

A probabilistic Taylor expansion with Gaussian processes

Anonymous authors

Paper under double-blind review

Abstract

We study a class of Gaussian processes for which the posterior mean, for a particular choice of data, replicates a truncated Taylor expansion of any order. The data consist of derivative evaluations at the expansion point and the prior covariance kernel belongs to the class of *Taylor kernels*, which can be written in a certain power series form. We discuss and prove some results on maximum likelihood estimation of parameters of Taylor kernels. The proposed framework is a special case of Gaussian process regression based on data that is orthogonal in the reproducing kernel Hilbert space of the covariance kernel.

1 Introduction

Taylor’s theorem is among the most fundamental results in analysis. In one dimension, Taylor’s theorem states that any function $f: \mathbb{R} \rightarrow \mathbb{R}$ that is sufficiently smooth at $a \in \mathbb{R}$ can be written as

$$f(x) = T_{n,a}(x) + P_{n,a}(x), \quad (1.1)$$

where $T_{n,a}(x) = \sum_{p=0}^n \frac{1}{p!} f^{(p)}(a)(x-a)^p$ is the n th order Taylor polynomial and $P_{n,a}(x)$ a remainder term which has the property that $P_{n,a}(x) = \mathcal{O}(|x-a|^{n+1})$ as $|x-a| \rightarrow 0$. Multidimensional generalisations are readily available and will be introduced in Section 2. Approximations derived from (1.1), in particular the first and second order Taylor approximations

$$f(x) \approx f(a) + f'(a)(x-a) \quad \text{and} \quad f(x) \approx f(a) + f'(a)(x-a) + \frac{1}{2}f''(a)(x-a)^2,$$

play an important role in numerical algorithms for a number of reasons. Firstly, Taylor approximations provide a straightforward and principled means of *linearising* a function of interest, which can often dramatically accelerate otherwise costly computations. Secondly, they require only information about a function and its derivatives at a single point; information that particular algorithms may already be collecting. Particular applications of Taylor’s theorem in numerical algorithms include optimisation (Moré, 1978; Conn et al., 2000), state estimation (Särkkä, 2013, Ch. 5), ordinary differential equations (Hairer et al., 1993, Ch. II), and approximation of exponential integrals in Bayesian statistics (Raudenbush et al., 2000), to name but a few. A crucial challenge when applying Taylor series in this way, however, is their locality. The approximation is valid only near a , and apart from trivial examples approximation quality decays rapidly away from this point. When a numerical algorithm attempts to use a Taylor approximation to explore function behaviour around a particularly novel point, far from a , the behaviour of the algorithm can be difficult to predict and control.

This paper proposes a remedy for this by introducing a Gaussian process (GP) model (Rasmussen and Williams, 2006) whose posterior mean, given the *derivative data* $(f(a), f'(a), \dots, f^{(n)}(a))$, is exactly the Taylor polynomial $T_{n,a}$, and whose posterior variance plays a role analogous to the remainder term $P_{n,a}$. In the spirit of probabilistic numerics (Diaconis, 1988; Cockayne et al., 2019b; Hennig et al., 2022), the posterior variance can then be used for principled probabilistic quantification of epistemic uncertainty in the Taylor approximation $f(x) \approx T_{n,a}(x)$ at $x \neq a$, which can be exploited and propagated forward. In effect, the variance may be used to encode into algorithms a degree of “scepticism” about the validity of the Taylor approximation away from a . Taylor approximation thus joins the ranks of classical numerical methods, such as algorithms for spline interpolation (Diaconis, 1988; Kimeldorf and Wahba, 1970), numerical quadrature (Diaconis, 1988; Karvonen and Särkkä, 2017; Karvonen et al., 2018), differential equations (Schober et al., 2014; 2019; Teymur et al., 2016), and linear algebra (Cockayne et al., 2019a; Hennig, 2015), that can be cast as statistical inference.

1.1 Related Literature

The Gaussian process priors we use to construct a probabilistic Taylor expansion are determined by positive-definite and non-stationary *Taylor kernels* which, at an expansion point a , take the form

$$K_a(x, y) = K(x - a, y - a), \quad \text{where} \quad K(x, y) = \sigma^2 \sum_{p=0}^{\infty} \frac{c_p \lambda^p}{(p!)^2} (xy)^p \quad (1.2)$$

for non-negative constants c_p and positive parameters σ and λ . For multidimensional input points, the index $p \in \mathbb{N}_0$ is replaced with a multi-index $\alpha \in \mathbb{N}_0^d$ and xy usually with the Euclidean inner product; see Section 2.1 for details. The canonical example is the *exponential kernel*

$$K(x, y) = \sigma^2 \sum_{p=0}^{\infty} \frac{\lambda^p}{p!} (xy)^p = \sigma^2 \exp(\lambda xy), \quad (1.3)$$

which is obtained by setting $c_p = p!$ in (1.2). The exponential kernel is closely connected to the popular Gaussian kernel $K(x, y) = \sigma^2 \exp(-\lambda^2(x - y)^2/2)$.

Taylor kernels, often under the name *power series kernels*, have been used—and their approximation properties analysed—in the numerical analysis and scattered data approximation literature; see Dick (2006); Zwicknagl (2009); De Marchi and Schaback (2010); Zwicknagl and Schaback (2013) and Fasshauer and McCourt (2015, Sec. 3.3.1). Section 1 in Zwicknagl and Schaback (2013) has been of particular inspiration for the present work. The *Szegő kernel* and the *Bergman kernel* $K(x, y) = 1/(1 - xy)$ and $K(x, y) = 1/(1 - xy)^2$, respectively, are particularly well studied in the approximation theory literature because their reproducing kernel Hilbert spaces (RKHSs) are important in complex analysis; see, for example, Larkin (1970); Richter-Dyn (1971b;a) and Oettershagen (2017, Sec. 6.2). These kernels are defined on the open interval $(-1, 1)$ and obtained from (1.2) by setting (Szegő) $c_p = (p!)^2$ and (Bergman) $c_{2p} = (p!)^2$ and $c_{2p+1} = 0$ for $p \in \mathbb{N}_0$ in (1.2).

Taylor kernels occasionally appear in the machine learning and statistics literature but, to the best of our knowledge, have not been used in conjunction with derivative data in the way proposed here. We refer to Minka (2000, Sec. 4); Steinwart and Christmann (2008, Example 4.9) and (Liang and Rakhlin, 2020) for a few of their appearances. Gaussian process regression with derivative data has been explored (e.g., Solak et al., 2002; Prüher and Särkkä, 2016; Wu et al., 2017; Eriksson et al., 2018), though typically for “standard” kernels such as the Gaussian kernel.

1.2 Contributions

The main contributions of the paper are contained in Sections 2 and 3. In Section 2, we derive a probabilistic Taylor expansion and a basic error bound; these results are given in Theorems 2.1 and 2.2. We also discuss how inclusion of observation noise affects probabilistic Taylor expansions. In Section 3, we derive expressions for maximum likelihood estimates of the Taylor kernel parameters σ and λ . Perhaps the most interesting result that we obtain is Theorem 3.2, which states that derivative data that could have been generated by a constant function yield the estimate $\lambda_{\text{ML}} = 0$. As mentioned above, the exponential kernel is related to the Gaussian kernel. In Section 4, we show how to derive closed form expression for the posterior mean and covariance given derivative data when the covariance kernel is Gaussian. Section 5 outlines generalisations of probabilistic Taylor expansions derived in Section 2 for data that are orthogonal in the reproducing kernel Hilbert space of the covariance kernel. Fourier coefficients constitute an example of orthogonal data when the kernel is periodic. Some simple numerical examples are included in Section 6.

1.3 Notation

We use \mathbb{N}_0 to denote the set of non-negative integers and \mathbb{N}_0^d to denote the collection of non-negative d -dimensional multi-indices $\alpha = (\alpha(1), \dots, \alpha(d))$, where $\alpha(j) \in \mathbb{N}_0$ is the j th index of α . We also use the standard notation $|\alpha| = \alpha(1) + \dots + \alpha(d)$ and $\alpha! = \alpha(1)! \times \dots \times \alpha(d)!$.

2 A Probabilistic Taylor Expansion

In this section, we derive a probabilistic Taylor expansion using Gaussian processes. We discuss a generalisation of this derivation for orthogonal data in Section 5.

2.1 Taylor Kernels

Let $\mathbf{a} \in \mathbb{R}^d$ and $r \in (0, \infty]$. Define $\Omega_{\mathbf{a},r} = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x} - \mathbf{a}\|_2 < r\}$. A *multidimensional Taylor kernel* on $\Omega_{\mathbf{a},r} \times \Omega_{\mathbf{a},r}$ is defined as

$$K_{\mathbf{a}}(\mathbf{x}, \mathbf{y}) = K(\mathbf{x} - \mathbf{a}, \mathbf{y} - \mathbf{a}) \quad \text{for} \quad K(\mathbf{x}, \mathbf{y}) = \sigma^2 \sum_{\boldsymbol{\alpha} \in \mathbb{N}_0^d} \frac{c_{\boldsymbol{\alpha}} \boldsymbol{\lambda}^{\boldsymbol{\alpha}}}{(\boldsymbol{\alpha}!)^2} \mathbf{x}^{\boldsymbol{\alpha}} \mathbf{y}^{\boldsymbol{\alpha}}, \quad (2.1)$$

where $\sigma > 0$ and $\boldsymbol{\lambda} \in \mathbb{R}_+^d$ are kernel hyperparameters. The coefficients $c_{\boldsymbol{\alpha}}$ are non-negative constants such that $c_{\boldsymbol{\alpha}} > 0$ for infinitely many $\boldsymbol{\alpha} \in \mathbb{N}_0^d$ and

$$\sum_{\boldsymbol{\alpha} \in \mathbb{N}_0^d} \frac{c_{\boldsymbol{\alpha}} \boldsymbol{\lambda}^{\boldsymbol{\alpha}}}{(\boldsymbol{\alpha}!)^2} r^{2|\boldsymbol{\alpha}|} < \infty \quad \text{if} \quad r < \infty \quad \text{or} \quad \sum_{\boldsymbol{\alpha} \in \mathbb{N}_0^d} \frac{c_{\boldsymbol{\alpha}} \boldsymbol{\lambda}^{\boldsymbol{\alpha}}}{\boldsymbol{\alpha}! \sqrt{\boldsymbol{\alpha}!}} e^{|\boldsymbol{\alpha}|} < \infty \quad \text{if} \quad r = \infty. \quad (2.2)$$

The conditions (2.2) are sufficient to ensure the series defining $K_{\mathbf{a}}$ via (2.1) converges absolutely for all $\mathbf{x}, \mathbf{y} \in \Omega_{\mathbf{a},r}$, which, together with $c_{\boldsymbol{\alpha}} > 0$ for infinitely many $\boldsymbol{\alpha}$, guarantees that Taylor kernels are positive-definite (Zwinnagl and Schaback, 2013, Thm. 2.2). However, $c_{\boldsymbol{\alpha}} \geq 0$ is not necessary for positive-definiteness of K in (2.1) (see Zwinnagl, 2009, Sec. 2).

An important subclass of Taylor kernels are *inner product kernels*, defined by

$$K(\mathbf{x}, \mathbf{y}) = \sigma^2 \sum_{p=0}^{\infty} \frac{c_p}{(p!)^2} \langle \mathbf{x}, \mathbf{y} \rangle_{\boldsymbol{\lambda}}^p, \quad \text{where} \quad \langle \mathbf{x}, \mathbf{y} \rangle_{\boldsymbol{\lambda}} = \sum_{i=1}^d \lambda_i x_i y_i. \quad (2.3)$$

It is easy to show that inner product kernels are Taylor kernels: From the multinomial theorem we have

$$\begin{aligned} K(\mathbf{x}, \mathbf{y}) &= \sigma^2 \sum_{p=0}^{\infty} \frac{c_p}{(p!)^2} \langle \mathbf{x}, \mathbf{y} \rangle_{\boldsymbol{\lambda}}^p = \sigma^2 \sum_{p=0}^{\infty} \frac{c_p}{(p!)^2} \left(\sum_{i=1}^d \lambda_i x_i y_i \right)^p = \sigma^2 \sum_{p=0}^{\infty} \frac{c_p}{(p!)^2} \sum_{|\boldsymbol{\alpha}|=p} \frac{p!}{\boldsymbol{\alpha}!} \boldsymbol{\lambda}^{\boldsymbol{\alpha}} \mathbf{x}^{\boldsymbol{\alpha}} \mathbf{y}^{\boldsymbol{\alpha}} \\ &= \sigma^2 \sum_{\boldsymbol{\alpha} \in \mathbb{N}_0^d} \frac{c_{|\boldsymbol{\alpha}|}}{\boldsymbol{\alpha}! |\boldsymbol{\alpha}!} \boldsymbol{\lambda}^{\boldsymbol{\alpha}} \mathbf{x}^{\boldsymbol{\alpha}} \mathbf{y}^{\boldsymbol{\alpha}}, \end{aligned}$$

which we recognise as a Taylor kernel in (2.1) with $c_{\boldsymbol{\alpha}} = c_{|\boldsymbol{\alpha}|} \boldsymbol{\alpha}! / |\boldsymbol{\alpha}!|$. We will discuss estimation of the parameters σ and $\boldsymbol{\lambda}$ (as well as the coefficients c_p) in Section 3; for now, we assume the parameters are given and proceed to show how Taylor kernels may be used to derive a probabilistic Taylor expansion.

The multidimensional version of the exponential kernel in (1.3) is

$$K(\mathbf{x}, \mathbf{y}) = \sigma^2 \exp(\langle \mathbf{x}, \mathbf{y} \rangle_{\boldsymbol{\lambda}}) = \sigma^2 \sum_{p=0}^{\infty} \frac{1}{p!} \langle \mathbf{x}, \mathbf{y} \rangle_{\boldsymbol{\lambda}}^p. \quad (2.4)$$

The exponential kernel is defined on $\Omega_{\mathbf{a},r} = \mathbb{R}^d$. In Section 4, we discuss a close connection that the exponential kernel has to the commonly used Gaussian kernel. By setting $c_p = 1$ we obtain the *Bessel kernel* $K(\mathbf{x}, \mathbf{y}) = \sigma^2 \sum_{p=0}^{\infty} \langle \mathbf{x}, \mathbf{y} \rangle_{\boldsymbol{\lambda}}^p / (p!)^2 = \text{I}_0(2\langle \mathbf{x}, \mathbf{y} \rangle_{\boldsymbol{\lambda}}^{1/2})$, where I_0 is the modified Bessel function of the first kind, which is another Taylor kernel defined on the whole of \mathbb{R}^d .

2.2 Gaussian Process Regression with Derivative Data

A Gaussian process $f_{\text{GP}} \sim \text{GP}(m, R)$ characterised by mean function $m: \Omega_{\mathbf{a},r} \rightarrow \mathbb{R}$ and covariance kernel $R: \Omega_{\mathbf{a},r} \times \Omega_{\mathbf{a},r} \rightarrow \mathbb{R}$ is a stochastic process such that $\mathbb{E}[f_{\text{GP}}(\mathbf{x})] = m(\mathbf{x})$ and $\text{Cov}[f(\mathbf{x}), f(\mathbf{y})] = R(\mathbf{x}, \mathbf{y})$ for

all $\mathbf{x}, \mathbf{y} \in \Omega_{\mathbf{a},r}$ (Rasmussen and Williams, 2006). For any points $\mathbf{x}_1, \dots, \mathbf{x}_N \in \Omega_{\mathbf{a},r}$ the joint distribution of $(f_{\text{GP}}(\mathbf{x}_1), \dots, f_{\text{GP}}(\mathbf{x}_N))$ is an N -dimensional Gaussian with mean vector $(m(\mathbf{x}_1), \dots, m(\mathbf{x}_N)) \in \mathbb{R}^N$ and covariance matrix $\mathbf{R} = (R(\mathbf{x}_i, \mathbf{x}_j))_{i,j=1}^N \in \mathbb{R}^{N \times N}$. Let $f: \Omega_{\mathbf{a},r} \rightarrow \mathbb{R}$ be an n times differentiable function on $\Omega_{\mathbf{a},r}$, meaning that the partial derivatives

$$\mathbf{D}^{\boldsymbol{\alpha}} f = \frac{\partial^{|\boldsymbol{\alpha}|} f}{\partial x_1^{\alpha(1)} \dots \partial x_d^{\alpha(d)}}$$

exist for all $\boldsymbol{\alpha} \in \mathbb{N}_0^d$ such that $|\boldsymbol{\alpha}| \leq n$. Suppose also that the prior mean m is n times differentiable and that R is n times differentiable in both arguments. The *noiseless* derivative data are

$$\mathbf{f}_{\mathbf{a}} = (\mathbf{D}^{\boldsymbol{\alpha}} f(\mathbf{a}))_{|\boldsymbol{\alpha}| \leq n} = (\mathbf{D}^{\boldsymbol{\alpha}_1} f(\mathbf{a}), \dots, \mathbf{D}^{\boldsymbol{\alpha}_{N_n^d}} f(\mathbf{a})),$$

where we use an arbitrary ordering of the set $\{\boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_{N_n^d}\} = \{\boldsymbol{\alpha} \in \mathbb{N}_0^d : |\boldsymbol{\alpha}| \leq n\}$, which contains

$$N_n^d = \binom{n+d}{n} = \frac{(n+d)!}{n!d!}$$

elements. When conditioned on these derivative data, the posterior $f_{\text{GP}} | \mathbf{f}_{\mathbf{a}}$ is a Gaussian process (Särkkä, 2011; Travelletti and Ginsbourger, 2022). That is, $f_{\text{GP}} | \mathbf{f}_{\mathbf{a}} \sim \text{GP}(s_{n,\mathbf{a}}, P_{n,\mathbf{a}})$ with mean and covariance

$$s_{n,\mathbf{a}}(\mathbf{x}) = m(\mathbf{x}) + \mathbf{r}_{\mathbf{a}}(\mathbf{x})^{\top} \mathbf{R}_{\mathbf{a}}^{-1} (\mathbf{f}_{\mathbf{a}} - \mathbf{m}_{\mathbf{a}}) \quad \text{and} \quad P_{n,\mathbf{a}}(\mathbf{x}, \mathbf{y}) = R(\mathbf{x}, \mathbf{y}) - \mathbf{r}_{\mathbf{a}}(\mathbf{x})^{\top} \mathbf{R}_{\mathbf{a}}^{-1} \mathbf{r}_{\mathbf{a}}(\mathbf{y}). \quad (2.5)$$

Here $\mathbf{R}_{\mathbf{a}} \in \mathbb{R}^{N_n^d \times N_n^d}$ and $\mathbf{r}_{\mathbf{a}}(\mathbf{x}) \in \mathbb{R}^{N_n^d}$ are each given by

$$(\mathbf{R}_{\mathbf{a}})_{ij} = \mathbf{D}_{\mathbf{y}}^{\boldsymbol{\alpha}_j} \mathbf{D}_{\mathbf{x}}^{\boldsymbol{\alpha}_i} R(\mathbf{x}, \mathbf{y}) \Big|_{\mathbf{x}=\mathbf{a}} \Big|_{\mathbf{y}=\mathbf{a}} \quad \text{and} \quad (\mathbf{r}_{\mathbf{a}}(\mathbf{x}))_i = \mathbf{D}_{\mathbf{y}}^{\boldsymbol{\alpha}_i} R(\mathbf{x}, \mathbf{y}) \Big|_{\mathbf{y}=\mathbf{a}}, \quad (2.6)$$

where subscripts denote the differentiation variable, and $\mathbf{m}_{\mathbf{a}} = (\mathbf{D}^{\boldsymbol{\alpha}_1} m(\mathbf{a}), \dots, \mathbf{D}^{\boldsymbol{\alpha}_{N_n^d}} m(\mathbf{a}))$. When f has multidimensional range, one may model each of its components independently, though we note that this modelling choice may be readily generalised using vector-valued Gaussian processes (Álvarez et al., 2012).

2.3 Replicating the Taylor Expansion Using Taylor Kernels

The next theorem combines what was described in Sections 2.1 and 2.2 to give a probabilistic version of the Taylor expansion.

Theorem 2.1. *Let $K_{\mathbf{a}}$ be a Taylor kernel defined as in (2.1). Let $f_{\text{GP}} \sim \text{GP}(m, K_{\mathbf{a}})$ and $\mathbf{f}_{\mathbf{a}} = (\mathbf{D}^{\boldsymbol{\alpha}} f(\mathbf{a}))_{|\boldsymbol{\alpha}| \leq n}$. Then $f_{\text{GP}} | \mathbf{f}_{\mathbf{a}} \sim \text{GP}(s_{n,\mathbf{a}}, P_{n,\mathbf{a}})$, where*

$$s_{n,\mathbf{a}}(\mathbf{x}) = m(\mathbf{x}) + \sum_{|\boldsymbol{\alpha}| \leq n} \frac{\mathbf{D}^{\boldsymbol{\alpha}} [f(\mathbf{a}) - m(\mathbf{a})]}{\boldsymbol{\alpha}!} (\mathbf{x} - \mathbf{a})^{\boldsymbol{\alpha}} \quad \text{and} \quad P_{n,\mathbf{a}}(\mathbf{x}, \mathbf{y}) = \sigma^2 \sum_{|\boldsymbol{\alpha}| > n} \frac{c_{\boldsymbol{\alpha}} \boldsymbol{\lambda}^{\boldsymbol{\alpha}}}{(\boldsymbol{\alpha}!)^2} (\mathbf{x} - \mathbf{a})^{\boldsymbol{\alpha}} (\mathbf{y} - \mathbf{a})^{\boldsymbol{\alpha}}. \quad (2.7)$$

If m is a polynomial of degree at most n , then

$$s_{n,\mathbf{a}}(\mathbf{x}) = \sum_{|\boldsymbol{\alpha}| \leq n} \frac{\mathbf{D}^{\boldsymbol{\alpha}} f(\mathbf{a})}{\boldsymbol{\alpha}!} (\mathbf{x} - \mathbf{a})^{\boldsymbol{\alpha}}, \quad (2.8)$$

which is identical to the multidimensional version of the Taylor polynomial in (1.1).

Proof. It is straightforward to compute that, for any $\boldsymbol{\beta}, \boldsymbol{\gamma} \in \mathbb{N}_0^d$,

$$\mathbf{D}_{\mathbf{y}}^{\boldsymbol{\beta}} \mathbf{D}_{\mathbf{x}}^{\boldsymbol{\gamma}} K_{\mathbf{a}}(\mathbf{x}, \mathbf{y}) = \sigma^2 \sum_{\boldsymbol{\alpha} \geq \boldsymbol{\beta} \wedge \boldsymbol{\gamma}} \frac{c_{\boldsymbol{\alpha}} \boldsymbol{\lambda}^{\boldsymbol{\alpha}} (\mathbf{x} - \mathbf{a})^{\boldsymbol{\alpha} - \boldsymbol{\beta}} (\mathbf{y} - \mathbf{a})^{\boldsymbol{\alpha} - \boldsymbol{\gamma}}}{(\boldsymbol{\alpha} - \boldsymbol{\beta})! (\boldsymbol{\alpha} - \boldsymbol{\gamma})!}, \quad (2.9)$$

where $\boldsymbol{\beta} \wedge \boldsymbol{\gamma} = (\max\{\boldsymbol{\beta}(1), \boldsymbol{\gamma}(1)\}, \dots, \max\{\boldsymbol{\beta}(d), \boldsymbol{\gamma}(d)\})$. If $\mathbf{x} = \mathbf{a}$ or $\mathbf{y} = \mathbf{a}$, all terms with $\boldsymbol{\alpha} - \boldsymbol{\beta} \neq \mathbf{0}$ or $\boldsymbol{\alpha} - \boldsymbol{\gamma} \neq \mathbf{0}$, respectively, in (2.9) vanish. Therefore in the context of (2.6) we have

$$(\mathbf{R}_{\mathbf{a}})_{ij} = \sigma^2 c_{\boldsymbol{\alpha}_i} \boldsymbol{\lambda}^{\boldsymbol{\alpha}_i} \delta_{ij} \quad \text{and} \quad (\mathbf{r}_{\mathbf{a}}(\mathbf{x}))_i = \sigma^2 \frac{c_{\boldsymbol{\alpha}_i} \boldsymbol{\lambda}^{\boldsymbol{\alpha}_i}}{\boldsymbol{\alpha}_i!} (\mathbf{x} - \mathbf{a})^{\boldsymbol{\alpha}_i}. \quad (2.10)$$

The matrix \mathbf{R}_a is thus diagonal and the i th element of the row vector $\mathbf{r}_a(\mathbf{x})^\top \mathbf{R}_a^{-1}$ in (2.5) is $(\alpha_i!)^{-1}(\mathbf{x} - \mathbf{a})^{\alpha_i}$. It follows that the posterior mean and covariance are

$$s_{n,\mathbf{a}}(\mathbf{x}) = m(\mathbf{x}) + \sum_{|\alpha| \leq n} \frac{D^\alpha [f(\mathbf{a}) - m(\mathbf{a})]}{\alpha!} (\mathbf{x} - \mathbf{a})^\alpha,$$

$$P_{n,\mathbf{a}}(\mathbf{x}, \mathbf{y}) = K_a(\mathbf{x}, \mathbf{y}) - \sigma^2 \sum_{|\alpha| \leq n} \frac{c_\alpha \lambda^\alpha}{(\alpha!)^2} (\mathbf{x} - \mathbf{a})^\alpha (\mathbf{y} - \mathbf{a})^\alpha,$$

which proves (2.7). From (2.1) we recognise the covariance $P_{n,\mathbf{a}}$ as the remainder in the kernel expansion. To prove (2.8) it is sufficient to observe that $m(\mathbf{x}) = \sum_{|\alpha| \leq n} D^\alpha m(\mathbf{a}) (\alpha!)^{-1} (\mathbf{x} - \mathbf{a})^\alpha$ when m is a polynomial of degree at most n . By inspection it is clear that $s_{n,\mathbf{a}}$ in (2.8) is identical to the Taylor expansion given in (1.1) (and its multivariate version), which completes the proof. \square

Note that the covariance is not identically zero—in fact, $P_{n,\mathbf{a}}(\mathbf{x}, \mathbf{x}) \rightarrow \infty$ as $\|\mathbf{x}\|_2 \rightarrow \infty$. Furthermore, while $P_{n,\mathbf{a}}(\mathbf{x}, \mathbf{y})$ takes the form of an infinite sum, provided K_a has a closed form it can be computed by subtracting the terms with $|\alpha| \leq n$ in the summation form of K_a from that closed form. For illustration, some posterior processes are displayed in Figure 1.

2.4 Error Bounds

Each positive-semidefinite kernel R on $\Omega_{\mathbf{a},r} \times \Omega_{\mathbf{a},r}$ is associated to a unique *reproducing kernel Hilbert space* (RKHS), $\mathcal{H}(R)$, equipped with inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}(R)}$ and norm $\|\cdot\|_{\mathcal{H}(R)}$. The RKHS is a Hilbert space of functions $f: \Omega_{\mathbf{a},r} \rightarrow \mathbb{R}$ such that $R(\cdot, \mathbf{x}) \in \mathcal{H}(R)$ for every $\mathbf{x} \in \Omega_{\mathbf{a},r}$ and in which the kernel R has the *reproducing property*

$$f(\mathbf{x}) = \langle f, R(\cdot, \mathbf{x}) \rangle_{\mathcal{H}(R)} \quad \text{for all } f \in \mathcal{H}(R) \text{ and } \mathbf{x} \in \Omega_{\mathbf{a},r}.$$

See [Berlinet and Thomas-Agnan \(2004\)](#) for more information on RKHSs. It is often difficult to characterise the functions which lie in the RKHS. Fortunately, for Taylor kernels one may use results such as Theorem 9 in [Minh \(2010\)](#) to show that

$$\mathcal{H}(K_a) = \left\{ f(\mathbf{x}) = \sigma \sum_{\alpha \in \mathbb{N}_0^d} f_\alpha \frac{\sqrt{c_\alpha \lambda^\alpha}}{\alpha!} (\mathbf{x} - \mathbf{a})^\alpha : \|f\|_{\mathcal{H}(K)}^2 = \sum_{\alpha \in \mathbb{N}_0^d} f_\alpha^2 < \infty \right\}.$$

See also [Zwacknagl and Schaback \(2013\)](#) and [Paulsen and Raghupathi \(2016, Sec. 2.1\)](#). For example, all polynomials are contained in the RKHS of any Taylor kernel for any $\mathbf{a} \in \mathbb{R}^d$.

The next theorem shows that the posterior variance has a similar interpretation to the Taylor remainder term if f is in $\mathcal{H}(K_a)$.

Theorem 2.2. *Let $f_{\text{GP}} | \mathbf{f}_a$ be as in Theorem 2.1, and let the assumptions of that theorem hold. If $f \in \mathcal{H}(K_a)$, then $s_{n,\mathbf{a}}$ and $P_{n,\mathbf{a}}$ satisfy*

$$|f(\mathbf{x}) - s_{n,\mathbf{a}}(\mathbf{x})| \leq \|f\|_{\mathcal{H}(K_a)} P_{n,\mathbf{a}}(\mathbf{x}, \mathbf{x})^{1/2} \leq C_{n,r} \sigma \|f\|_{\mathcal{H}(K_a)} \|\mathbf{x} - \mathbf{a}\|_2^{n+1}$$

for all $\mathbf{x} \in \Omega_{\mathbf{a},r}$, where $(C_{n,r})_{n=0}^\infty$ is a positive sequence such that $C_{n,r} \rightarrow 0$ as $n \rightarrow \infty$.

Proof. By the standard equivalence between Gaussian process regression in the noiseless setting and worst-case optimal approximation (e.g., [Scheuerer et al., 2013](#); [Kanagawa et al., 2018](#)) the posterior mean $s_{n,\mathbf{a}} \in \mathcal{H}(K_a)$ is the minimum-norm approximant of f such that $D^\alpha s_{n,\mathbf{a}}(\mathbf{a}) = D^\alpha f(\mathbf{a})$ for every $|\alpha| \leq n$ and the posterior standard deviation at \mathbf{x} , $P_{n,\mathbf{a}}(\mathbf{x}, \mathbf{x})^{1/2}$, equals the worst-case approximation error at \mathbf{x} in the RKHS. Hence

$$|f(\mathbf{x}) - s_{n,\mathbf{a}}(\mathbf{x})| \leq \|f - s_a\|_{\mathcal{H}(K_a)} P_{n,\mathbf{a}}(\mathbf{x}, \mathbf{x})^{1/2} \leq \|f\|_{\mathcal{H}(K_a)} P_{n,\mathbf{a}}(\mathbf{x}, \mathbf{x})^{1/2}$$

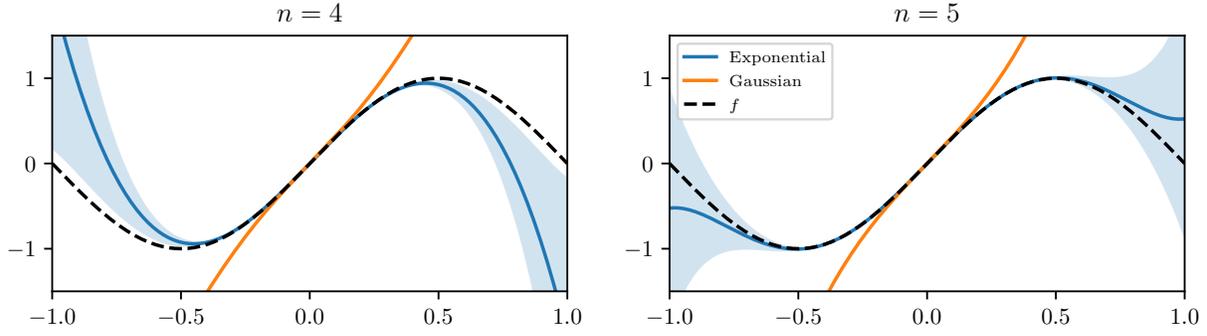


Figure 1: Gaussian process posterior means and 95% credible intervals given derivative data for $f(x) = \sin(\pi x)$ at $a = 0$. The priors have $m \equiv 0$ and use either (blue) the exponential kernel $K(x, y) = \sigma^2 \exp(\lambda xy)$ with $\lambda = 3/2$ or (orange) the Gaussian kernel $K(x, y) = \sigma^2 \exp(-\lambda^2(x - y)^2/2)$. Maximum likelihood estimation was used to select the scaling parameter σ . We shall revisit this example in Section 6.

for all $\mathbf{x} \in \Omega_{\mathbf{a}, r}$ if $f \in \mathcal{H}(K_{\mathbf{a}})$ (e.g., Wendland, 2005, Thm. 16.3). To prove the upper bound for the posterior variance observe that from the general inequality $\|\mathbf{x}\|_{\infty}^{\alpha} \leq \|\mathbf{x}\|_{\infty}^{|\alpha|} \leq \|\mathbf{x}\|_2^{|\alpha|}$ for $\mathbf{x} \in \mathbb{R}^d$ and $\alpha \in \mathbb{N}_0^d$ it follows that, for any $\mathbf{x} \in \Omega_{\mathbf{a}, r} = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x} - \mathbf{a}\|_2 < r\}$,

$$\begin{aligned}
P_{n, \mathbf{a}}(\mathbf{x}, \mathbf{x}) &= \sigma^2 \sum_{|\alpha| > n} \frac{c_{\alpha} \lambda^{\alpha}}{(\alpha!)^2} (\mathbf{x} - \mathbf{a})^{2\alpha} \\
&= \sigma^2 \left(\sum_{|\alpha|=n+1} \frac{c_{\alpha} \lambda^{\alpha}}{(\alpha!)^2} (\mathbf{x} - \mathbf{a})^{2\alpha} + \sum_{|\alpha| > n+1} \frac{c_{\alpha} \lambda^{\alpha}}{(\alpha!)^2} (\mathbf{x} - \mathbf{a})^{2\alpha} \right) \\
&\leq \sigma^2 \left(\sum_{|\alpha|=n+1} \frac{c_{\alpha} \lambda^{\alpha}}{(\alpha!)^2} \|\mathbf{x} - \mathbf{a}\|_2^{2(n+1)} + \sum_{|\alpha| > n+1} \frac{c_{\alpha} \lambda^{\alpha}}{(\alpha!)^2} \|\mathbf{x} - \mathbf{a}\|_2^{2|\alpha|} \right) \\
&= \sigma^2 \|\mathbf{x} - \mathbf{a}\|_2^{2(n+1)} \left(\sum_{|\alpha|=n+1} \frac{c_{\alpha} \lambda^{\alpha}}{(\alpha!)^2} + \sum_{|\alpha| > n+1} \frac{c_{\alpha} \lambda^{\alpha}}{(\alpha!)^2} \|\mathbf{x} - \mathbf{a}\|_2^{2|\alpha| - 2(n+1)} \right) \\
&\leq \sigma^2 \|\mathbf{x} - \mathbf{a}\|_2^{2(n+1)} \underbrace{\left(\sum_{|\alpha|=n+1} \frac{c_{\alpha} \lambda^{\alpha}}{(\alpha!)^2} + r^{-2(n+1)} \sum_{|\alpha| > n+1} \frac{c_{\alpha} \lambda^{\alpha}}{(\alpha!)^2} r^{2|\alpha|} \right)}_{=C_{n, r}}.
\end{aligned}$$

The summability assumption (2.2) ensures that $C_{n, r}$ is finite and that $C_{n, r} \rightarrow 0$ as $n \rightarrow \infty$. \square

2.5 Noisy Observations

Suppose that the observations are corrupted by Gaussian noise. That is, the data vector is $\mathbf{f}_{\mathbf{a}} = (y_{\alpha})_{|\alpha| \leq n}$, where $y_{\alpha} = D^{\alpha} f(\mathbf{a}) + z_{\alpha}$ with independent $z_{\alpha} \sim N(0, \varepsilon_{\alpha}^2)$ for $\varepsilon_{\alpha} > 0$. In this setting the posterior mean and covariance for a general sufficiently differentiable prior mean m and covariance kernel R are

$$s_{n, \mathbf{a}}(\mathbf{x}) = m(\mathbf{x}) + \mathbf{r}_{\mathbf{a}}(\mathbf{x})^{\top} (\mathbf{R}_{\mathbf{a}} + \mathbf{E})^{-1} (\mathbf{f}_{\mathbf{a}} - \mathbf{m}_{\mathbf{a}}) \quad \text{and} \quad P_{n, \mathbf{a}}(\mathbf{x}, \mathbf{y}) = R(\mathbf{x}, \mathbf{y}) - \mathbf{r}_{\mathbf{a}}(\mathbf{x})^{\top} (\mathbf{R}_{\mathbf{a}} + \mathbf{E})^{-1} \mathbf{r}_{\mathbf{a}}(\mathbf{y}), \quad (2.11)$$

where \mathbf{E} is a diagonal $N_n^d \times N_n^d$ matrix containing the noise variances ε_{α}^2 and $\mathbf{r}_{\mathbf{a}}(\mathbf{x})$, $\mathbf{R}_{\mathbf{a}}$, and $\mathbf{m}_{\mathbf{a}}$ were defined in Section 2.2. Recall then from Section 2.3 that when R is a Taylor kernel $K_{\mathbf{a}}$ we have

$$(\mathbf{R}_{\mathbf{a}})_{ij} = \sigma^2 c_{\alpha_i} \lambda^{\alpha_i} \delta_{ij} \quad \text{and} \quad (\mathbf{r}_{\mathbf{a}}(\mathbf{x}))_i = \sigma^2 \frac{c_{\alpha_i} \lambda^{\alpha_i}}{\alpha_i!} (\mathbf{x} - \mathbf{a})^{\alpha_i}.$$

Therefore $(\mathbf{R}_a + \mathbf{E})^{-1} = (\sigma^2 c_\alpha \boldsymbol{\lambda}^{\alpha_i} + \varepsilon_{\alpha_i}^2)^{-1} \delta_{ij}$, and plugging this into (2.11) yields

$$s_{n,\mathbf{a}}(\mathbf{x}) = m(\mathbf{x}) + \sigma^2 \sum_{|\alpha| \leq n} \frac{c_\alpha \boldsymbol{\lambda}^\alpha [y_\alpha - \mathbf{D}^\alpha m(\mathbf{a})]}{\alpha! (\sigma^2 c_\alpha \boldsymbol{\lambda}^\alpha + \varepsilon_\alpha^2)} (\mathbf{x} - \mathbf{a})^\alpha$$

and

$$\begin{aligned} P_{n,\mathbf{a}}(\mathbf{x}, \mathbf{y}) &= K_a(\mathbf{x}, \mathbf{y}) - \sigma^2 \sum_{|\alpha| \leq n} \frac{\sigma^2 c_\alpha^2 \boldsymbol{\lambda}^{2\alpha}}{(\alpha!)^2 (\sigma^2 c_\alpha \boldsymbol{\lambda}^\alpha + \varepsilon_\alpha^2)} (\mathbf{x} - \mathbf{a})^\alpha (\mathbf{y} - \mathbf{a})^\alpha \\ &= \sigma^2 \left[\sum_{\alpha \in \mathbb{N}_0^d} \frac{c_\alpha \boldsymbol{\lambda}^\alpha}{(\alpha!)^2} (\mathbf{x} - \mathbf{a})^\alpha (\mathbf{y} - \mathbf{a})^\alpha - \sum_{|\alpha| \leq n} \frac{\sigma^2 c_\alpha^2 \boldsymbol{\lambda}^{2\alpha}}{(\alpha!)^2 (\sigma^2 c_\alpha \boldsymbol{\lambda}^\alpha + \varepsilon_\alpha^2)} (\mathbf{x} - \mathbf{a})^\alpha (\mathbf{y} - \mathbf{a})^\alpha \right] \\ &= \sigma^2 \left[\sum_{|\alpha| \leq n} \frac{c_\alpha \boldsymbol{\lambda}^\alpha \varepsilon_\alpha^2}{(\alpha!)^2 (\sigma^2 c_\alpha \boldsymbol{\lambda}^\alpha + \varepsilon_\alpha^2)} (\mathbf{x} - \mathbf{a})^\alpha (\mathbf{y} - \mathbf{a})^\alpha + \sum_{|\alpha| > n} \frac{c_\alpha \boldsymbol{\lambda}^\alpha}{(\alpha!)^2} (\mathbf{x} - \mathbf{a})^\alpha (\mathbf{y} - \mathbf{a})^\alpha \right]. \end{aligned}$$

Note that by setting $\varepsilon_\alpha = 0$ for every $\alpha \in \mathbb{N}_0^d$ such that $|\alpha| \leq n$ we recover the noiseless posterior mean and covariance in (2.7).

3 Parameter Estimation

Observe from (2.7) that, although they do not affect the posterior mean, proper selection of the Taylor kernel parameters $\boldsymbol{\lambda}$ and σ is a prerequisite for useful and meaningful uncertainty quantification via the posterior variance $P_{n,\mathbf{a}}(\mathbf{x}, \mathbf{x})$. In this section we consider maximum likelihood estimation of these parameters. For the Gaussian process model in Section 2.2, the negative log-likelihood function that is to be minimised with respect to a generic vector of kernel hyperparameters $\boldsymbol{\theta}$ is

$$\ell(\boldsymbol{\theta}) = \frac{1}{2} (\mathbf{f}_a - \mathbf{m}_a)^\top \mathbf{R}_a^{-1} (\mathbf{f}_a - \mathbf{m}_a) + \frac{1}{2} \log \det \mathbf{R}_a + \frac{N_n^d}{2} \log(2\pi).$$

By discarding terms and coefficients that do not depend on $\boldsymbol{\theta}$ and using (2.10) we see that for a Taylor kernel the maximum likelihood estimate $\boldsymbol{\theta}_{\text{ML}}$ is any minimiser of the function

$$\tilde{\ell}(\boldsymbol{\theta}) = \frac{1}{\sigma^2} \sum_{|\alpha| \leq n} \frac{(\mathbf{D}^\alpha [f(\mathbf{a}) - m(\mathbf{a})])^2}{c_\alpha \boldsymbol{\lambda}^\alpha} + N_n^d \log \sigma^2 + \sum_{|\alpha| \leq n} \log \boldsymbol{\lambda}^\alpha + \sum_{|\alpha| \leq n} c_\alpha. \quad (3.1)$$

In principle, every coefficient c_α of a Taylor kernel in (2.1) is a free parameter. However, maximum likelihood estimation of these coefficient is either useless or impossible: From (2.7) we see that the posterior process depends on c_α only for $|\alpha| > n$. However the objective function (3.1) does not depend on these c_α , making it impossible to estimate those parameters that actually influence posterior uncertainty.

3.1 Estimation of σ

From (3.1) it is easy to calculate σ_{ML} , the maximum likelihood estimate of σ , for any fixed $\boldsymbol{\lambda} \in \mathbb{R}_+^d$ and $n \in \mathbb{N}_0$ by differentiating $\tilde{\ell}$ and setting its derivative to zero. This gives

$$\sigma_{\text{ML}}^2 = \frac{1}{N_n^d} \sum_{|\alpha| \leq n} \frac{(\mathbf{D}^\alpha [f(\mathbf{a}) - m(\mathbf{a})])^2}{c_\alpha \boldsymbol{\lambda}^\alpha}. \quad (3.2)$$

3.2 Estimation of $\boldsymbol{\lambda}$

Estimation of $\boldsymbol{\lambda}$ for a fixed σ is more complicated and the maximum likelihood estimate does not appear to admit a closed form expression akin to that for σ_{ML} in (3.2). However, something interesting can be said. We write $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_d)$.

Lemma 3.1. *Suppose that $n \geq 1$ and let $1 \leq i \leq d$. Then $\lim_{\lambda_i \rightarrow 0} \tilde{\ell}(\boldsymbol{\lambda}) = -\infty$ if $D^\alpha[f(\mathbf{a}) - m(\mathbf{a})] = 0$ for every $|\alpha| \leq n$ such that $\alpha(i) > 0$ and $\lim_{\lambda_i \rightarrow 0} \tilde{\ell}(\boldsymbol{\lambda}) = \infty$ otherwise.*

Proof. Assume first that $D^\alpha[f(\mathbf{a}) - m(\mathbf{a})] = 0$ for every $|\alpha| \leq n$ such that $\alpha(i) > 0$. It follows that the first term in (3.1) does not depend on λ_i . Because $\sum_{|\alpha| \leq n} \log \lambda^\alpha \rightarrow -\infty$ as $\lambda_i \rightarrow 0$, we have $\tilde{\ell}(\boldsymbol{\lambda}) \rightarrow -\infty$ as $\lambda_i \rightarrow 0$. Assume then that $D^\beta[f(\mathbf{a}) - m(\mathbf{a})] \neq 0$ for some $|\beta| \leq n$ such that $\beta(i) > 0$. Therefore

$$\begin{aligned} \frac{1}{\sigma^2} \sum_{|\alpha| \leq n} \frac{(D^\alpha[f(\mathbf{a}) - m(\mathbf{a})])^2}{c_\alpha \lambda^\alpha} + \sum_{|\alpha| \leq n} \log \lambda^\alpha &\geq \frac{1}{\sigma^2} \frac{(D^\beta[f(\mathbf{a}) - m(\mathbf{a})])^2}{c_\beta \lambda^\beta} + \sum_{|\alpha| \leq n} \log \lambda^\alpha \\ &\geq \frac{C}{\lambda_i} + \sum_{|\alpha| \leq n} \log \lambda^\alpha \end{aligned}$$

for a certain positive constant C and $\lambda_i \leq 1$. Since $1/x + a \log x \rightarrow \infty$ for any $a > 0$ as $x \rightarrow 0$ from the right, we conclude from the above lower bound that $\tilde{\ell}(\boldsymbol{\lambda}) \rightarrow \infty$ as $\lambda_i \rightarrow 0$. This concludes the proof. \square

Lemma 3.1 states that $\tilde{\ell}(\boldsymbol{\lambda})$ attains a minimum at $\lambda_i = 0$ when the data are consistent with m being equal to f up to constant along dimension i . From Lemma 3.1 and the fact that $\tilde{\ell}(\boldsymbol{\lambda})$ can tend to negative infinity only if a component of $\boldsymbol{\lambda}$ tends to zero we obtain the following theorem, which is essentially a special case of Theorem 5.2 in Karvonen and Oates (2023). See also Proposition 4.3 in Ben Salem et al. (2019).

Theorem 3.2. *Suppose that $n \geq 1$ and let $1 \leq i \leq d$. Fix λ_j for $j \neq i$. Then*

$$\lambda_{i,\text{ML}} = \arg \min_{\lambda_i \geq 0} \tilde{\ell}((\lambda_1, \dots, \lambda_d)) = 0$$

if and only if $D^\alpha[f(\mathbf{a}) - m(\mathbf{a})] = 0$ for every $|\alpha| \leq n$ such that $\alpha(i) > 0$. In particular, if $d = 1$, then

$$\lambda_{\text{ML}} = \arg \min_{\lambda \geq 0} \tilde{\ell}(\lambda) = 0$$

if and only if $f^{(p)}(a) - m^{(p)}(a) = 0$ for every $1 \leq p \leq n$.

For simplicity, let $d = 1$. If $\lambda = 0$, we see from (2.7) that $P_{n,a}(x, y) = 0$ for all $x, y \in \mathbb{R}$. Moreover, if $f^{(p)}(a) - m^{(p)}(a) = 0$ for every $1 \leq p \leq n$, then

$$s_{n,a}(x) = \sum_{p=0}^n \frac{f^{(p)}(a) - m^{(p)}(a)}{p!} (x - a)^p = f(a) - m(a)$$

for every $x \in \mathbb{R}$. That is, $s_{n,a}$ is a constant function. The interpretation of Theorem 3.2 is thus that when the data look like they could have been generated by the function $f(x) = m(x) + c$ for some $c \in \mathbb{R}$ (i.e., by a constant shift of the prior), maximum likelihood estimation returns $\lambda_{\text{ML}} = 0$ because this value of λ both explains the data and yield the simplest model, one of zero variance. When the posterior covariance is identically zero, the resulting degenerate posterior $f_{\text{GP}} | \mathbf{f}_a \sim \text{GP}(f(a) - m(a), 0)$ does not provide useful uncertainty quantification as it is unreasonable to expect perfect predictions from a finite set of data.

3.3 On Simultaneous Estimation of σ and λ

The purpose of this section is to demonstrate that simultaneous maximum likelihood estimation of σ and $\boldsymbol{\lambda}$ is likely to cause problems. We consider inner product kernels of the form (2.3) with coefficients $c_\alpha = c_{|\alpha|} \alpha! / |\alpha|!$ and $n = 1$. Let $\partial_i^p f(\mathbf{x})$ denote the p th order partial derivative of f at \mathbf{x} with respect to the i th coordinate.

Note from $c_\alpha = c_{|\alpha|} \alpha! / |\alpha|!$ that $c_\alpha = c_0$ for $\alpha = \mathbf{0}$ and $c_\alpha = c_1$ when $|\alpha| = 1$. By differentiating (3.1) with respect to the i th component of $\boldsymbol{\lambda}$ we see that to obtain $\boldsymbol{\lambda}_{\text{ML}} = (\lambda_{1,\text{ML}}, \dots, \lambda_{d,\text{ML}})$ we need to solve

$$\frac{1}{\sigma^2} \sum_{|\alpha| \leq 1} \frac{\alpha(i) (D^\alpha[f(\mathbf{a}) - m(\mathbf{a})])^2}{c_\alpha \lambda^\alpha} - \sum_{|\alpha| \leq 1} \alpha(i) = \frac{(\partial_i[f(\mathbf{a}) - m(\mathbf{a})])^2}{\sigma^2 c_1 \lambda_i} - 1 = 0 \quad (3.3)$$

for each $i = 1, \dots, d$. Equation (3.3) readily gives the maximum likelihood estimates

$$\lambda_{i,\text{ML}} = \frac{(\partial_i[f(\mathbf{a}) - m(\mathbf{a})])^2}{\sigma^2 c_1} \quad (3.4)$$

for a fixed $\sigma > 0$. Inserting these to the expression for the maximum likelihood estimate σ_{ML}^2 in (3.2) yields

$$\sigma_{\text{ML}}^2 = \frac{1}{d+1} \left(\frac{[f(\mathbf{a}) - m(\mathbf{a})]^2}{c_0} + \sum_{i=1}^d \frac{(\partial_i[f(\mathbf{a}) - m(\mathbf{a})])^2}{c_1 \lambda_{i,\text{ML}}} \right) = \frac{1}{d+1} \left(\frac{[f(\mathbf{a}) - m(\mathbf{a})]^2}{c_0} + d\sigma_{\text{ML}}^2 \right),$$

which is solved by $\sigma_{\text{ML}}^2 = [f(\mathbf{a}) - m(\mathbf{a})]^2 / c_0$. By plugging this in (3.4) we obtain the final estimates

$$\sigma_{\text{ML}}^2 = \frac{[f(\mathbf{a}) - m(\mathbf{a})]^2}{c_0} \quad \text{and} \quad \lambda_{i,\text{ML}} = \frac{c_0}{c_1} \left(\frac{\partial_i[f(\mathbf{a}) - m(\mathbf{a})]}{f(\mathbf{a}) - m(\mathbf{a})} \right)^2. \quad (3.5)$$

It is clear what the problem is: If $|f(\mathbf{a}) - m(\mathbf{a})|$ is small relative to $|\partial_i[f(\mathbf{a}) - m(\mathbf{a})]|$ for some i (i.e., $m \approx f$ but $\partial_i m \not\approx \partial_i f$ at \mathbf{a}), the estimate for $\boldsymbol{\lambda}$ becomes large, which may cause numerical problems and yields a large posterior variance. For example, let $d = 1$ and insert the estimates in (3.5) to the posterior variance in (2.7). This gives

$$P_{n,a}(x, x) = \sum_{p=2}^{\infty} \left(\frac{c_0}{[f(a) - m(a)]^2} \right)^{p-1} \left(\frac{[f'(a) - m'(a)]^2}{c_1} \right)^p \frac{c_p}{(p!)^2} (x - a)^{2p}.$$

4 Comparison to the Gaussian Kernel

Let $\|\cdot\|_{\boldsymbol{\lambda}} = \langle \cdot, \cdot \rangle_{\boldsymbol{\lambda}}$. It is well known that the ubiquitous Gaussian kernel

$$R(\mathbf{x}, \mathbf{y}) = \exp\left(-\frac{1}{2} \|\mathbf{x} - \mathbf{y}\|_{\boldsymbol{\lambda}}^2\right)$$

is closely connected to the exponential kernel in (2.4) via the equation

$$R(\mathbf{x}, \mathbf{y}) = \exp\left(-\frac{1}{2} [\langle \mathbf{x}, \mathbf{x} \rangle_{\boldsymbol{\lambda}} - 2\langle \mathbf{x}, \mathbf{y} \rangle_{\boldsymbol{\lambda}} + \langle \mathbf{y}, \mathbf{y} \rangle_{\boldsymbol{\lambda}}]\right) = \exp\left(-\frac{1}{2} \|\mathbf{x}\|_{\boldsymbol{\lambda}}^2\right) \exp(\langle \mathbf{x}, \mathbf{y} \rangle_{\boldsymbol{\lambda}}) \exp\left(-\frac{1}{2} \|\mathbf{y}\|_{\boldsymbol{\lambda}}^2\right).$$

Given such a relationship to a Taylor kernel it should come as no surprise that, given derivative data, the posterior mean and covariance in (2.5) are available in closed form for the Gaussian kernel—even though the matrix \mathbf{R}_a is not diagonal.

Let us consider the univariate case. We get

$$\mathrm{D}_y^i \mathrm{D}_x^j R(x, y) \Big|_{y=a}^{x=a} = (-1)^i \mathrm{D}_z^{i+j} \exp\left(-\frac{\lambda^2}{2} z^2\right) \Big|_{z=0} = (-1)^i \sum_{p=0}^{\infty} (-1)^p \frac{\lambda^{2p}}{2^p p!} \mathrm{D}_z^{i+j} z^{2p} \Big|_{z=0}.$$

When $i + j$ is odd, all derivatives in the sum vanish, so that in this case $\mathrm{D}_y^i \mathrm{D}_x^j R(x, y) \Big|_{y=a}^{x=a} = 0$. If $i + j = 2k$ for $k \in \mathbb{N}_0$, we have

$$\mathrm{D}_y^i \mathrm{D}_x^j R(x, y) \Big|_{y=a}^{x=a} = (-1)^{i+k} \frac{\lambda^{2k} (2k)!}{2^k k!}.$$

This provides us with a relatively straightforward expression for the matrix \mathbf{R}_a in (2.6). From the Rodrigues' formula $\mathrm{H}_p(x) = (-1)^p e^{x^2/2} \mathrm{D}_x^p e^{-x^2/2}$ for the probabilist's Hermite polynomials it is easy to compute $\mathbf{r}_a(x)$:

$$(\mathbf{r}_a(x))_i = \mathrm{D}_y^i R(x, y) \Big|_{y=a} = \mathrm{D}_y^i \exp\left(-\frac{\lambda^2}{2} (x - y)^2\right) \Big|_{y=a} = \lambda^i \exp\left(-\frac{\lambda^2}{2} (x - a)^2\right) \mathrm{H}_i(\lambda(x - a)). \quad (4.1)$$

Finding $s_{n,a}$ and $P_{n,a}$ requires the inverse of \mathbf{R}_a . Fortunately, [Xu and Stein \(2017, Proposition 3.2\)](#) have computed the Cholesky decomposition of the inverse of \mathbf{R}_a .¹ The inverse of \mathbf{R}_a has the Cholesky decomposition $\mathbf{R}_a^{-1} = \mathbf{L}\mathbf{L}^\top$, where $\mathbf{L} \in \mathbb{R}^{(n+1) \times (n+1)}$ is a lower triangular matrix with non-zero elements $(\mathbf{L})_{ij} = \sqrt{i!}/(\lambda^j j!(i-j)!!)$ when $i \geq j$ and $i+j$ is even. Here $i!!$ is the double factorial, the product of positive integers up to i that have the same parity as i . Thus

$$(\mathbf{R}_a^{-1})_{ij} = (\mathbf{L}\mathbf{L}^\top)_{ij} = \lambda^{-(i+j)}Q(i,j), \quad \text{where} \quad Q(i,j) = \sum_{\substack{p=\max\{i,j\} \\ p+i \text{ is even}}}^n \frac{p!}{i!j!(p-i)!(p-j)!}. \quad (4.2)$$

Consequently, inserting (4.1) and (4.2) in (2.5) gives the convenient closed form expressions

$$s_{n,a}(x) = \sum_{i=0}^n \sum_{j=0}^n (\mathbf{r}_a(x))_i (\mathbf{R}_a^{-1})_{ij} f^{(j)}(a) = \exp\left(-\frac{\lambda^2}{2}(x-a)^2\right) \sum_{j=0}^n \frac{f^{(j)}(a)}{\lambda^j j!} \sum_{i=0}^n \frac{Q(i,j)}{i!} \mathbf{H}_i(\lambda(x-a))$$

and

$$P_{n,a}(x,y) = \exp\left(-\frac{\lambda^2}{2}(x-y)^2\right) \left(1 - e^{-\lambda^2(x-a)(y-a)} \sum_{i=0}^n \sum_{j=0}^n \frac{Q(i,j)}{i!j!} \mathbf{H}_i(\lambda(x-a)) \mathbf{H}_j(\lambda(y-a))\right).$$

In particular, the posterior variance is

$$P_{n,a}(x,x) = 1 - e^{-\lambda^2(x-a)^2} \sum_{i=0}^n \sum_{j=0}^n \frac{Q(i,j)}{i!j!} \mathbf{H}_i(\lambda(x-a)) \mathbf{H}_j(\lambda(x-a)).$$

These expressions resemble those in (2.7) for Taylor kernels. However, a notable difference is that for Taylor kernels the posterior variance explodes as $|x-a|$ grows but for the Gaussian kernel the variance tends to a constant as $|x-a| \rightarrow \infty$. Similarly, the posterior mean for Taylor kernels is unbounded, while for the Gaussian kernel the mean reverts to zero (i.e., the prior mean; recall that we set $m \equiv 0$) far away from a .

5 General Orthogonal Data

In this section we discuss how simple posterior formulae analogous to those derived in Section 2.3 are available for any data that are orthogonal.

5.1 Generic Construction

Let Ω be an arbitrary non-empty set, \mathcal{P} a countable index set, and $(\phi_p)_{p \in \mathcal{P}}$ a collection of linearly independent basis functions on Ω such that $\sum_{p \in \mathcal{P}} |\phi_p(x)|^2 < \infty$ for every $x \in \Omega$. We may then define (at least formally) a Gaussian process f_{GP} on Ω by setting $f_{\text{GP}}(x) = \sum_{p \in \mathcal{P}} Z_p \phi_p(x)$ for every $x \in \Omega$, where Z_p are i.i.d standard normal random variables. It is then straightforward to compute that

$$\mathbb{E}[f_{\text{GP}}(x)] = 0 \quad \text{and} \quad R(x,y) = \text{Cov}[f_{\text{GP}}(x), f_{\text{GP}}(y)] = \sum_{p \in \mathcal{P}} \phi_p(x) \phi_p(y). \quad (5.1)$$

The kernel R is positive-semidefinite. Assume that $(\phi_p)_{p \in \mathcal{P}}$ are an orthonormal basis of the RKHS $\mathcal{H}(R)$ (see Section 2.4 for RKHSs).² Then each $f \in \mathcal{H}(R)$ has the pointwise convergent expansion $f(x) = \sum_{p \in \mathcal{P}} f_p \phi_p(x)$ for the coefficients $f_p = \langle f, \phi_p \rangle_{\mathcal{H}(R)}$. Suppose that one observes $\gamma_p f_p$ for p in a finite collection $\mathcal{N} \subset \mathcal{P}$ of indices and some constants γ_p . These data are *orthogonal* because they are obtained by taking inner products of f with a collection of pairwise orthogonal functions $\gamma_p \phi_p$. A derivation similar to that in Section 2.3 then shows that the posterior mean and covariance are simply

$$s(x) = \sum_{p \in \mathcal{N}} f_p \phi_p(x) = \sum_{p \in \mathcal{N}} \gamma_p f_p \frac{\phi_p(x)}{\gamma_p} \quad \text{and} \quad P(x,y) = \sum_{p \in \mathcal{P} \setminus \mathcal{N}} \phi_p(x) \phi_p(y) \quad (5.2)$$

¹Note that the denominator in Equation (3.1) of [Xu and Stein \(2017\)](#) should have $(i-j)!!$ in the place of $(i-j)$.

²Any functions $(\phi_p)_{p \in \mathcal{P}}$ for which the expansion of the kernel R in (5.1) converges pointwise are a *Parseval frame* for the RKHS $\mathcal{H}(R)$ ([Paulsen and Raghupathi, 2016](#), Thm. 2.10).

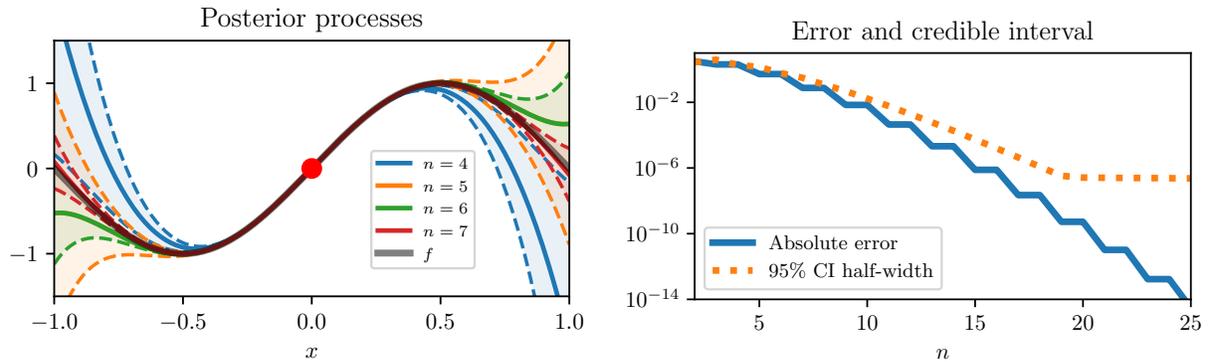


Figure 2: *Left*: Posterior means and 95% credible intervals given derivative data for $f(x) = \sin(\pi x)$ at $a = 0$. The zero-mean prior uses the exponential kernel $K(x, y) = \sigma^2 \exp(\lambda xy)$ with $\lambda = 3/2$ and scale σ set using maximum likelihood. *Right*: Maximal absolute errors $\max_{x \in [-1, 1]} |f(x) - s_{n,a}(x)|$ and half-widths $1.96 \times \max_{x \in [-1, 1]} \sqrt{P_{n,a}(x, x)}$ of the 95% credible interval over the domain $\Omega = [-1, 1]$.

for all $x, y \in \Omega$. See (Wendland, 2005, Ch. 16) and (Oettershagen, 2017, Cor. 3.6) for general formulae when the data consists of applications to f of arbitrary linear functionals. Orthogonal data are known to be optimal in a certain sense and settings (Novak and Woźniakowski, 2008, Sec. 4.2.3). To connect (5.2) to the derivations for Taylor kernels, suppose for instance that $f: \mathbb{R} \rightarrow \mathbb{R}$ has the Taylor expansion

$$f(x) = \sum_{p=0}^{\infty} \frac{f^{(p)}(0)}{p!} x^p = \sum_{p=0}^{\infty} f_p \phi_p(x), \quad \text{where} \quad \phi_p(x) = \frac{\sigma \sqrt{c_p \lambda^p}}{p!} x^p \quad \text{and} \quad f_p = \frac{f^{(p)}(0)}{\sigma \sqrt{c_p \lambda^p}}.$$

In this case we therefore have the index set $\mathcal{P} = \mathbb{N}_0$. Observing the N first derivatives $f^{(p)}(0) = \gamma_p f_p$, so that $\mathcal{N} = \{0, 1, \dots, N\}$ and $\gamma_p = \sigma \sqrt{c_p \lambda^p}$, and using (5.2) yields (2.7) with $d = 1$, $a = 0$, and $m \equiv 0$. Moreover, $\text{Cov}[f_{\text{GP}}(x), f_{\text{GP}}(y)] = \sum_{p=0}^{\infty} \phi_p(x) \phi_p(y) = K(x, y)$ for K the univariate Taylor kernel in (1.2). We mention two other examples orthogonal data.

5.2 Mehler Kernel

Let $\mathcal{P} = \mathbb{N}_0$ and $\Omega = \mathbb{R}$. Let H_p be the p th probabilist's Hermite polynomial and $\rho \in (0, 1)$. Set $\phi_p(x) = \sigma \sqrt{\rho^p (p!)^{-1}} H_p(x)$. Then Mehler's formula yields the *Mehler kernel*

$$R(x, y) = \sum_{p=0}^{\infty} \phi_p(x) \phi_p(y) = \sigma^2 \sum_{p=0}^{\infty} \frac{\rho^p}{p!} H_p(x) H_p(y) = \frac{\sigma^2}{\sqrt{1 - \rho^2}} \exp\left(-\frac{\rho^2(x^2 + y^2) - 2\rho xy}{2(1 - \rho^2)}\right).$$

See Irrgeher and Leobacher (2015) and Oettershagen (2017, Sec. 3.6.4) for the Mehler kernel in the context of kernel-based approximation. If $f(x) = \sum_{p=0}^{\infty} f_p \phi_p(x)$, then by the orthogonality with respect to Gaussian integration, other basic properties of the Hermite polynomials, and properties of Gaussian integrals of derivatives (e.g., Bogachev, 1998, Rmk. 1.3.5),

$$\gamma_p f_p = \frac{1}{\sqrt{2\pi}} \int_{\mathbb{R}} f(x) H_p(x) \exp\left(-\frac{x^2}{2}\right) dx = \frac{1}{\sqrt{2\pi}} \int_{\mathbb{R}} f^{(p)}(x) \exp\left(-\frac{x^2}{2}\right) dx, \quad \text{where} \quad \gamma_p = \sigma \rho^p.$$

Here orthogonal data are therefore obtained by Gaussian integration of the derivatives of f .

5.3 Periodic Kernel

Let $\mathcal{P} = \mathbb{Z}$, $\Omega = [0, 1]$, and $s \in \mathbb{N}$. Set $\phi_p(x) = \sigma \sqrt{2}(2\pi p)^{-s} \cos(2\pi p x)$ and $\phi_{-p}(x) = \sigma \sqrt{2}(2\pi p)^{-s} \sin(2\pi p x)$ for $p \in \mathbb{N}$. Moreover, set $\phi_0 \equiv \sigma$. Then we obtain the *periodic Sobolev kernel* (or the *Korobov kernel*)

$$R(x, y) = \sum_{p \in \mathbb{Z}} \phi_p(x) \phi_p(y) = \sigma^2 \left(1 + 2 \sum_{p=1}^{\infty} \frac{1}{(2\pi p)^{2s}} \cos(2\pi p(x - y))\right) = \sigma^2 \left(1 + \frac{(-1)^{s+1}}{(2s)!} B_{2s}(|x - y|)\right) \quad (5.3)$$

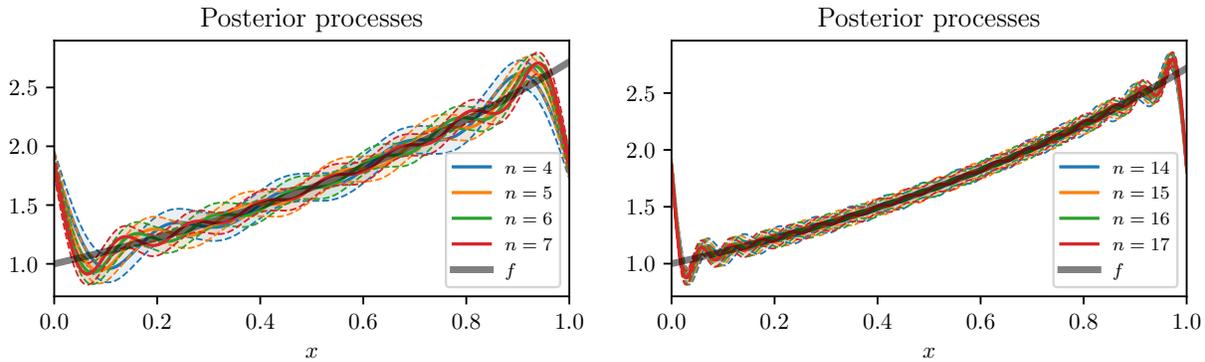


Figure 3: Posterior means and 95% credible intervals given Fourier data for $f(x) = \exp(x)$ on $\Omega = [0, 1]$. The zero-mean prior uses the periodic kernel in (5.3) with $s = 2$ and scale σ set using maximum likelihood.

for $x, y \in [0, 1]$, where B_{2s} is the Bernoulli polynomial of degree $2s$; see, for example, Wahba (1990, Sec. 2.1). If $f: [0, 1] \rightarrow \mathbb{R}$ has the expansion $f(x) = \sum_{p \in \mathbb{Z}} f_p \phi_p(x)$, then it is straightforward to compute that

$$\gamma_p f_p = 2 \int_0^1 f(x) \cos(2\pi p x) dx \quad \text{and} \quad f_{-p} = 2 \int_0^1 f(x) \sin(2\pi p x) dx \quad \text{for } p \in \mathbb{N} \quad (5.4)$$

for $\gamma_p = \sigma \sqrt{2}(2\pi p)^{-s}$ and $\gamma_0 f_0 = \int_0^1 f(x) dx$ for $\gamma = \sigma$. These orthogonal data are the Fourier coefficients.

6 Numerical Examples

This section contains two numerical examples. Figure 2 displays a number of posterior processes and the behaviour of maximal error and standard deviation when a zero-mean Gaussian process with the Taylor kernel $K(x, y) = \sigma^2 \exp(\lambda xy)$ with $\lambda = 3/2$ is used to infer the function $f(x) = \sin(\pi x)$ based on noiseless derivative evaluations at $a = 0$, as described in Section 2. See also Figure 1. The scaling parameter σ was taken to be the maximum likelihood estimate in (3.2). From the right panel we see that the Gaussian process model is well-calibrated in the weak sense that, except for small n , $f(x)$ is never further away from the posterior mean than maximal half-width of the 95% credible interval over the domain $\Omega = [-1, 1]$ of interest: $\max_{x \in [-1, 1]} |f(x) - s_{n,a}(x)| \leq 1.96 \times \max_{x \in [-1, 1]} \sqrt{P_{n,a}(x, x)}$.

Our second example uses the periodic kernel (5.3) with $s = 2$ and scaled Fourier data in (5.4), so that the posterior mean and covariance are given by (5.2). We use index sets of the form $\mathcal{N} = \{-n, \dots, -1, 0, 1, \dots, n\}$ for $n \in \mathbb{N}$ and again use maximum likelihood to set the scaling parameter, which in this case simply yields $\sigma_{\text{ML}}^2 = \frac{1}{2n+1} \sum_{p=-n}^n (\gamma_p f_p)^2$. The function being inferred is $f(x) = \exp(x)$, and we compute that

$$\gamma_p f_p = s_p [2e\pi p \sin(2\pi p) + e \cos(2\pi p) - 1] \quad \text{and} \quad \gamma_{-p} f_{-p} = s_p [2\pi p + e \sin(2\pi p) - 2e\pi p \cos(2\pi p)]$$

for $p \in \mathbb{N}$, where $s_p = 2/(4\pi^2 p^2 + 1)$, and $\gamma_0 f_0 = e - 1$. Figure 3 depicts some of the resulting posterior processes. Except at the boundaries where the Gibbs phenomenon caused by the non-periodicity of f occurs, the posteriors fare well and appear to provide reasonable quantification of predictive uncertainty.

7 Conclusion

We have proposed a GP model based on Taylor kernels which gives rise to a probabilistic version of the classical Taylor expansion when the data consist of derivative evaluations. Using Taylor kernels in Bayesian optimisation (Snoek et al., 2012) would be an interesting future application, where they might be expected to inherit properties from both standard Bayesian optimisation algorithms based on commonly used stationary kernels, such as the Gaussian and Matérns, and classical optimisation algorithms. Because their uncertainty explodes away from the expansion point, Taylor kernels might prove a useful alternative to stationary kernels which have a tendency to be over-exploitative in Bayesian optimisation (Bull, 2011).

References

- Álvarez, M. A., Rosasco, L., and Lawrence, N. D. (2012). Kernels for vector-valued functions: A review. *Foundations and Trends® in Machine Learning*, 4(3):195–266.
- Ben Salem, M., Bachoc, F., Roustant, O., Gamboa, F., and Tomaso, L. (2019). Gaussian process-based dimension reduction for goal-oriented sequential design. *SIAM/ASA Journal on Uncertainty Quantification*, 7(4):1369–1397.
- Berlinet, A. and Thomas-Agnan, C. (2004). *Reproducing Kernel Hilbert Spaces in Probability and Statistics*. Springer.
- Bogachev, V. I. (1998). *Gaussian Measures*. Number 62 in Mathematical Surveys and Monographs. American Mathematical Society.
- Bull, A. D. (2011). Convergence rates of efficient global optimization algorithms. *Journal of Machine Learning Research*, 12:2879–2904.
- Cockayne, J., Oates, C. J., Ipsen, I. C. F., and Girolami, M. (2019a). A Bayesian conjugate gradient method (with discussion). *Bayesian Analysis*, 14(3):937–1012.
- Cockayne, J., Oates, C. J., Sullivan, T., and Girolami, M. (2019b). Bayesian probabilistic numerical methods. *SIAM Review*, 61(4):756–789.
- Conn, A. R., Gould, N. I., and Tointi, P. L. (2000). *Trust-Region Methods*, volume 1 of *MPS-SIAM Series on Optimization*. Society for Industrial and Applied Mathematics.
- De Marchi, S. and Schaback, R. (2010). Nonstandard kernels and their applications. *Dolomites Research Notes on Approximation*, 2:16–43.
- Diaconis, P. (1988). Bayesian numerical analysis. In *Statistical decision theory and related topics IV*, volume 1, pages 163–175. Springer-Verlag New York.
- Dick, J. (2006). A Taylor space for multivariate integration. *Monte Carlo Methods and Applications*, 12(2):99–112.
- Eriksson, D., Dong, K., Lee, E., Bindel, D., and Wilson, A. G. (2018). Scaling Gaussian process regression with derivatives. In *Advances in Neural Information Processing Systems*, volume 31, pages 6867–6877.
- Fasshauer, G. and McCourt, M. (2015). *Kernel-Based Approximation Methods Using MATLAB*. Number 19 in Interdisciplinary Mathematical Sciences. World Scientific Publishing.
- Hairer, E., Nørsett, S. P., and Wanner, G. (1993). *Solving Ordinary Differential Equations I: Nonstiff Problems*, volume 8 of *Springer Series in Computational Mathematics*. Springer.
- Hennig, P. (2015). Probabilistic interpretation of linear solvers. *SIAM Journal on Optimization*, 25(1):234–260.
- Hennig, P., Osborne, M. A., and Kersting, H. P. (2022). *Probabilistic Numerics*. Cambridge University Press.
- Irrgeher, C. and Leobacher, G. (2015). High-dimensional integration on \mathbb{R}^d , weighted Hermite spaces, and orthogonal transforms. *Journal of Complexity*, 31(2):174–205.
- Kanagawa, M., Hennig, P., Sejdinovic, D., and Sriperumbudur, B. K. (2018). Gaussian processes and kernel methods: A review on connections and equivalences. *arXiv:1807.02582v1*.
- Karvonen, T. and Oates, C. J. (2023). Maximum likelihood estimation in Gaussian process regression is ill-posed. *Journal of Machine Learning Research*, 24(120):1–47.
- Karvonen, T., Oates, C. J., and Särkkä, S. (2018). A Bayes–Sard cubature method. In *Advances in Neural Information Processing Systems*, volume 31, pages 5882–5893.

- Karvonen, T. and Särkkä, S. (2017). Classical quadrature rules via Gaussian processes. In *27th IEEE International Workshop on Machine Learning for Signal Processing*.
- Kimeldorf, G. S. and Wahba, G. (1970). A correspondence between Bayesian estimation on stochastic processes and smoothing by splines. *The Annals of Mathematical Statistics*, 41(2):495–502.
- Larkin, F. M. (1970). Optimal approximation in Hilbert spaces with reproducing kernel functions. *Mathematics of Computation*, 24(112):911–921.
- Liang, T. and Rakhlin, A. (2020). Just interpolate: Kernel “ridgeless” regression can generalize. *Annals of Statistics*, 48(3):1329–1347.
- Minh, H. Q. (2010). Some properties of Gaussian reproducing kernel Hilbert spaces and their implications for function approximation and learning theory. *Constructive Approximation*, 32(2):307–338.
- Minka, T. (2000). Deriving quadrature rules from Gaussian processes. Technical report, Statistics Department, Carnegie Mellon University.
- Moré, J. J. (1978). The Levenberg-Marquardt algorithm: Implementation and theory. In *Numerical Analysis*, volume 630 of *Lecture Notes in Mathematics*, pages 105–116. Springer.
- Novak, E. and Woźniakowski, H. (2008). *Tractability of Multivariate Problems, Volume I: Linear Information*. Number 6 in EMS Tracts in Mathematics. European Mathematical Society.
- Oettershagen, J. (2017). *Construction of Optimal Cubature Algorithms with Applications to Econometrics and Uncertainty Quantification*. PhD thesis, Institut für Numerische Simulation, Universität Bonn.
- Paulsen, V. I. and Raghupathi, M. (2016). *An Introduction to the Theory of Reproducing Kernel Hilbert Spaces*. Number 152 in Cambridge Studies in Advanced Mathematics. Cambridge University Press.
- Prüher, J. and Särkkä, S. (2016). On the use of gradient information in Gaussian process quadratures. In *26th IEEE International Workshop on Machine Learning for Signal Processing*.
- Rasmussen, C. E. and Williams, C. K. I. (2006). *Gaussian Processes for Machine Learning*. Adaptive Computation and Machine Learning. MIT Press.
- Raudenbush, S. W., Yang, M.-L., and Yosef, M. (2000). Maximum likelihood for generalized linear models with nested random effects via high-order, multivariate Laplace approximation. *Journal of Computational and Graphical Statistics*, 9(1):141–157.
- Richter-Dyn, N. (1971a). Minimal interpolation and approximation in Hilbert spaces. *SIAM Journal on Numerical Analysis*, 8(3):583–597.
- Richter-Dyn, N. (1971b). Properties of minimal integration rules. II. *SIAM Journal on Numerical Analysis*, 8(3):497–508.
- Särkkä, S. (2013). *Bayesian Filtering and Smoothing*, volume 3 of *IMS Textbooks*. Cambridge University Press.
- Scheuerer, M., Schaback, R., and Schlather, M. (2013). Interpolation of spatial data – A stochastic or a deterministic problem? *European Journal of Applied Mathematics*, 24(4):601–629.
- Schober, M., Duvenaud, D. K., and Hennig, P. (2014). Probabilistic ODE solvers with Runge-Kutta means. In *Advances in Neural Information Processing Systems*, volume 27, pages 739–747.
- Schober, M., Särkkä, S., and Hennig, P. (2019). A probabilistic model for the numerical solution of initial value problems. *Statistics and Computing*, 29:99–122.
- Snoek, J., Larochelle, H., and Adams, R. P. (2012). Practical Bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems*, volume 25, pages 2951–2959.

- Solak, E., Murray-Smith, R., Leithead, W. E., Leith, D. J., and Rasmussen, C. (2002). Derivative observations in Gaussian process models of dynamic systems. In *Advances in Neural Information Processing Systems*, volume 15, pages 1057–1064.
- Steinwart, I. and Christmann, A. (2008). *Support Vector Machines*. Information Science and Statistics. Springer.
- Särkkä, S. (2011). Linear operators and stochastic partial differential equations in Gaussian process regression. In *International Conference on Artificial Neural Networks*, pages 151–158.
- Teymur, O., Zygalakis, K., and Calderhead, B. (2016). Probabilistic linear multistep methods. In *Advances in Neural Information Processing Systems*, volume 29, pages 4321–4328.
- Travelletti, C. and Ginsbourger, D. (2022). Disintegration of Gaussian measures for sequential assimilation of linear operator data. *arXiv:2207.13581v1*.
- Wahba, G. (1990). *Spline Models for Observational Data*. Number 59 in CBMS-NSF Regional Conference Series in Applied Mathematics. Society for Industrial and Applied Mathematics.
- Wendland, H. (2005). *Scattered Data Approximation*. Number 17 in Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press.
- Wu, J., Poloczek, M., Wilson, A. G., and Frazier, P. I. (2017). Bayesian optimization with gradients. In *Advances in Neural Information Processing Systems*, volume 30, pages 5267–5278.
- Xu, W. and Stein, M. L. (2017). Maximum likelihood estimation for a smooth Gaussian random field model. *SIAM/ASA Journal on Uncertainty Quantification*, 5(1):138–175.
- Zwacknagl, B. (2009). Power series kernels. *Constructive Approximation*, 29(1):61–84.
- Zwacknagl, B. and Schaback, R. (2013). Interpolation and approximation in Taylor spaces. *Journal of Approximation Theory*, 171:65–83.