\|_{\mathcal{F}}$, where \mathcal{F} is a countable dense subset of the unit ball of $L_\infty[0, 1]^d$, $G(h) = \sum_{i=1}^n g(X_i)$, and $g(t) = \int_{[0,1]^d} L_x(t)h(x)dx$. Letting \mathcal{G} denote the class of functions $\{g(t) = \int_{[0,1]^d} L_x(t)h(x)dx, h \in \mathcal{F}\}$, one has $\|G\|_{\mathcal{F}} = \sup_{g \in \mathcal{G}} |\sum_{i=1}^n g(X_i)|$. Putting together, $W = \sup_{g \in \mathcal{G}} |\sum_{i=1}^n g(X_i)|$ and $\mathbb{E}_X g(X_1) = 0$ by Tonelli’s theorem. We now aim to apply Proposition 6 to bound $\mathbb{P}_X(W > n\epsilon_n/2)$. In order to apply Proposition 6, we need to estimate $K_1, \sigma_{\mathcal{G}}^2, K_2$ and $\mathbb{E}_P(W)$ which is carried out below.

Fix $g \in \mathcal{G}$. Then, there exists $h \in \mathcal{F}$ such that $g(t) = \int_{[0,1]^d} L_x(t)h(x)dx = f(t) \int_{[0,1]^d} \psi_{\sigma_n}(t-x)h(x)dx - \int \psi_{\sigma_n} \star f(x)h(x)dx$. Using the triangle inequality,

$$|g(t)| \leq |f(t)| \int_{[0,1]^d} |\psi_{\sigma_n}(t-x)h(x)|dx + \int_{[0,1]^d} |\psi_{\sigma_n} \star f(x)||h(x)|dx.$$

Using $\|h\|_\infty \leq 1$, the first term in the above display can be bounded above by $C_1\|f\|_\infty$ where $C_1 = \int |\psi(t)|dt$. Similarly, the second term can be bounded above by $\|\psi_{\sigma_n} \star f\|_1 \leq \|\psi_{\sigma_n} \star f\|_\infty \leq \|f\|_\infty + \epsilon_n$, where the final inequality follows from (BS). Noting that for any $f \in \mathcal{P}_n$, $\|f\|_\infty \leq 2M_n$ (since the Hilbert space norm is stronger than the $\|\cdot\|_\infty$ norm), we have $K_1 \leq CM_n$.

Next we bound $\sigma_{\mathcal{G}}^2 = \sup_{g \in \mathcal{G}} \int_{[0,1]^d} g(t)^2 dt$. Fix $g \in \mathcal{G}$. Using the expression for $g(t)$ in the previous paragraph, $|g(t)| \leq |f(t)| \int |\psi_{\sigma_n}(x-t)|dx + \int |\psi_{\sigma_n} \star f(x)|dx$. As before, we can bound $\int |\psi_{\sigma_n}(x-t)|dx$ from above by C_1 and also $\int |\psi_{\sigma_n} \star f(x)|dx \leq C_1 \int_{s \in [0,1]^d} |f(s)|ds$. Using $(|a| + |b|)^2 \leq 2(|a|^2 + |b|^2)$ and the Cauchy–Schwarz inequality, $|g(t)|^2 \leq C|f(t)|^2 + C\{\int_{s \in [0,1]^d} |f(s)|ds\}^2 \leq C\{|f(t)|^2 + \|f\|_2^2\}$. Thus, we have $\sigma_{\mathcal{G}}^2 \leq C\|f\|_2^2$ for some absolute constant C . Using the bound for $\sup_{f \in \mathcal{P}_n} \|f\|_2^2$ in the following Proposition 7, we conclude that $\sigma_{\mathcal{G}}^2 \leq C$ for some absolute constant $C > 0$.

Proposition 7 *Recall \mathcal{P}_n from (5). Then, $\sup_{f \in \mathcal{P}_n} \|f\|_2^2 \leq C$ for some absolute constant $C > 0$.*

Proof Let $f \in \mathcal{P}_n$. Then, there exists $h \in \mathbb{H}^{a_n}$ with $\|h\|_{\mathbb{H}^{a_n}} \leq M_n$ such that $\|f-h\|_\infty \leq \epsilon_n$. Hence, $\|f\|_2^2 \leq 2(\|h\|_2^2 + \epsilon_n^2)$ and it is enough to bound $\|h\|_2^2$. Recalling that $\|\cdot\|_{2,d}$ denotes the L_2 norm of \mathbb{R}^d , we have $\|h\|_2^2 \leq \|h\|_{2,d}^2$. We provide a bound for $\|h\|_{2,d}^2$ below.

There exists $\psi \in L_2(\mu_{a_n})$ such that $h(t) = \int \exp(i\langle \lambda, t \rangle)\xi(\lambda)\omega_{a_n}(\lambda)d\lambda$. Letting \hat{h} denote the Fourier transform of h , one has from the Fourier inversion theorem that $\hat{h}(\lambda) =$

$\xi(-\lambda)\omega_{a_n}(\lambda)$. By Parseval's theorem, $\|h\|_{2,d}^2 = \|\hat{h}\|_{2,d}^2 = \int \xi^2(\lambda)\omega_{a_n}^2(\lambda)d\lambda^1$. Observe that $\omega_{a_n}^2(\lambda) = a_n^{-2d} \exp\{-\|\lambda\|^2/(2a_n^2)\}/C^2$, where $C = 2^d\pi^{d/2}$. Hence,

$$\begin{aligned} \|h\|_{2,d}^2 &= \frac{a_n^{-2d}}{C^2} \int \xi^2(\lambda) \exp\{-\|\lambda\|^2/(2a_n^2)\}d\lambda \leq \frac{a_n^{-2d}}{C^2} \int \xi^2(\lambda) \exp\{-\|\lambda\|^2/(4a_n^2)\}d\lambda \\ &= \frac{a_n^{-d}}{C} \int \xi^2(\lambda)\omega_{a_n}(\lambda)d\lambda = \frac{\|h\|_{\mathbb{H}^{a_n}}^2}{Ca_n^d} \leq \frac{M_n^2}{Ca_n^d} = \frac{1}{C}, \end{aligned}$$

since $\|h\|_{\mathbb{H}^{a_n}}^2 = \left\| \xi\omega_{a_n}^{1/2} \right\|_{2,d}^2$ and $M_n = a_n^{d/2}$. ■

Finally, we proceed to bound $\mathbb{E}_X W$, where $W = \int_{[0,1]^d} |\sum_{i=1}^n L_x(X_i)|dx$. Using Jensen's inequality and the integral version of Minkowski's inequality, one has

$$\begin{aligned} \mathbb{E}_X W &\leq (\mathbb{E}_X W^2)^{1/2} = \left[\int_{\prod_{i=1}^n [0,1]^d} \left\{ \int_{[0,1]^d} \left| \sum_{i=1}^n L_x(t_i) \right| dx \right\}^2 dt_1 \dots dt_n \right]^{1/2} \\ &\leq \int_{[0,1]^d} \left\{ \int_{\prod_{i=1}^n [0,1]^d} \sum_{i=1}^n |L_x(t_i)|^2 dt_1 \dots dt_n \right\}^{1/2} dx. \end{aligned}$$

Clearly, $\int_{\prod_{i=1}^n [0,1]^d} |\sum_{i=1}^n L_x(t_i)|^2 dt_1 \dots dt_n = \text{Var}_X \{ \sum_{i=1}^n L_x(X_i) \} = n \text{Var}_X \{ L_x(X_1) \}$, since $\mathbb{E}_X L_x(X_1) = 0$. Further, $\text{Var}_X \{ L_x(X_1) \} \leq \mathbb{E}_X \{ \psi_{\sigma_n}(x - X_1) f(X_1) \}^2 = \int_{[0,1]^d} \psi_{\sigma_n}(x - t)^2 f(t)^2 dt \leq \frac{1}{\sigma_n^d} \psi_{\sigma_n} \star f^2$. Substituting this in the above display

$$\begin{aligned} \mathbb{E}_X W &\leq \left(\frac{n}{\sigma_n^d} \right)^{1/2} \int_{[0,1]^d} \{ \psi_{\sigma_n} \star f^2(x) \}^{1/2} dx \\ &\leq \left(\frac{n}{\sigma_n^d} \right)^{1/2} \left\{ \int_{[0,1]^d} |\psi_{\sigma_n}| \star f^2(x) dx \right\}^{1/2} \\ &\leq \left(\frac{n}{\sigma_n^d} \right)^{1/2} \left\{ \int_{[0,1]^d} \int_{[0,1]^d} |\psi_{\sigma_n}(x-t)| f^2(t) dt dx \right\}^{1/2} \\ &\leq \left(\frac{n}{\sigma_n^d} \right)^{1/2} \left[\int_{[0,1]^d} f^2(t) \left\{ \int_{\mathbb{R}^d} |\psi_{\sigma_n}(x-t)| dx \right\} dt \right]^{1/2} \\ &\leq C \left(\frac{n}{\sigma_n^d} \right)^{1/2} = C n^{\frac{\alpha+d}{2\alpha+d}} \log^{t_2 d/2} n \leq C n \epsilon_n. \end{aligned}$$

From the penultimate line to the last line of the above display, we invoked Proposition 7 to bound $\|f\|_2$ by a constant. We have thus obtained $K_1 \leq C M_n$ and $K_2 \leq C n$. In Proposition 6, set $t = n\epsilon_n^2$. We have $K_1 t \leq C(n\epsilon_n M_n)\epsilon_n \leq n\epsilon_n$ for sufficiently large n provided $\alpha > d/2$. Further, $K_2 t \leq n^2 \epsilon_n^2 + K_1 \mathbb{E}_P(W)t = n^{\frac{2\alpha+2d}{2\alpha+d}} \log^{3t_1} n + n^{\frac{\alpha+2d+d/2}{2\alpha+d}} \log^{2t_1+t_2 d/2} n \leq 2n^{\frac{2\alpha+2d}{2\alpha+d}} \log^{3t_1} n$ for sufficiently large n if $\alpha > d/2$. Therefore, $(K_2 t)^{1/2} \leq n\epsilon_n$. ■

1. ω_{a_n} is symmetric about zero.

We next show that (BS) holds. Fix $f \in \mathcal{P}_n \cap U_n$. Since $f \in \mathcal{P}_n$, there exists $h \in \mathbb{H}^{a_n}$ with $\|h\|_{\mathbb{H}^{a_n}} \leq M_n$ such that $\|f - h\|_\infty \leq \epsilon_n$. By the triangle inequality, $\|\psi_{\sigma_n} \star f - f\|_1 \leq \|\psi_{\sigma_n} \star f - \psi_{\sigma_n} \star h\|_1 + \|\psi_{\sigma_n} \star h - h\|_1 + \|h - f\|_1$. Using $\|\psi_{\sigma_n} \star g\|_1 \leq \|g\|_1$ for any L_1 function g , we can further bound $\|\psi_{\sigma_n} \star f - f\|_1$ from above by $2\epsilon_n + \|\psi_{\sigma_n} \star h - h\|_\infty$. It thus remains to show that $\|\psi_{\sigma_n} \star h - h\|_\infty \leq \epsilon_n$.

There exists $\xi \in L_2(\mu_{a_n})$ such that $h(t) = \int \exp(i \langle \lambda, t \rangle) \xi(\lambda) \omega_{a_n}(\lambda) d\lambda$. Clearly, $\hat{h}(\lambda) = \xi(-\lambda) \omega_a(\lambda)$. Since the Fourier transform of $(\psi_{\sigma_n} \star h)$ is $(2\pi)^d \hat{\psi}_{\sigma_n} \hat{h}$ and $\hat{\psi}_{\sigma_n}(\lambda) = \hat{\psi}(\sigma_n \lambda)$, we have $\psi_{\sigma_n} \star h(t) = (2\pi)^d \int \exp(-i \langle \lambda, t \rangle) \hat{\psi}(\sigma_n \lambda) \hat{h}(\lambda) d\lambda$. We can choose ψ in a manner such that $\hat{\psi}$ is compactly supported, equals $(2\pi)^{-d}$ on $[-1, 1]^d$ and is bounded above by this constant everywhere; see proof of Lemma 4.3 in van der Vaart and van Zanten (2009). Putting together,

$$\begin{aligned} |\psi_{\sigma_n} \star h(t) - h(t)|^2 &\leq \left\{ \int_{\|\lambda\| > \sigma_n^{-1}} |\hat{h}(\lambda)| \right\}^2 \leq \left\{ \int \xi(\lambda)^2 \omega_{a_n}(\lambda) d\lambda \right\} \int_{\|\lambda\| > \sigma_n^{-1}} \omega_{a_n}(\lambda) d\lambda \\ &\leq C \|h\|_{\mathbb{H}^{a_n}}^2 \exp\{-\sigma_n^{-2}/(4a_n^2)\} \leq C M_n^2 \exp\{-\sigma_n^{-2}/(4a_n^2)\} = C a_n^d \exp\{-\log^{2t_2} n/4\}, \end{aligned}$$

where C is an absolute constant. The proof follows by noting that $C a_n^d \exp\{-\log^{2t_2} n/4\} \leq \epsilon_n^2$ whenever $t_2 > 1/2$ (holds for $t_2 = 1/(2 - \kappa)$, for $0 < \kappa < 1$).

3.2 Proof of Corollary 2

Case $\beta < \alpha$: Setting $\sigma_n = n^{-1/(2\beta+d)} \log^{-t_2} n$ for some constant $t_2 \geq 1/(2 - \kappa)$, for $0 < \kappa < 2$, $M_n = a_n^{d/2}$, \tilde{f}_n same as in (4), $\mathcal{P}_n = M_n \mathbb{H}_1^{a_n} + \epsilon_n \mathbb{B}_1$ with $\epsilon_n = n^{-\beta/(2\beta+d)} \log^{3t_1/2} n$, $t_1 \geq t_2 d/2$, and $\delta_n = \epsilon_n = \epsilon'_n$, one can verify (PCS), (BT), (DT), (BS), (DS) exactly as in the proof of Theorem 1. (PCN) follows from Lemma 4.3 of van der Vaart and van Zanten (2009).

Case $\beta > \alpha$: Same as before with $\epsilon_n = n^{-\alpha/(2\beta+d)} \log^{3t_1/2} n$ for $t_1 \geq t_2 d/2$.

4. Discussion

The article extends upon previous results on random design regression using Gaussian process priors. A limitation of the current exposition is the requirement of the knowledge of the smoothness parameter to construct the rescaling sequence. A natural question is whether one can find a suitable prior on the bandwidth parameter which adapts to the unknown smoothness level as in the fixed design case in van der Vaart and van Zanten (2009). We propose to resolve this issue as a part of future research. Also, our current proof technique would lead to a sub-optimal rate of posterior convergence for L_p norms with $p \neq 1$. We believe this is due to the use of Talagrand’s inequality to construct the test function. A key requirement to obtain optimal convergence rate is that the *variance term* $\sigma_{\mathcal{F}}^2$ in the application of Talagrand’s inequality should be at most $O(n)$. This assertion is true only when $p = 1$. Obtaining convergence rates for integrated L_p norms with $p \neq 1$ is a topic of future research.

Acknowledgement

Dr. Pati and Dr. Bhattacharya acknowledge support for this project from the Office of Naval Research (ONR BAA 14-0001). Dr. Cheng acknowledges support from the National Science Foundation (NSF CAREER, DMS – 1151692, DMS – 1418042), Simons Foundation (305266) and warm hosting at SAMSI.

Appendix A. Proof of Theorem 3

Let $\|f\|_{2,n}$ denote the empirical L_2 norm of f , so that $\|f\|_{2,n}^2 = n^{-1} \sum_{i=1}^n f^2(X_i)$. Also, define

$$L_n(f, f_0) = \frac{q^{(n)}(Y_{1:n}, X_{1:n} | f)}{q^{(n)}(Y_{1:n}, X_{1:n} | f_0)}.$$

Lemma 8 *Let A_n denote the following event in the sigma-field generated by $(Y_{1:n}, X_{1:n})$:*

$$A_n = \left\{ (Y_{1:n}, X_{1:n}) : \int L_n(f, f_0) \Pi(df) \geq e^{-n\delta_n^2} \Pi(\|f - f_0\|_\infty \leq \delta_n) \right\}. \quad (11)$$

Then, $\mathbb{P}_{Y,X}^{f_0}(A_n) \geq 1 - e^{-Cn\delta_n^2}$.

Proof Clearly, $\mathbb{P}_{Y,X}^{f_0}(A_n) = \mathbb{E}_X^{f_0}[\mathbb{P}_{Y|X}^{f_0}(A_n)]$. By Lemma 14 of van der Vaart and van Zanten (2011), $\mathbb{P}_{Y|X}^{f_0}\{\int L_n(f, f_0) \Pi(df) \geq e^{-n\delta_n^2} \Pi(\|f - f_0\|_{2,n} \leq \delta_n)\} \geq 1 - e^{-n\delta_n^2/8}$. The conclusion follows by noting that $\Pi(\|f - f_0\|_\infty < \delta_n) \leq \Pi(\|f - f_0\|_{2,n} < \delta_n)$. ■

Lemma 9 *There exists a test function Φ_n for $H_0 : f = f_0$ vs $H_1 : f \in U_n \cap \mathcal{P}_n$ such that*

$$\mathbb{E}_{Y,X}^{f_0} \Phi_n \leq e^{-Cn\delta_n^2}, \quad (12)$$

$$\sup_{f \in U_n \cap \mathcal{P}_n} \mathbb{E}_{Y,X}^f (1 - \Phi_n) \leq e^{-Cn\delta_n^2}. \quad (13)$$

for some absolute constant C .

Proof Let $\Phi_n = 1(\|\tilde{f}_n - f_0\|_1 > M\epsilon_n/2)$. The error bounds follow from (BT), (DT) and (BS), (DS). ■

Using a standard line of argument for establishing convergence rates in Bayesian non-parametric models (Ghosal et al., 2000), we have $\mathbb{E}_{Y,X}^{f_0} \Pi(U_n | Y_{1:n}, X_{1:n}) \leq \sum_{i=1}^4 b_{in}$, where $b_{1n} = \mathbb{E}_{Y,X}^{f_0} \Phi_n$, $b_{2n} = e^{n\delta_n^2} \sup_{f \in U_n \cap \mathcal{P}_n} \mathbb{E}_{Y,X}^f (1 - \Phi_n) / \Pi(\|f - f_0\|_\infty < \delta_n)$, $b_{3n} = e^{n\delta_n^2} \Pi(\mathcal{P}_n^c) / \Pi(\|f - f_0\|_\infty < \delta_n)$ and $b_{4n} = \mathbb{P}_{Y,X}^{f_0}(A_n^c)$. The Theorem then follows from Lemmas 8, 9 and Conditions (PCS) and (PCN).

References

- R. J. Adler. *An Introduction to Continuity, Extrema, and Related Topics for General Gaussian processes*, volume 12. Institute of Mathematical Statistics, 1990.
- Y. Baraud. Model selection for regression on a random design. *ESAIM: Probability and Statistics*, 6:127–146, 2002.
- A. Bhattacharya, D. Pati, and D. B. Dunson. Anisotropic function estimation using multi-bandwidth gaussian processes. *The Annals of Statistics*, 42(1):352–381, 2014.
- L. Birgé. Sur un théorème de minimax et son application aux tests. *Probability and Mathematical Statistics*, 3:259–282, 1984.
- L. Birgé. Model selection for gaussian regression with random design. *Bernoulli*, 10(6):1039–1051, 2004.
- D. Bontemps. Bernstein-von Mises theorems for Gaussian regression with increasing number of regressors. *The Annals of Statistics*, 39(5):2557–2584, 2011.
- O. Bousquet. Concentration inequalities for sub-additive functions using the entropy method. *Progress in Probability*, pages 213–248, 2003.
- L. D. Brown, T. T. Cai, M. G. Low, and C-H. Zhang. Asymptotic equivalence theory for nonparametric regression with random design. *The Annals of Statistics*, pages 688–707, 2002.
- T. Choi and M.J. Schervish. On posterior consistency in nonparametric regression problems. *Journal of Multivariate Analysis*, 98(10):1969–1987, 2007.
- S. Ghosal and A. Roy. Posterior consistency of Gaussian process prior for nonparametric binary regression. *The Annals of Statistics*, 34(5):2413–2429, 2006.
- S. Ghosal, J. K. Ghosh, and A. W. van der Vaart. Convergence rates of posterior distributions. *The Annals of Statistics*, 28(2):500–531, 2000.
- E. Giné and R. Nickl. Rates of contraction for posterior distributions in L^r -metrics, $1 \leq r \leq \infty$. *The Annals of Statistics*, 39(6):2883–2911, 2011.
- B. J. K. Kleijn and A. W. van der Vaart. Misspecification in infinite-dimensional Bayesian statistics. *The Annals of Statistics*, 34(2):837–877, 2006.
- L. Le Cam. *Asymptotic methods in statistical decision theory*. New York, 1986.
- D. Pati, A. Bhattacharya, N. S. Pillai, and D. B. Dunson. Posterior contraction in sparse Bayesian factor models for massive covariance matrices. *The Annals of Statistics*, 42(3):1102–1130, 2014.
- C. E. Rasmussen. Gaussian processes in machine learning. *Advanced Lectures on Machine Learning*, pages 63–71, 2004.

- C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- K. Ray. Bayesian inverse problems with non-conjugate priors. *Electronic Journal of Statistics*, 7:2516–2549, 2013.
- M. Seeger. Gaussian processes for machine learning. *International Journal of Neural Systems*, 14(02):69–106, 2004.
- M. Seeger, S. M. Kakade, and D. P. Foster. Information consistency of nonparametric gaussian process methods. *IEEE Transactions on Information Theory*, 54(5):2376–2382, 2008.
- Z. Shang and G. Cheng. Nonparametric Bernstein-von Mises phenomenon: A tuning prior perspective. *ArXiv Preprint ArXiv:1411.3686*, 2014.
- B. T. Szabó, A. W. van der Vaart, and J. H. van Zanten. Empirical bayes scaling of gaussian priors in the white noise model. *Electronic Journal of Statistics*, 7:991–1018, 2013.
- S. T. Tokdar and J. K. Ghosh. Posterior consistency of logistic Gaussian process priors in density estimation. *Journal of Statistical Planning and Inference*, 137(1):34–42, 2007.
- A. W. van der Vaart and J. H. van Zanten. Bayesian inference with rescaled Gaussian process priors. *Electronic Journal of Statistics*, 1:433–448, 2007. ISSN 1935-7524.
- A. W. van der Vaart and J. H. van Zanten. Rates of contraction of posterior distributions based on gaussian process priors. *The Annals of Statistics*, 36(3):1435–1463, 2008a.
- A. W. van der Vaart and J. H. van Zanten. Reproducing kernel Hilbert spaces of Gaussian priors. *IMS Collections*, 3:200–222, 2008b.
- A. W. van der Vaart and J. H. van Zanten. Adaptive Bayesian estimation using a Gaussian random field with inverse Gamma bandwidth. *The Annals of Statistics*, 37(5B):2655–2675, 2009.
- A. W. van der Vaart and J. H. van Zanten. Information rates of nonparametric Gaussian process methods. *Journal of Machine Learning Research*, 12:2095–2119, 2011.