

Sample Path Regularity of Gaussian Processes from the Covariance Kernel

Nathaël Da Costa¹
Marvin Pförtner¹
Lancelot Da Costa^{2,3,4}
Philipp Hennig¹

NATHAEL.DA-COSTA@UNI-TUEBINGEN.DE
MARVIN.PFOERTNER@UNI-TUEBINGEN.DE
LANCE.DACOSTA@VERSES.AI
PHILIPP.HENNIG@UNI-TUEBINGEN.DE

¹*Tübingen AI Center, University of Tübingen, Tübingen, Germany*

²*VERSES AI Research Lab, Los Angeles, USA*

³*Department of Mathematics, Imperial College, London, UK*

⁴*Wellcome Centre for Human Neuroimaging, UCL, London, UK*

Abstract

Gaussian processes (GPs) are the most common formalism for defining probability distributions over spaces of functions. While applications of GPs are myriad, a comprehensive understanding of GP sample paths, i.e. the function spaces over which they define a probability measure, is lacking. In practice, GPs are not constructed through a probability measure, but instead through a mean function and a covariance kernel. In this paper we provide necessary and sufficient conditions on the covariance kernel for the sample paths of the corresponding GP to attain a given regularity. We use the framework of Hölder regularity as it grants particularly straightforward conditions, which simplify further in the cases of stationary and isotropic GPs. We then demonstrate that our results allow for novel and unusually tight characterisations of the sample path regularities of the GPs commonly used in machine learning applications, such as the Matérn GPs.

Keywords: Gaussian processes, Gaussian random fields, differentiability, Hölder regularity, Sobolev regularity, stationary kernels, isotropic kernels, Matérn kernels

1 Introduction

Gaussian processes (GPs) provide a formalism to assign probability distributions over spaces of functions. That distribution is principally characterised and controlled by the process’ covariance function, which is a positive definite kernel. Such kernels are also associated with reproducing kernel Hilbert spaces (RKHSs). However, it is relatively widely known that the *sample paths* of the associated GP are not generally elements of the RKHS, but form a “larger” space of typically less regular functions (Kanagawa et al., 2018, Section 4). This sample path space is *much harder to characterise* than the RKHS.

GP regression is widely applied in statistics and machine learning for inference from physical observations. In such settings, practitioners mostly concern themselves with the posterior mean function (which *is* an element of the RKHS) and the marginal variance (which is *not* in the RKHS, but inherits its regularity) while largely ignoring the regularity of the sample paths and other properties of the support of the prior probability measure. But recently relevant use cases for GPs in computational tasks urgently require a more careful, and ideally tight analysis of the sample path regularity, which we provide in this

work. For instance, Pförtner et al. (2022) recently showed that a large class of classic solution methods for linear partial differential equations (PDEs) can be interpreted as GP inference, i.e. as base instances of probabilistic numerical methods. To infer the solution to a PDE using a probabilistic numerical method, we want to construct a GP prior for it. It should be tailored to the problem as tightly as possible, for the following reasons:

1. We need to ensure that the PDE’s differential operator is well-defined on all samples of the GP. Hence, the sample paths must be **regular enough**. Here, regularity typically refers to the existence of a number of strong or weak partial derivatives, which is encapsulated in the frameworks of Hölder and Sobolev regularity respectively.
2. But we also want the GP posterior to provide a useful uncertainty estimate over the solution of the PDE. Hence, we do not want to needlessly impose *additional* regularity constraints on the sample paths. The sample paths should be **as irregular as possible**, to avoid overconfidence. Such cautious models may also be advisable on numerical grounds, to avoid instabilities such as Gibbs or Runge phenomena.

The above is just one example of a setting in which one would like to characterise the regularity of GP sample paths as tightly as possible. Further examples include generalised coordinates, a method that underlies the analysis and filtering of stochastic differential equations driven by noise admitting a high number of derivatives (Friston et al., 2023; Heins et al., 2023). These noise signals are usually modelled with GPs, and hence characterising GP sample path regularity is crucial for understanding in which situations these methods can be applied.

1.1 Summary of Contributions

A consequence of our main result, Theorem 7, may be written as follows:

Corollary 1 *Let $k: \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ be a symmetric positive definite kernel. If either*

- *all partial derivatives of the form $\frac{\partial^{2n} k}{\partial \mathbf{x}^\alpha \partial \mathbf{y}^\alpha}$ for multi-indices $\alpha \in \mathbb{N}_0^d$ with $|\alpha| = n$ exist and are (locally) Lipschitz,*
- *for stationary $k(\mathbf{x}, \mathbf{y}) = k_\delta(\mathbf{x} - \mathbf{y})$, $\frac{\partial^{2n} k_\delta}{\partial \mathbf{x}^\alpha}$ for $|\alpha| = 2n$ exist and are (locally) Lipschitz,*
- *for isotropic $k(\mathbf{x}, \mathbf{y}) = k_r(\|\mathbf{x} - \mathbf{y}\|)$, $k_r^{(2n)}$ exists and is (locally) Lipschitz,*

then the sample paths of $f \sim \mathcal{GP}(0, k)$ are n times continuously differentiable.

Note that, by the mean value theorem, the existence of one additional continuous derivative of k is sufficient for local Lipschitz continuity of the lower derivatives.

The actual Theorem 7 is sharper than this. It provides necessary and sufficient conditions on the regularity of the kernel for the sample paths to attain a given Hölder regularity. Just as the above corollary, the theorem encompasses statements for stationary and isotropic GPs.

We apply Theorem 7 in particular to the Matérn GPs to obtain Proposition 10, a consequence of which is the following:

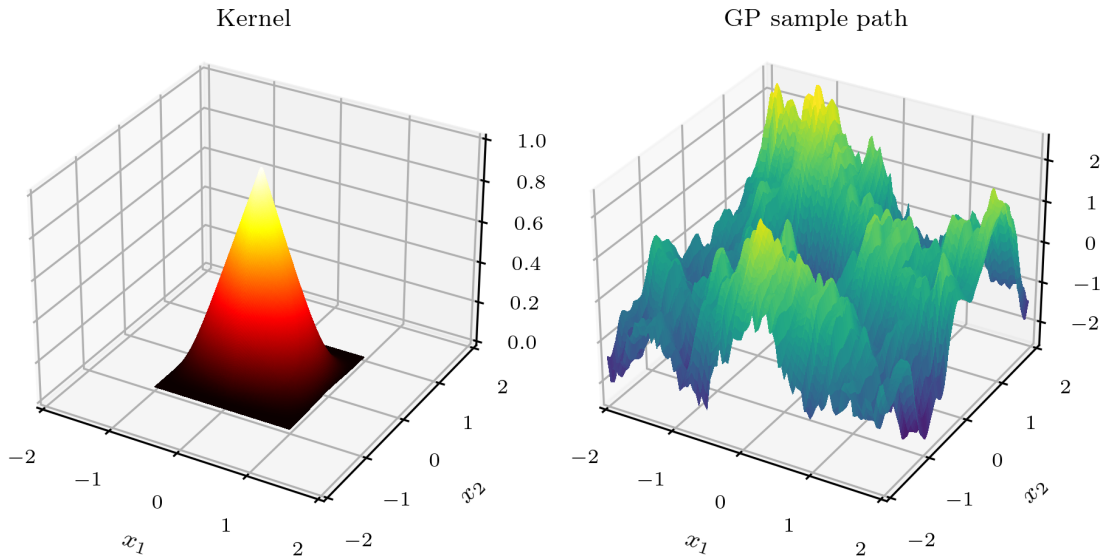


Figure 1: *Left:* surface plot of a kernel defined as the tensor product of two Wendland kernels (with $d = 1, n = 0$ on the x_1 -axis and $d = 1, n = 1$ on the x_2 -axis). *Right:* surface plot of a sample path from the corresponding centered GP. The sample path displays different regularity along each axis (non-differentiable along the x_1 -axis; once differentiable along the x_2 -axis), reflecting the rotational asymmetry in the regularity of the kernel.

Corollary 2 *The sample paths of a centered Matérn GP with smoothness parameter $\nu \notin \mathbb{N}_0$ are $\lfloor \nu \rfloor$ times continuously differentiable and no more. In particular, when $\nu = n + 1/2$ for some $n \in \mathbb{N}_0$, the sample paths are n times continuously differentiable and no more.*

We then obtain similar results for other GPs, such as the Wendland (Proposition 12), squared exponential and rational quadratic GPs (Remark 14).

In Section 4.6 we describe how to apply Theorem 7 to covariance kernels constructed through algebraic transformations of other kernels, including conic combinations (Proposition 17), products (Proposition 18), tensor products (Proposition 21) and coordinate transformations (Proposition 23). We note in Section 4.7 how the results in this work may be generalised to manifold GPs.

Theorem 25 summarises results about GP sample path Sobolev regularity. Comparing these with Theorem 7, we argue that one should not expect the existence of a greater number of weak (Sobolev) derivatives than of strong derivatives for GP sample paths. This provides further evidence that Theorem 7 contains all that is needed in practice.

We then plot a number of stationary kernel functions as well as sample paths from the corresponding centered GPs, as in Figure 1.

All regularity properties considered in this work are *local*, in the sense that they capture the infinitesimal behaviour of functions around each point in their domain. Note that our proofs can be extended and our assumptions strengthened to consider global regularity properties, such as global Hölder continuity.

1.2 Related Work

Some of the earliest work on sample path regularity can be attributed to Kolmogorov and Chentsov (1956). Later, Fernique (1975) proved the first necessary and sufficient conditions for sample path continuity of stationary GPs. Since then, many authors have extended these continuity results. Potthoff (2009) provides an approachable overview of how such results can be obtained. It also investigates conditions for uniform continuity, which is a *global* regularity property. Azmoodeh et al. (2014), which we rely on in this work, proves a converse to the Kolmogorov continuity theorem for GPs, providing necessary and sufficient conditions for Hölder continuity.

Differentiability has also been investigated by many authors, including Scheuerer (2010b) for general random fields, Adler and Taylor (2007), Potthoff (2010) and Henderson (2024) for Sobolev regularity. The sample path regularity of Matérn GPs was studied in Scheuerer (2010a) (see specifically Scheuerer (2010a, Examples 5.3.19, 5.5.6 & 5.5.12)) and of tensor product of one dimensional Matérn GPs in Wang et al. (2021). Pierce (1970, Proposition I.3) combines differentiability results with the Kolmogorov continuity theorem, and can be seen as the global version of the forward implication of our Theorem 7 (1).

Compared to the existing literature, the present work provides *general, versatile, easily applicable* and *unusually tight* results for the study of GP sample path continuity and differentiability.

2 Preliminaries

Definition 3 *A Gaussian process on a set O is a map $f: O \times \Omega \rightarrow \mathbb{R}$, where Ω is a probability space, such that for all $x \in O$, $f(x, \cdot): O \rightarrow \mathbb{R}$ is measurable, and such that for each $X := (x_1, \dots, x_N) \in O^N$, the map $f(X, \cdot): \Omega \rightarrow \mathbb{R}^N$ given by $f(X, \omega) = (f(x_1, \omega), \dots, f(x_N, \omega))$ is a multivariate Gaussian random variable.*

The maps $f(\cdot, \omega): O \rightarrow \mathbb{R}$ for $\omega \in \Omega$ are the sample paths of the GP.

We would like to study the regularity of the sample paths of a GP. However, Definition 3 is not convenient to work with in practice: one rarely constructs the probability space Ω of a GP. Instead, one characterises a GP by its *mean* and its *covariance kernel*.

Definition 4 *The mean μ of a GP $f: O \times \Omega \rightarrow \mathbb{R}$ is the map*

$$\mu: O \rightarrow \mathbb{R}, x \mapsto \mathbb{E}(f(x, \cdot)).$$

The covariance kernel k is the map

$$k: O \times O \rightarrow \mathbb{R}, (x, y) \mapsto \text{cov}(f(x, \cdot), f(y, \cdot)).$$

The GP f is said to be centered if $\mu = 0$.

Conversely, given maps $\mu: O \rightarrow \mathbb{R}$ and $k: O \times O \rightarrow \mathbb{R}$ with k symmetric positive definite – meaning that for any finite set of points $\{x_1, \dots, x_N\} \subset O$, the matrix $(k(x_i, x_j))_{i,j}$ is symmetric positive semi-definite – one can construct a probability space Ω , and a GP $f: O \times \Omega \rightarrow \mathbb{R}$ with mean μ and covariance kernel k (Klenke, 2014, Theorem 14.36). We then write $f \sim \mathcal{GP}(\mu, k)$.

At this point it is important to note that $f \sim \mathcal{GP}(\mu, k)$ does not uniquely specify f or the probability space Ω . However it uniquely specifies the finite dimensional distributions of f – since multivariate Gaussians are entirely characterised by their first two moments. In machine learning applications all that will ever be observed are evaluations of f at finitely many points. Therefore we would like to study sample path regularity of f *up to modification*. By this we mean that $f \sim \mathcal{GP}(\mu, k)$ will be said to have samples in $\mathcal{F}(O)$, where $\mathcal{F}(O)$ is some space of functions $O \rightarrow \mathbb{R}$, if there exists a construction of the GP f , say $f: O \times \Omega \rightarrow \mathbb{R}$, such that $f(\cdot, \omega) \in \mathcal{F}(O)$ for all $\omega \in \Omega$.

Finally, note that $f \sim \mathcal{GP}(\mu, k)$ having samples in $\mathcal{F}(O)$ is equivalent to $\tilde{f} + \mu$ having samples in $\mathcal{F}(O)$, where $\tilde{f} \sim \mathcal{GP}(0, k)$. For this it suffices that $\mu \in \mathcal{F}(O)$ and \tilde{f} has samples in $\mathcal{F}(O)$. Therefore, in what follows, we will consider only centered GPs.

Thus we are studying sample path regularity of GPs *from the covariance kernel*, and our goal is to link the regularity of the kernel to the regularity of the GP sample paths.

2.1 Setup

We study GPs defined on open subsets of Euclidean spaces: $O \subset \mathbb{R}^d$ is open, for some $d \in \mathbb{N}$, and $f \sim \mathcal{GP}(0, k)$ is a centered GP on O . An important special case is when k is stationary, i.e. $O = \mathbb{R}^d$ and there is an even function $k_\delta: \mathbb{R}^d \rightarrow \mathbb{R}$ such that

$$k(\mathbf{x}, \mathbf{y}) = k_\delta(\mathbf{x} - \mathbf{y}) \tag{1}$$

for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$. A further special case of k stationary is when k is isotropic, i.e. $O = \mathbb{R}^d$ and there is an even function $k_r: \mathbb{R} \rightarrow \mathbb{R}$ such that

$$k(\mathbf{x}, \mathbf{y}) = k_r(\|\mathbf{x} - \mathbf{y}\|) \tag{2}$$

for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, where $\|\cdot\|$ denotes the standard Euclidean norm on \mathbb{R}^d .

3 Hölder Regularity

The main type of sample path regularity we would like to investigate is continuity and the order of continuous differentiability. These are encapsulated in the framework of *Hölder regularity*.

We choose O to be an open set in \mathbb{R}^d in order to capture differentiability of the sample paths. However this implies that Hölder spaces on O are inconvenient for our purposes, as they do not solely characterise local regularity properties of functions on O . Indeed, the Hölder spaces not only constrain the local behaviour of functions, but also their behaviour at infinity, or near the boundary of O . To retain solely local constraints we employ local Hölder spaces.

Definition 5 (Local Hölder and almost-Hölder spaces) *Let $n \in \mathbb{N}_0$ and $\gamma \in [0, 1]$.*

- (1) *The local Hölder space $C_{loc}^{n, \gamma}(O)$ is the space of functions f on O for which $\partial^\alpha f$ exists for all multi-indices $\alpha = (\alpha_1, \dots, \alpha_d) \in \mathbb{N}_0^d$ with $|\alpha| := \alpha_1 + \dots + \alpha_d \leq n$, and such that the highest order partial derivatives satisfy a Hölder condition of the form: for all compact subsets $K \subset O$ there is a constant $C_K > 0$ such that*

$$|\partial^\alpha f(\mathbf{x}) - \partial^\alpha f(\mathbf{y})| \leq C_K \|\mathbf{x} - \mathbf{y}\|^\gamma$$

for all $\mathbf{x}, \mathbf{y} \in K$ and $|\boldsymbol{\alpha}| = n$.

(2) The local almost-Hölder space $C_{loc}^{(n+\gamma)^-}(O)$ is defined as $\bigcap_{n'+\gamma' < n+\gamma} C_{loc}^{n',\gamma'}(O)$.

Remark 6 For $n, n' \in \mathbb{N}_0$ and $\gamma, \gamma' \in [0, 1]$, $n'+\gamma' < n+\gamma$ implies $C_{loc}^{n',\gamma'}(O) \supsetneq C_{loc}^{(n+\gamma)^-}(O) \supsetneq C_{loc}^{n,\gamma}(O)$. Moreover $C^n(O) = C_{loc}^{n,0}(O)$.

The key point in the definition of $C_{loc}^{(n+\gamma)^-}(O)$ is the strict inequality $n'+\gamma' < n+\gamma$, since allowing for equality we have $\bigcap_{n'+\gamma' \leq n+\gamma} C_{loc}^{n',\gamma'}(O) = C_{loc}^{n,\gamma}(O)$.

As will become clear, the local almost-Hölder spaces are natural for characterising GP sample path regularity.

To describe the regularity of the covariance kernel k , we will require derivatives of the form $\partial^{\boldsymbol{\alpha},\boldsymbol{\beta}}k$. Here $\partial^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ stands for the application of $\partial^{\boldsymbol{\alpha}}$ with respect to the first variable, followed by the application of $\partial^{\boldsymbol{\beta}}$ with respect to the second variable. In fact the order in which they are applied will not matter, by continuity of these partial derivatives. We moreover write $C^{n \otimes n}(O \times O)$ for the space of functions k on $O \times O$ for which $\partial^{\boldsymbol{\alpha},\boldsymbol{\beta}}k$ exists and is continuous for all $|\boldsymbol{\alpha}|, |\boldsymbol{\beta}| \leq n$.

Also recall the “big O ” notation: the functions $f, g: \mathbb{R}^d \setminus \{\mathbf{0}\} \rightarrow \mathbb{R}$ satisfy $f(\mathbf{h}) = \mathcal{O}(g(\mathbf{h}))$ as $\mathbf{h} \rightarrow \mathbf{0}$ if and only if there is $C > 0$ such that $\limsup_{\mathbf{h} \rightarrow \mathbf{0}} |f(\mathbf{h})/g(\mathbf{h})| \leq C$. Moreover, for a family of functions $(f_{\mathbf{x}})_{\mathbf{x} \in O}$, we say $f_{\mathbf{x}}(\mathbf{h}) = \mathcal{O}(g(\mathbf{h}))$ as $\mathbf{h} \rightarrow \mathbf{0}$ locally uniformly in $\mathbf{x} \in O$ if, for every compact subset $K \subset O$, $C > 0$ may be chosen independently of $\mathbf{x} \in K$.

We are now in a position to state the main result of this paper.

Theorem 7 (Sample path Hölder regularity) Let $n \in \mathbb{N}_0$ and $\gamma \in (0, 1]$. The process $f \sim \mathcal{GP}(0, k)$ has samples in $C_{loc}^{(n+\gamma)^-}(O)$ if and only if,

(1) for general k ,

- $k \in C^{n \otimes n}(O \times O)$,
- $|\partial^{\boldsymbol{\alpha},\boldsymbol{\beta}}k(\mathbf{x}+\mathbf{h}, \mathbf{x}+\mathbf{h}) - \partial^{\boldsymbol{\alpha},\boldsymbol{\beta}}k(\mathbf{x}+\mathbf{h}, \mathbf{x}) - \partial^{\boldsymbol{\alpha},\boldsymbol{\beta}}k(\mathbf{x}, \mathbf{x}+\mathbf{h}) + \partial^{\boldsymbol{\alpha},\boldsymbol{\beta}}k(\mathbf{x}, \mathbf{x})| = \mathcal{O}(\|\mathbf{h}\|^{2\epsilon})$ as $\mathbf{h} \rightarrow \mathbf{0}$, locally uniformly in $\mathbf{x} \in O$, for all $\epsilon \in (0, \gamma)$ and $|\boldsymbol{\alpha}| = |\boldsymbol{\beta}| = n$.

(2) for stationary $k(\mathbf{x}, \mathbf{y}) = k_{\delta}(\mathbf{x} - \mathbf{y})$,

- $k_{\delta} \in C^{2n}(\mathbb{R}^d)$,
- $|\partial^{\boldsymbol{\alpha}}k_{\delta}(\mathbf{h}) - \partial^{\boldsymbol{\alpha}}k_{\delta}(\mathbf{0})| = \mathcal{O}(\|\mathbf{h}\|^{2\epsilon})$ as $\mathbf{h} \rightarrow \mathbf{0}$ for all $\epsilon \in (0, \gamma)$ and $|\boldsymbol{\alpha}| = 2n$.

(3) for isotropic $k(\mathbf{x}, \mathbf{y}) = k_r(\|\mathbf{x} - \mathbf{y}\|)$,

- $k_r \in C^{2n}(\mathbb{R})$,
- $|k_r^{(2n)}(h) - k_r^{(2n)}(0)| = \mathcal{O}(|h|^{2\epsilon})$ as $h \rightarrow 0$ for all $\epsilon \in (0, \gamma)$.

In each case, differentiating sample-wise we have $\partial^{\boldsymbol{\alpha}}f \sim \mathcal{GP}(0, \partial^{\boldsymbol{\alpha},\boldsymbol{\alpha}}k)$ for all $|\boldsymbol{\alpha}| \leq n$.

The proof can be found in Appendix A.

Remark 8 *The positive definiteness of k allows us to extrapolate regularity: k , k_δ or k_r have no more derivatives around the diagonal $\{(\mathbf{x}, \mathbf{x}) \in O \times O : \mathbf{x} \in O\}$, $\mathbf{0}$ or 0 respectively as they do on the rest of the domain (see Lemma A.3). Therefore, when applying Theorem 7, it suffices to check the existence of the respective derivatives around the diagonal, $\mathbf{0}$ or 0 respectively.*

Another way to see that information about k in a neighbourhood of the diagonal must be sufficient to characterise any local sample path regularity property, as it must suffice to consider the covariances at arbitrarily close by points.

Theorem 7 combines results about sample Hölder continuity of GPs (Azmoodeh et al., 2014) to results about sample differentiability of GPs (Potthoff, 2010), applied inductively on the partial derivatives. It provides necessary and sufficient conditions for f to have almost-Hölder continuous samples of a certain degree. Let us note however that it does not quite give us necessary conditions for f to have samples in $C^n(O)$, as the latter is not an almost-Hölder space. To achieve this, we can adapt the proof of Theorem 7 combining necessary and sufficient conditions for sample continuity of GPs with the results about sample differentiability from Potthoff (2010) (see Scheuerer (2010a, Theorem 5.3.16) for a result of this flavour). Characterising necessary and sufficient conditions on the kernel for the sample continuity of the GP is much more involved than almost-Hölder continuity, and therefore we do not expand our results to this setting (see Fernique (1975) for the stationary case, Talagrand (1987) for the general case). Moreover we demonstrate in the examples below that Theorem 7 is all we need.

4 Examples

In this section we demonstrate how Theorem 7 can be applied in practice. We investigate the sample path regularity of the most widely used GP families, recovering known results as well as proving novel ones.

The reverse implications in Theorem 7 allow us to obtain sharp sample path regularity characterisations in the following sense: f is said to have samples in $C_{loc}^{(n+\gamma)^-}(O)$ and no more, if f is not sample $C_{loc}^{(n'+\gamma')^-}(O)$ for any $n' + \gamma' > n + \gamma$.

4.1 Wiener Kernel

The Wiener process is a centered GP on $O = \mathbb{R}_{>0}$ with covariance kernel

$$k(x, y) = \min(x, y)$$

for $x, y \in \mathbb{R}_{>0}$. $\min(\cdot, \cdot)$ is Lipschitz but non-differentiable in both its arguments. So by Theorem 7(1) the Wiener process has samples in $C_{loc}^{1/2^-}(\mathbb{R}_{>0})$ and no more.

Remark 9 *By the fundamental theorem of calculus, the n -times integrated Wiener process (Schober et al., 2014, Section B) has samples in $C_{loc}^{(n+1/2)^-}(\mathbb{R}_{>0})$ and no more.*

4.2 Matérn Kernels

The Matérn kernels are isotropic kernels on \mathbb{R}^d given by

$$k(\mathbf{x}, \mathbf{y}) = k_r(\|\mathbf{x} - \mathbf{y}\|) = \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu} \|\mathbf{x} - \mathbf{y}\| \right)^\nu K_\nu \left(\sqrt{2\nu} \|\mathbf{x} - \mathbf{y}\| \right) \quad (3)$$

for $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, where K_ν is the modified Bessel function of the second kind and $\nu > 0$ is the smoothness parameter. The cases where $\nu = n + 1/2$ for some $n \in \mathbb{N}_0$ are particularly interesting because the expression in Equation (3) simplifies to a product of a polynomial and an exponential of the radial distance (Rasmussen and Williams, 2005, Equation 4.16).

Proposition 10 *A centered Matérn GP with smoothness parameter $\nu > 0$ has samples in $C_{loc}^{\nu-}(O)$ and no more. In particular, if $\nu = n + 1/2$ for some $n \in \mathbb{N}_0$, the GP has samples in $C_{loc}^{(n+1/2)-}(\mathbb{R}^d)$ and no more.*

This follows from Theorem 7 (3) and the proof can be found in Appendix B.1.

Remark 11 *Proposition 10 also covers the (multivariate) Ornstein-Uhlenbeck process, the centered GP with covariance kernel given by $k(\mathbf{x}, \mathbf{y}) = \exp(-\|\mathbf{x} - \mathbf{y}\|)$ which corresponds to the Matérn GP with $\nu = 1/2$. So the (multivariate) Ornstein-Uhlenbeck process has samples in $C_{loc}^{1/2-}(\mathbb{R}^d)$ and no more.*

4.3 Wendland Kernels

The Wendland kernels are isotropic kernels which are compactly supported and piecewise polynomial in the radial distance. On \mathbb{R}^d , they are defined as follows (Wendland, 2004, Definition 9.11):

$$k(\mathbf{x}, \mathbf{y}) = k_r(\|\mathbf{x} - \mathbf{y}\|) = \mathcal{I}^n \phi_{\lfloor d/2 \rfloor + n + 1}(\|\mathbf{x} - \mathbf{y}\|)$$

for $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, where $\phi_j(\rho) := \max((1 - \rho)^j, 0)$ and $\mathcal{I}\phi(\rho) := \int_\rho^\infty t\phi(t) dt / \int_0^\infty t\phi(t) dt$ for $\rho \geq 0$. $n \in \mathbb{N}_0$ is a parameter controlling the degree of the polynomial (precisely, $k_r(x)$ has degree $\lfloor d/2 \rfloor + 3n + 1$ in ρ around 0).

Proposition 12 *A centered Wendland GP with degree parameter n has samples in $C_{loc}^{(n+1/2)-}(O)$ and no more.*

This follows from Theorem 7 (3) and the proof can be found in Appendix B.2.

4.4 Smooth Kernels

From Theorem 7 immediately follows the following important corollary:

Corollary 13 *The process $f \sim \mathcal{GP}(0, k)$ has samples in $C^\infty(O)$ if and only if,*

- (1) *for general k , $k \in C^\infty(O \times O)$.*
- (2) *for stationary $k(\mathbf{x}, \mathbf{y}) = k_\delta(\mathbf{x} - \mathbf{y})$, $k_\delta \in C^\infty(\mathbb{R}^d)$.*
- (3) *for isotropic $k(\mathbf{x}, \mathbf{y}) = k_r(\|\mathbf{x} - \mathbf{y}\|)$, $k_r \in C^\infty(\mathbb{R})$.*

Remark 14 *The centered GPs with squared exponential ($k(\mathbf{x}, \mathbf{y}) = \exp(-\|\mathbf{x} - \mathbf{y}\|^2)$), rational quadratic ($k(\mathbf{x}, \mathbf{y}) = (1 + \|\mathbf{x} - \mathbf{y}\|^2)^{-a}$ for some $a > 0$) and periodic ($k(x, y) = \exp(-\sin(\pi(x - y))^2)$ for $x, y \in \mathbb{R}$) covariance kernels are therefore sample $C^\infty(\mathbb{R}^d)$.*

4.5 Feature Kernels

For any map $\phi: O \rightarrow \mathbb{R}^m$ is associated a feature kernel

$$k(\mathbf{x}, \mathbf{y}) = \phi(\mathbf{x})^\top \phi(\mathbf{y})$$

for $\mathbf{x}, \mathbf{y} \in O$. From Theorem 7(1) we obtain the following result:

Proposition 15 *For $n \in \mathbb{N}_0$ and $\gamma \in (0, 1]$, if $\phi_i \in C_{loc}^{(n+\gamma)^-}(O)$ for all $1 \leq i \leq m$ then $f \sim \mathcal{GP}(0, k)$ has samples in $C_{loc}^{(n+\gamma)^-}(O)$.*

The proof can be found in Appendix B.3.

Remark 16 *As a special case, linear kernels ($k(\mathbf{x}, \mathbf{y}) = \langle \mathbf{x}, \mathbf{y} \rangle$) and polynomial kernels ($k(\mathbf{x}, \mathbf{y}) = (1 + \langle \mathbf{x}, \mathbf{y} \rangle)^m$ for some $m \in \mathbb{N}$) are feature kernels, with the feature maps being the coordinate maps or polynomials of the coordinate maps respectively. Therefore the associated centered GPs have samples in $C^\infty(\mathbb{R}^d)$.*

4.6 Kernel Algebra

The versatility of the GP framework in applications is largely due to the fact that covariance kernels can be recursively combined using algebraic operations like pointwise sums or products which allows for constructing highly flexible and complex prior models in a systematic and interpretable fashion. It is therefore important to understand how the sample path regularity of a GP depends on the recursive structure of the covariance kernel.

In the following the k_i are always assumed to be symmetric positive definite kernels. Note that in this section we work solely with general kernels, leveraging Theorem 7(1). It would then be straightforward to similarly leverage Theorem 7(2) and Theorem 7(3) for stationary and isotropic kernels respectively.

4.6.1 CONIC COMBINATIONS

Conic combinations of kernels are pointwise linear combinations of the kernel functions with positive weights, i.e.

$$k(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^m a_i k_i(\mathbf{x}, \mathbf{y})$$

for $\mathbf{x}, \mathbf{y} \in O$, where $a_i > 0$ for all $1 \leq i \leq m$. Sums of kernels and positive scaling of kernels are special cases of conic combinations with $a_i = 1$ and $m = 1$, respectively.

Proposition 17 *For $n \in \mathbb{N}_0$ and $\gamma \in (0, 1]$, if the k_i satisfy the condition in Theorem 7(1) for all $1 \leq i \leq m$ then $f \sim \mathcal{GP}(0, k)$ has samples in $C_{loc}^{(n+\gamma)^-}(O)$.*

This follows from the fact that for $f \sim \mathcal{GP}(0, k)$, we can express the sample paths as $f(\cdot, \omega) = \sum_{i=1}^m \sqrt{a_i} f_i(\cdot, \omega)$ where $f_i \sim \mathcal{GP}(0, k_i)$, and that $C_{loc}^{(n+\gamma)^-}(O)$ is a vector space.

4.6.2 PRODUCTS

Define the product kernel

$$k(\mathbf{x}, \mathbf{y}) = \prod_{i=1}^m k_i(\mathbf{x}, \mathbf{y})$$

for $\mathbf{x}, \mathbf{y} \in O$. k is also positive definite (Berg et al., 1984, Chapter 3 Theorem 1.12).

Proposition 18 *For $n \in \mathbb{N}_0$ and $\gamma \in (0, 1]$, if the k_i satisfy the condition in Theorem 7 (1) for all $1 \leq i \leq m$ then $f \sim \mathcal{GP}(0, k)$ has samples in $C_{loc}^{(n+\gamma)^-}(O)$.*

The proof can be found in Appendix B.4.

Remark 19 *Combining Proposition 18 with Proposition 15 we get that if \tilde{k} satisfies the condition in Theorem 7 (1) for some $n \in \mathbb{N}_0$ and $\gamma \in (0, 1]$, and $\phi \in C_{loc}^{(n+\gamma)^-}(O)$, then defining $k(\mathbf{x}, \mathbf{y}) = \phi(\mathbf{x})\tilde{k}(\mathbf{x}, \mathbf{y})\phi(\mathbf{y})$ for all $\mathbf{x}, \mathbf{y} \in O$, $f \sim \mathcal{GP}(0, k)$ has samples in $C_{loc}^{(n+\gamma)^-}(O)$. We could equivalently prove this by noting that we can express the sample paths as $f(\cdot, \omega) = \phi(\cdot)\tilde{f}(\cdot, \omega)$ where $\tilde{f} \sim \mathcal{GP}(0, \tilde{k})$.*

4.6.3 TENSOR PRODUCTS

Let $O_i \subset \mathbb{R}^{d_i}$ be open sets for $1 \leq i \leq m$ and $O_1 \times \dots \times O_m =: O \subset \mathbb{R}^d = \mathbb{R}^{d_1} \times \dots \times \mathbb{R}^{d_m}$. For $k_i: O_i \times O_i \rightarrow \mathbb{R}$ for $1 \leq i \leq m$, the tensor product kernel $k: O \times O \rightarrow \mathbb{R}$, is given by

$$k(\mathbf{x}, \mathbf{y}) = (k_1 \otimes \dots \otimes k_m)(\mathbf{x}, \mathbf{y}) := \prod_{i=1}^m k_i(\mathbf{x}_i, \mathbf{y}_i)$$

for $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, where $\mathbf{x}_i, \mathbf{y}_i \in \mathbb{R}^{d_i}$ are the orthogonal projections of \mathbf{x}, \mathbf{y} respectively onto \mathbb{R}^{d_i} .

GPs with tensor product covariance kernels are interesting as they allow *spherically asymmetric* sample path regularity. We need the following definition to formalise this.

Definition 20 *Let $n_1, \dots, n_m \in \mathbb{N}_0$ and $\gamma_1, \dots, \gamma_m \in [0, 1]$.*

- (1) $C_{loc}^{(n_1, \gamma_1) \otimes \dots \otimes (n_m, \gamma_m)}(O_1 \times \dots \times O_m)$ is the space of functions f on O for which $\partial^\alpha f$ exists for all multi-indices $\alpha \in \mathbb{N}_0^d$ with $|\alpha_i| \leq n_i$, where $\alpha_i \in \mathbb{N}_0^{d_i}$ is the projection of α onto $\mathbb{N}_0^{d_i}$, and such that the highest order partial derivatives satisfy a Hölder condition of the form: for all compact subsets $K \subset O$ there is a constant $C_K > 0$ such that

$$|\partial^\alpha f(\mathbf{x}) - \partial^\alpha f(\mathbf{y})| \leq C_K \|\mathbf{x}_i - \mathbf{y}_i\|^{\gamma_i}$$

for all $1 \leq i \leq d$, $\mathbf{x}, \mathbf{y} \in K$ with $\mathbf{x}_j = \mathbf{y}_j$ for $j \neq i$, and $|\alpha_i| = n_i$.

- (2) $C_{loc}^{(n_1+\gamma_1)^- \otimes \dots \otimes (n_m+\gamma_m)^-}(O_1 \times \dots \times O_m)$
 $:= \bigcap_{n'_i+\gamma'_i < n_i+\gamma_i} C_{loc}^{(n'_1, \gamma'_1) \otimes \dots \otimes (n'_m, \gamma'_m)}(O_1 \times \dots \times O_m)$.

The following proposition generalises Theorem 7.

Proposition 21 For $n_1, \dots, n_m \in \mathbb{N}_0$ and $\gamma_1, \dots, \gamma_m \in (0, 1]$, $f \sim \mathcal{GP}(0, k)$ has samples in $C_{loc}^{(n_1+\gamma_1)^- \otimes \dots \otimes (n_m+\gamma_m)^-} (O_1 \times \dots \times O_m)$ if and only if the k_i satisfy the condition in Theorem 7(1) with $n = n_i$ and $\gamma = \gamma_i$ for all $1 \leq i \leq m$.

The proof of Proposition 21 follows the same steps as the one of Theorem 7 in Appendix A.

Remark 22 We could generalise Theorem 7 to allow for even more flexibility in the asymmetrical derivatives in Theorem 7. This can be done by defining Hölder spaces extending the spaces $C^A(O)$, where $A \subset \mathbb{N}_0^d$ is a downward closed set of multi-indices which specifies which partial derivatives exist (Pförtner et al., 2022, Definition B.9).

4.6.4 COORDINATE TRANSFORMATIONS

For a map $\varphi: O \rightarrow \mathbb{R}^m$ and a symmetric positive definite kernel \tilde{k} defined on $\text{Im}(\varphi) \times \text{Im}(\varphi)$, we have a kernel

$$k(\mathbf{x}, \mathbf{y}) = \tilde{k}(\varphi(\mathbf{x}), \varphi(\mathbf{y})).$$

Proposition 23 For $n \in \mathbb{N}_0$ and $\gamma, \delta \in (0, 1]$, if \tilde{k} satisfies the condition in Theorem 7(1) and $\varphi_i \in C_{loc}^{(n+\delta)^-} (O)$ for all $1 \leq i \leq m$ then $f \sim \mathcal{GP}(0, k)$ has samples in $C_{loc}^{(n+\gamma\delta)^-} (O)$.

This follows from the fact that for $\tilde{f} \sim GP(0, \tilde{k})$ we have $\tilde{f} \circ \varphi \sim \mathcal{GP}(0, k)$, and that the composition of a $C_{loc}^{n, \zeta} (O)$ with a $C_{loc}^{m, \epsilon} (O)$ function is $C_{loc}^{m, \epsilon\zeta} (O)$.

4.7 Manifold Kernels

We can generalise the results in the present work to manifold domains; instead of considering GPs on open subsets $O \subset \mathbb{R}^d$, we consider GPs on manifolds M . There exists two ways of constructing positive definite kernels, and hence GPs, on manifolds M .

Extrinsic kernels are defined by viewing $M \subset \mathbb{R}^d$, taking a positive definite kernel k on \mathbb{R}^d , and restricting it to M . This restriction is analogous to a coordinate transformation (Section 4.6.4), since $k|_{M \times M}(\mathbf{x}, \mathbf{y}) = k(\iota(\mathbf{x}), \iota(\mathbf{y}))$ for $\mathbf{x}, \mathbf{y} \in M$, where $\iota: M \rightarrow \mathbb{R}^d$ is the inclusion map. Therefore, if the manifold is smooth, $k|_{M \times M}$ has the same sample regularity as k .

Intrinsic kernels are constructed directly on the manifold. Examples of these include the intrinsic Matérn and heat kernels (Borovitskiy et al., 2020), and the hyperbolic secant kernel (Da Costa et al., 2023). In this case, since the sample path regularity results in this work are all local, we may study sample path regularity in each coordinate patch $O \subset M$, treating them as open subsets of \mathbb{R}^d .

5 Sobolev Regularity

In this section we investigate how the regularity of the kernel affects the weak differentiability of the sample paths. Specifically, we consider the important L^2 -Sobolev regularity. The more general L^p -Sobolev regularity is studied in Henderson (2024). To characterise only the local regularity of the sample paths, as for Hölder spaces, we define the local pre-Sobolev spaces.

Definition 24 (Local pre-Sobolev spaces) Let $n \in \mathbb{N}_0$. The local pre-Sobolev space $\mathcal{H}_{loc}^n(O)$ is the space of functions¹ f on O for which for every compact $K \subset O$ the L^2 weak derivative $\partial_w^\alpha f$ exists on K for all multi-index $\alpha \in \mathbb{N}_0^d$ with $|\alpha| \leq n$, i.e. for all such α there is a function $\partial_w^\alpha f \in L^2(K)$ such that

$$\int_K \partial_w^\alpha f(\mathbf{x}) \varphi(\mathbf{x}) \, d\mathbf{x} = (-1)^{|\alpha|} \int_K f(\mathbf{x}) \partial^\alpha \varphi(\mathbf{x}) \, d\mathbf{x}$$

for all $\varphi \in C^\infty(K)$.

Theorem 25 (Sample path Sobolev regularity) Let $n \in \mathbb{N}$. The process $f \sim \mathcal{GP}(0, k)$ has samples in $\mathcal{H}_{loc}^n(O)$ if,

- (1) for general k , $k \in C^{m \otimes n}(O \times O)$.
- (2) for stationary $k(\mathbf{x}, \mathbf{y}) = k_\delta(\mathbf{x} - \mathbf{y})$, $\partial^\alpha k_\delta$ exists at $\mathbf{0}$ for all $|\alpha| \leq 2n$.
- (3) for isotropic $k(\mathbf{x}, \mathbf{y}) = k_r(\|\mathbf{x} - \mathbf{y}\|)$, $k_r^{(j)}$ exists at 0 for all $j \leq 2n$.

If we assume that k is continuous, or more generally that f has a measurable modification, then the converse holds in the scenarios (2) and (3).

Proof (1) follows from the general theorem for second order measurable random fields (Scheuerer, 2010b, Theorem 1). (2) and (3) follow from its corollary (Scheuerer, 2010b, Corollary 1). The converse in (2) follows from (Scheuerer, 2010b, Proposition 1). Finally, if (3) holds then, by the propagation of regularity from Gneiting (1999), we have $k_r \in C^{2n}(\mathbb{R})$. So by Lemma A.1 (3) \Rightarrow (2) (to be precise an adapted version without the Hölder conditions) we deduce (2), and hence we obtain the converse statement in this case too. \blacksquare

As observed in Scheuerer (2010b), Sobolev differentiability is a natural regularity notion for sample paths of general random fields as it can be deduced from continuous mean square differentiability. In the case of GPs however, if, in addition to the condition in Theorem 25 (1), we assume Hölder control on the highest order partial derivatives of the kernel at the diagonal as in Theorem 7 (1), we deduce strong differentiability of the sample paths. The Hölder condition in Theorem 7 (1) being a weak one, one can usually not obtain a greater number of weak derivatives using Theorem 25 than of strong derivatives using Theorem 7. The reverse implications in Theorem 25 (2) and Theorem 25 (3) show that, in the stationary and isotropic cases, this is not a limitation of the theorem, but an inherent property of stationary and isotropic GPs. For example, the proof of Proposition 10 Appendix B.1 shows that, by the converse to Theorem 25 (3), even in the edge case $\nu \in \mathbb{N}$, the samples of the Matérn GPs have as many Sobolev derivatives as of strong derivatives.

Note however that one way to construct non-stationary GPs with more weak derivatives than strong derivatives is through feature kernels, with features that admit more weak derivatives than strong derivatives. This also shows that it is not possible to obtain a converse for Theorem 25 (1) as we have for Theorem 25 (2) and Theorem 25 (3).

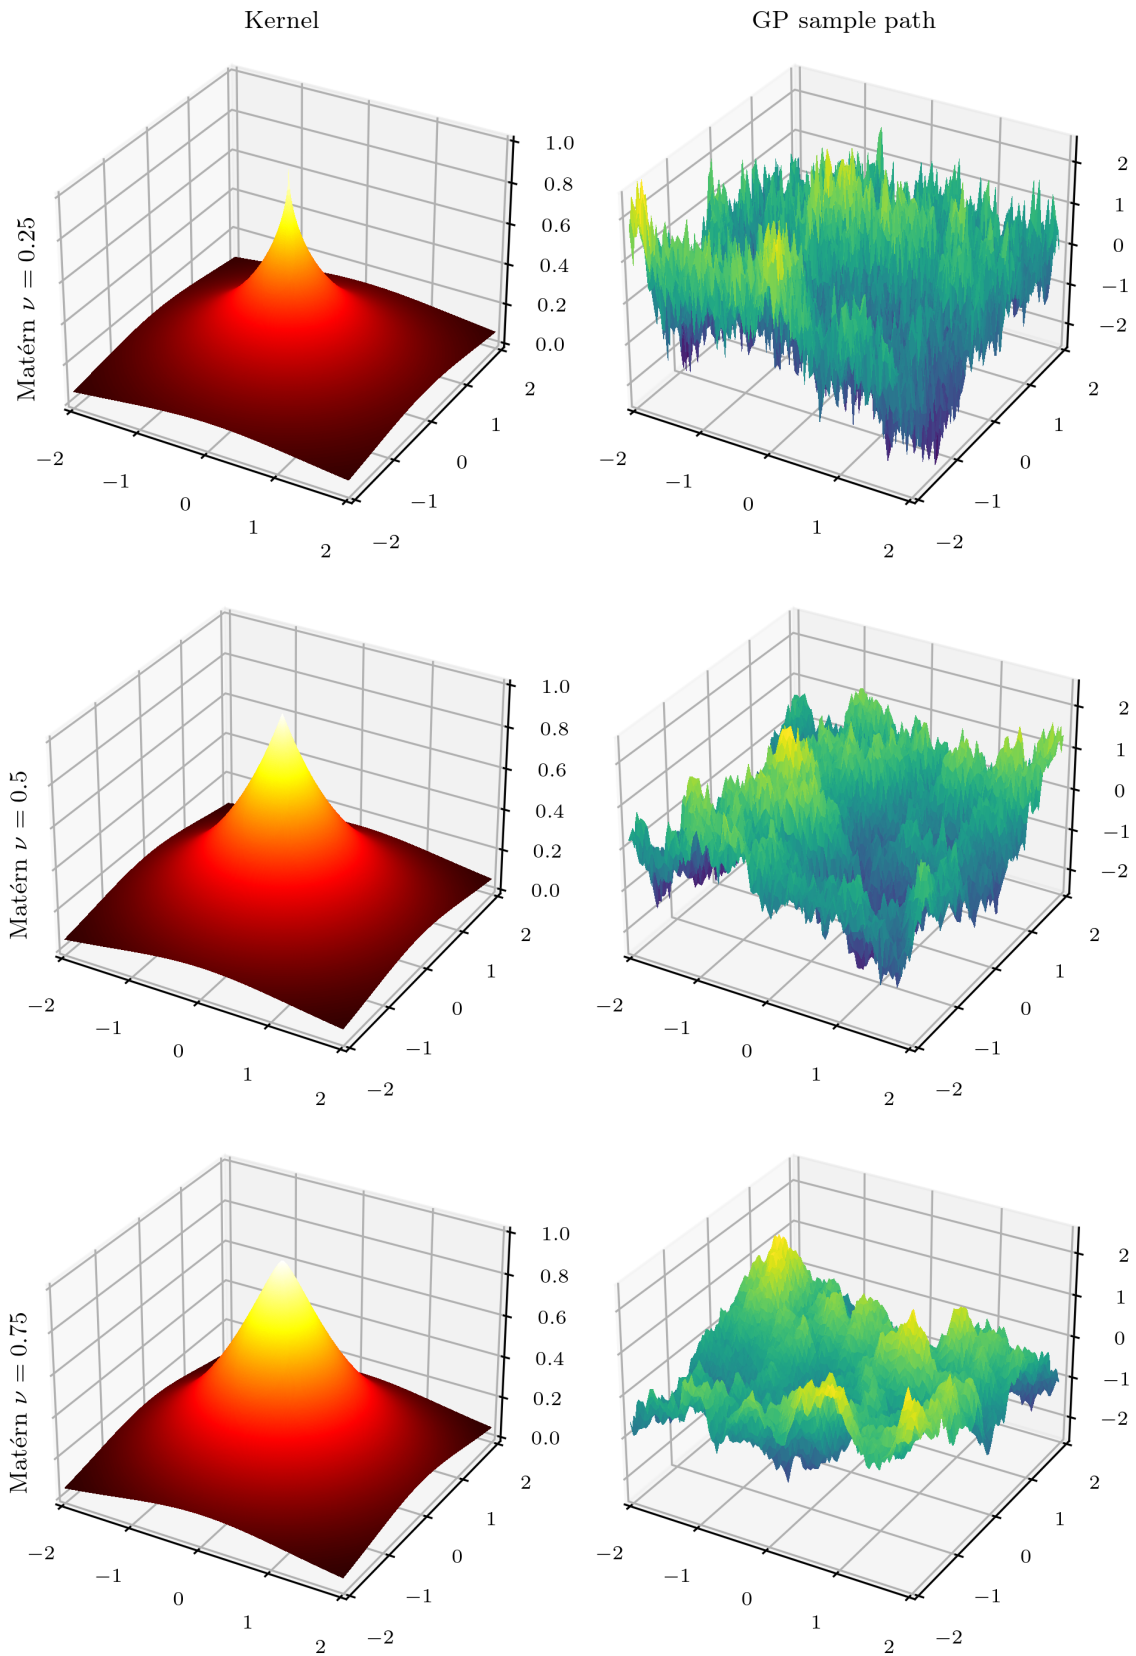
1. A (local) Sobolev space is an L^2 quotient of a (local) pre-Sobolev space. Since we are interested in function spaces, pre-Sobolev spaces are the right framework for us.

Also note that deducing sample path Hölder regularity from applying Theorem 25 and Sobolev embedding theorems introduces a superfluous dependence on dimension, and yields weaker results than the dimension independent Theorem 7.

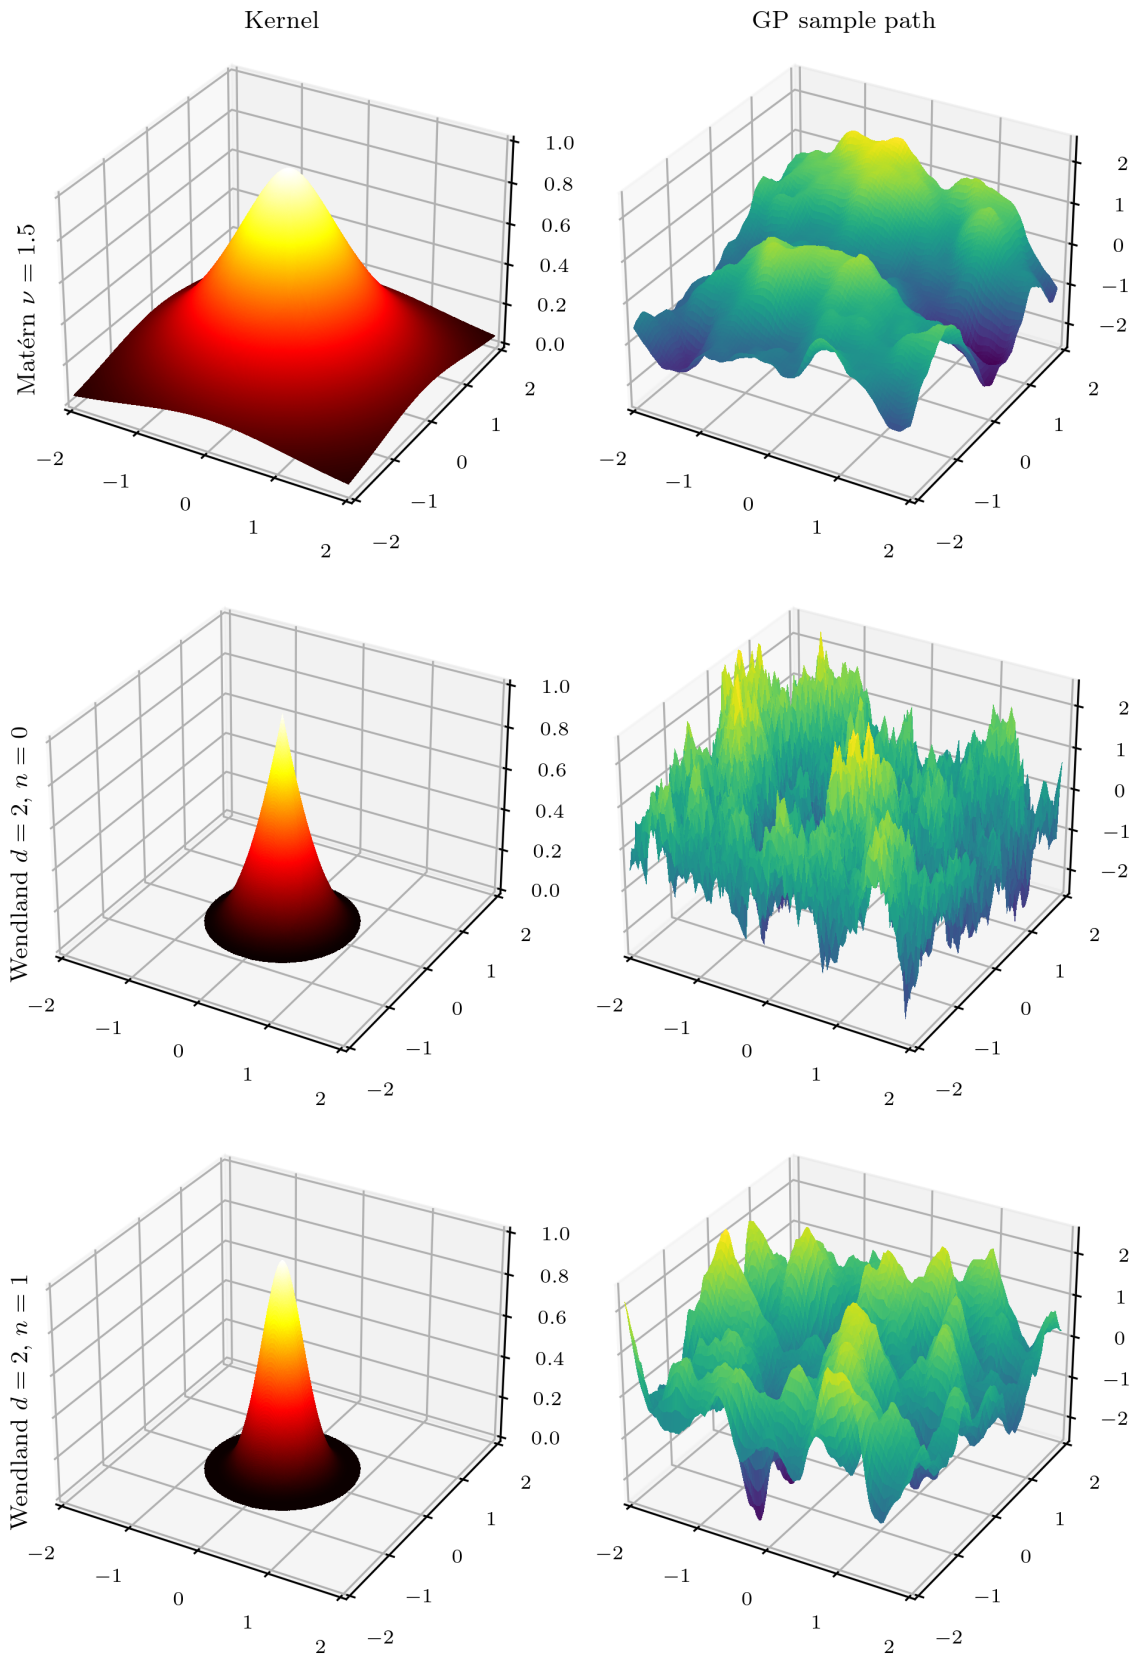
We therefore expect Theorem 25 to have less practical value for GPs as its general version (Scheuerer, 2010b, Theorem 1) for second order measurable random fields. For GPs we encourage the use of Theorem 7.

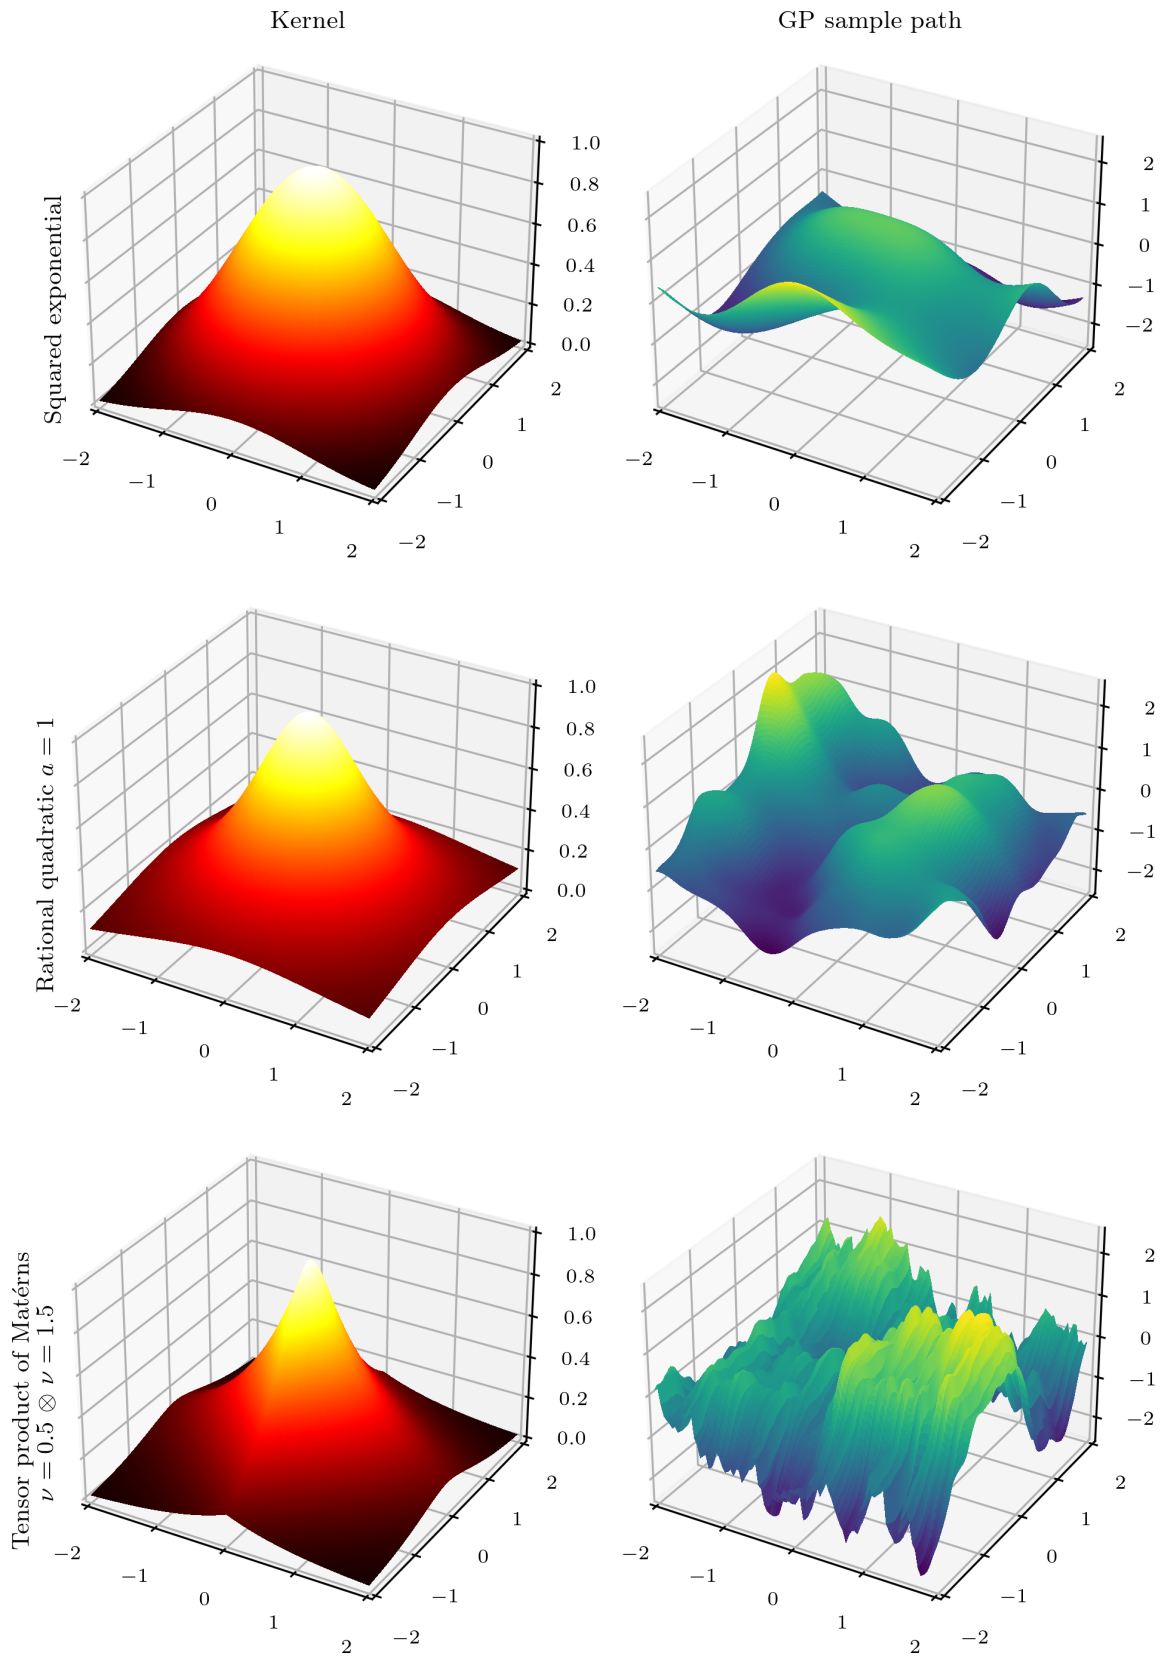
Acknowledgments and Disclosure of Funding

NDC, MP and PH gratefully acknowledge financial support by the European Research Council through ERC StG Action 757275 / PANAMA; the DFG Cluster of Excellence “Machine Learning - New Perspectives for Science”, EXC 2064/1, project number 390727645; the German Federal Ministry of Education and Research (BMBF) through the Tübingen AI Center (FKZ: 01IS18039A); and funds from the Ministry of Science, Research and Arts of the State of Baden-Württemberg. NDC acknowledges the support of the Fonds National de la Recherche, Luxembourg. The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting NDC and MP.



SAMPLE PATH REGULARITY OF GPs





Appendix A. Proof of Theorem 7

We start by showing the equivalences of the various kernel conditions in Theorem 7, namely (1) \Leftrightarrow (2) in the stationary case, and (1) \Leftrightarrow (2) \Leftrightarrow (3) in the isotropic case.

Lemma A.1 *Let $n \in \mathbb{N}_0$ and $\epsilon \in (0, 1]$. Then the condition*

- (1)
 - $k \in C^{n \otimes n}(O \times O)$,
 - $|\partial^{\alpha, \beta} k(\mathbf{x} + \mathbf{h}, \mathbf{x} + \mathbf{h}) - \partial^{\alpha, \beta} k(\mathbf{x} + \mathbf{h}, \mathbf{x}) - \partial^{\alpha, \beta} k(\mathbf{x}, \mathbf{x} + \mathbf{h}) + \partial^{\alpha, \beta} k(\mathbf{x}, \mathbf{x})| = \mathcal{O}(\|\mathbf{h}\|^{2\epsilon})$ as $\mathbf{h} \rightarrow \mathbf{0}$, locally uniformly in $\mathbf{x} \in O$, for all $|\alpha| = |\beta| = n$.

is equivalent to the following:

- (2) for stationary $k(\mathbf{x}, \mathbf{y}) = k_\delta(\mathbf{x}, \mathbf{y})$,
- $k_\delta \in C^{2n}(\mathbb{R}^d)$,
 - $|\partial^\alpha k_\delta(\mathbf{h}) - \partial^\alpha k_\delta(\mathbf{0})| = \mathcal{O}(\|\mathbf{h}\|^{2\epsilon})$ as $\mathbf{h} \rightarrow \mathbf{0}$ for all $|\alpha| = 2n$.
- (3) for isotropic $k(\mathbf{x}, \mathbf{y}) = k_r(\|\mathbf{x} - \mathbf{y}\|)$,
- $k_r \in C^{2n}(\mathbb{R})$,
 - $|k_r^{(2n)}(h) - k_r^{(2n)}(0)| = \mathcal{O}(|h|^{2\epsilon})$ as $h \rightarrow 0$.

Proof of Lemma A.1 (1) \Leftrightarrow (2) Let k be stationary, i.e. $k(\mathbf{x}, \mathbf{y}) = k_\delta(\mathbf{x} - \mathbf{y})$.

(2) \Rightarrow (1): For $|\alpha|, |\beta| \leq n$,

$$\partial^{\alpha, \beta} k(\mathbf{x}, \mathbf{y}) = (-1)^{|\beta|} \partial^{\alpha + \beta} k_\delta(\mathbf{x} - \mathbf{y})$$

for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, by the chain rule. So the existence and the continuity of $\partial^{\alpha + \beta} k_\delta$ implies the existence and the continuity of $\partial^{\alpha, \beta} k$. Moreover when $|\alpha| = |\beta| = n$,

$$\begin{aligned} & |\partial^{\alpha, \beta} k(\mathbf{x} + \mathbf{h}, \mathbf{x} + \mathbf{h}) - \partial^{\alpha, \beta} k(\mathbf{x} + \mathbf{h}, \mathbf{x}) - \partial^{\alpha, \beta} k(\mathbf{x}, \mathbf{x} + \mathbf{h}) + \partial^{\alpha, \beta} k(\mathbf{x}, \mathbf{x})| \\ &= 2|\partial^{\alpha + \beta} k_\delta(\mathbf{0}) - \partial^{\alpha + \beta} k_\delta(\mathbf{h})| \end{aligned} \quad (4)$$

for all $\mathbf{x}, \mathbf{h} \in \mathbb{R}^d$, so the Hölder conditions on the highest order partial derivatives of k at the diagonal and of k_δ at $\mathbf{0}$ correspond.

(1) \Rightarrow (2): For $|\alpha| \leq 2n$, say $\alpha = \beta + \gamma$ with $|\beta|, |\gamma| \leq n$, $\partial^\alpha k_\delta(\mathbf{x})$ exists and is equal to $(-1)^{|\gamma|} \partial^{\beta, \gamma} k(\mathbf{x}, \mathbf{0})$, for all $\mathbf{x} \in \mathbb{R}^d$. Moreover the Hölder conditions on the partial derivatives of k and k_δ correspond by Equation (4). ■

To prove the second equivalence in Lemma A.1, we need the following result:

Lemma A.2 *Let $m \in \mathbb{N}_0$, $g \in C^m(\mathbb{R} \setminus \{0\})$, and $\nu \in \mathbb{R}$ such that $g^{(j)}(h) = \mathcal{O}(|h|^{\nu-j})$ as $h \rightarrow 0$, for all $0 \leq j \leq m$. Define $f := g \circ \|\cdot\| : \mathbb{R}^d \setminus \{\mathbf{0}\} \rightarrow \mathbb{R}$. Then $f \in C^m(\mathbb{R}^d \setminus \{\mathbf{0}\})$ and $\partial^\alpha f(\mathbf{h}) = \mathcal{O}(\|\mathbf{h}\|^{\nu-j})$ as $\mathbf{h} \rightarrow \mathbf{0}$, for $0 \leq j \leq m$ and $|\alpha| = j$.*

Proof By induction on m . $m = 0$ is clear; the growth/decay of f at $\mathbf{0}$ is the same as that of g at 0. For $m > 0$, $\mathbf{x} \in \mathbb{R}^d \setminus \{\mathbf{0}\}$, and $1 \leq i \leq d$, note that

$$\partial^{e_i} f(\mathbf{x}) = g'(\|\mathbf{x}\|) \frac{x_i}{\|\mathbf{x}\|} = \tilde{g}(\|\mathbf{x}\|) x_i \quad (5)$$

where \mathbf{e}_i is the i^{th} unit vector and $\tilde{g}(x) := \frac{g'(x)}{x}$ for $x \in \mathbb{R} \setminus \{0\}$. Clearly $\tilde{g} \in C^{m-1}(\mathbb{R} \setminus \{0\})$ and

$$\tilde{g}^{(j)}(h) = \sum_{l=0}^j \binom{j}{l} g^{(l+1)}(h) (-1)^{j-l} (j-l)! h^{-1-j+l} = \mathcal{O}(|h|^{(\nu-2)-j})$$

as $h \rightarrow 0$, for all $0 \leq j \leq m-1$. Let $\tilde{f} = \tilde{g} \circ \|\cdot\|$. By the induction hypothesis, $\partial^\alpha \tilde{f}(\mathbf{h}) = \mathcal{O}(\|\mathbf{h}\|^{(\nu-2)-j})$ as $\mathbf{h} \rightarrow \mathbf{0}$, for $|\alpha| = j$ and $0 \leq j \leq m-1$. Thus, by (5),

$$\partial^\alpha \partial^{e_i} f(\mathbf{h}) = \begin{cases} h_i \partial^\alpha \tilde{f}(\mathbf{h}) & \text{if } \alpha_i = 0 \\ h_i \partial^\alpha \tilde{f}(\mathbf{h}) + \partial^{\alpha - e_i} \tilde{f}(\mathbf{h}) & \text{if } \alpha_i > 0 \end{cases} = \mathcal{O}(\|\mathbf{h}\|^{\nu-(j+1)})$$

as $\mathbf{h} \rightarrow \mathbf{0}$. Consequently, $\partial^\alpha f(\mathbf{h}) = \mathcal{O}(\|\mathbf{h}\|^{\nu-j})$ as $\mathbf{h} \rightarrow \mathbf{0}$ for $|\alpha| = j$ and $0 \leq j \leq m$. This concludes the induction step. \blacksquare

Proof of Lemma A.1 (1) \Leftrightarrow (3) Let k be isotropic, i.e. $k(\mathbf{x}, \mathbf{y}) = k_r(\|\mathbf{x} - \mathbf{y}\|)$. By Lemma A.1 (1) \Leftrightarrow (2), it suffices to show that the condition on k_δ in (2) is equivalent to the condition on k_r in (3).

(2) \Rightarrow (3): This follows from the fact that $x \mapsto k_\delta(x\mathbf{e}_1)$ is exactly k_r , and analogously $x \mapsto \partial^{j\mathbf{e}_1} k_\delta(x\mathbf{e}_1)$ is equal to $k_r^{(j)}$ for $0 \leq j \leq 2n$.

(3) \Rightarrow (2): $k_r: \mathbb{R} \rightarrow \mathbb{R}$ is even, so the odd derivatives of k_r which exist at 0 vanish there. Let

$$g(x) := k_r(x) - \sum_{j=0}^n \frac{x^{2j}}{j!} k_r^{(2j)}(0),$$

for all $x \in \mathbb{R}$. The Lagrange form of the remainder in an order $2n-j-1$ Taylor expansion of $g^{(j)}$ then reveals that

$$g^{(j)}(h) = \frac{h^{2n-j}}{(2n-j)!} (k_r^{(2n)}(\xi_j(h)) - k_r^{(2n)}(0)) = \mathcal{O}(|h|^{2n-j+2\epsilon})$$

as $h \rightarrow 0$, for all $0 \leq j \leq 2n$ and where $\xi_j(h) \in (0, h)$ (or $(h, 0)$ if $h < 0$). So $g|_{\mathbb{R} \setminus \{0\}}$ satisfies the conditions of Lemma A.2 with $m = 2n$ and $\nu = 2n + 2\epsilon$. Let $f := g \circ \|\cdot\|: \mathbb{R}^d \rightarrow \mathbb{R}$. Then, by Lemma A.2, $f|_{\mathbb{R}^d \setminus \{\mathbf{0}\}} \in C^{2n}(\mathbb{R}^d \setminus \{\mathbf{0}\})$, and $\partial^\alpha f(\mathbf{h}) = \mathcal{O}(\|\mathbf{h}\|^{2n+2\epsilon-j})$ as $\mathbf{h} \rightarrow \mathbf{0}$ for $|\alpha| = j$ and $0 \leq j \leq 2n$. In particular, $\partial^\alpha f(\mathbf{h}) \rightarrow 0$ as $\mathbf{h} \rightarrow \mathbf{0}$, and by the mean value theorem this is sufficient to deduce that $\partial^\alpha f(\mathbf{0})$ exists and is 0. Now by noting that $\|\mathbf{x}\|^{2j} = (x_1^2 + \dots + x_d^2)^j$ is smooth in \mathbf{x} for all j , we have that

$$k_\delta(\mathbf{x}) = f(\mathbf{x}) + \sum_{j=0}^n \frac{\|\mathbf{x}\|^{2j}}{j!} k_r^{(2j)}(0)$$

is in $C^{2n}(\mathbb{R})$. Moreover, for $\mathbf{h} \in \mathbb{R}^d$ and $|\alpha| = 2n$ there is a constant $C_\alpha \in \mathbb{R}$ such that $\partial^\alpha k_\delta(\mathbf{h}) = \partial^\alpha f(\mathbf{h}) + C_\alpha$, implying that

$$|\partial^\alpha k_\delta(\mathbf{h}) - \partial^\alpha k_\delta(\mathbf{0})| = |\partial^\alpha f(\mathbf{h}) - \underbrace{\partial^\alpha f(\mathbf{0})}_{=0}| = \mathcal{O}(\|\mathbf{h}\|^{2n+2\epsilon-2n}) = \mathcal{O}(\|\mathbf{h}\|^{2\epsilon})$$

as $\mathbf{h} \rightarrow \mathbf{0}$. Thus k_δ satisfies the Hölder condition in (2). \blacksquare

Note that Lemma A.1 is purely a result about stationary and isotropic functions; the proof does make use of the positive definiteness of k . In fact, the positive definiteness of k gives us additional information about its regularity, which can be seen through Lemma A.3 below.

For a GP f , the mean-square partial derivative $\partial_{ms}^{e_i} f$, when it exists, is a stochastic process defined on the same probability space such that $\partial_{ms}^{e_i} f(\mathbf{x})$ is the L^2 limit of $\frac{f(\mathbf{x}+h\mathbf{e}_i)-f(\mathbf{x})}{h}$ as $h \rightarrow 0$. The existence of a mean-square partial derivative is neither sufficient nor necessary for the existence of sample derivatives. However such derivatives are relevant for our work as they are directly related to the derivatives of the covariance kernel:

Lemma A.3 *Let $n \in \mathbb{N}_0$. For $f \sim \mathcal{GP}(0, k)$, the following conditions are equivalent:*

- (1) $k \in C^{n \otimes n}(O \times O)$,
- (2) $\partial^{\alpha, \alpha} k$ exists in a neighbourhood the diagonal and is continuous there for all $|\alpha| \leq n$,
- (3) the mean-square partial derivative $\partial_{ms}^\alpha f$ exists and is mean-square continuous for all $|\alpha| \leq n$.

Moreover, if either condition holds then we have $\partial_{ms}^\alpha f \sim \mathcal{GP}(0, \partial^{\alpha, \alpha} k)$ for all $|\alpha| \leq n$.

Proof We prove this for $n = 1$; for general n it then suffices to apply the same argument inductively on the partial derivatives.

(1) \Rightarrow (2): This implication is clear.

(2) \Rightarrow (3): Let $1 \leq i \leq d$. Note that

$$\begin{aligned}
 & \mathbb{E} \left[\left(\frac{f(\mathbf{x} + h\mathbf{e}_i) - f(\mathbf{x})}{h} \right) \left(\frac{f(\mathbf{y} + h'\mathbf{e}_i) - f(\mathbf{y})}{h'} \right) \right] \\
 &= \frac{\mathbb{E}[f(\mathbf{x} + h\mathbf{e}_i)f(\mathbf{y} + h'\mathbf{e}_i)] - \mathbb{E}[f(\mathbf{x} + h\mathbf{e}_i)f(\mathbf{y})] - \mathbb{E}[f(\mathbf{x})f(\mathbf{y} + h'\mathbf{e}_i)] + \mathbb{E}[f(\mathbf{x})f(\mathbf{y})]}{hh'} \\
 &= \frac{k(\mathbf{x} + h\mathbf{e}_i, \mathbf{y} + h'\mathbf{e}_i) - k(\mathbf{x} + h\mathbf{e}_i, \mathbf{y}) - k(\mathbf{x}, \mathbf{y} + h'\mathbf{e}_i) + k(\mathbf{x}, \mathbf{y})}{hh'} \\
 &\rightarrow \partial^{e_i, e_i} k(\mathbf{x}, \mathbf{y})
 \end{aligned} \tag{6}$$

as $h, h' \rightarrow 0$, for $\mathbf{x}, \mathbf{y} \in O$ close enough, since $\partial^{e_i, e_i} k$ is assumed to be continuous in a neighbourhood of the diagonal. Thus we have

$$\begin{aligned}
 & \mathbb{E} \left[\left(\frac{f(\mathbf{x} + h\mathbf{e}_i) - f(\mathbf{x})}{h} - \frac{f(\mathbf{x} + h'\mathbf{e}_i) - f(\mathbf{x})}{h'} \right)^2 \right] \\
 &= \mathbb{E} \left[\left(\frac{f(\mathbf{x} + h\mathbf{e}_i) - f(\mathbf{x})}{h} \right)^2 \right] - 2\mathbb{E} \left[\left(\frac{f(\mathbf{x} + h\mathbf{e}_i) - f(\mathbf{x})}{h} \right) \left(\frac{f(\mathbf{x} + h'\mathbf{e}_i) - f(\mathbf{x})}{h'} \right) \right] \\
 &\quad + \mathbb{E} \left[\left(\frac{f(\mathbf{x} + h'\mathbf{e}_i) - f(\mathbf{x})}{h'} \right)^2 \right] \\
 &\rightarrow \partial^{e_i, e_i} k(\mathbf{x}, \mathbf{x}) - 2\partial^{e_i, e_i} k(\mathbf{x}, \mathbf{x}) + \partial^{e_i, e_i} k(\mathbf{x}, \mathbf{x}) = 0
 \end{aligned} \tag{7}$$

where the limit is taken as $h, h' \rightarrow 0$. Equation (7) shows that $\lim_{h \rightarrow 0} \frac{f(\mathbf{x} + h\mathbf{e}_i) - f(\mathbf{x})}{h}$ exists in L^2 , i.e. that $\partial_{ms}^{e_i} f(\mathbf{x})$ exists. Moreover Equation (6) then implies that $\mathbb{E} [\partial_{ms}^{e_i} f(\mathbf{x}) \partial_{ms}^{e_i} f(\mathbf{y})] = \partial^{e_i, e_i} k(\mathbf{x}, \mathbf{y})$ for $\mathbf{x}, \mathbf{y} \in O$ close enough. Hence

$$\begin{aligned} \mathbb{E} [(\partial_{ms}^{e_i} f(\mathbf{x} + \mathbf{h}) - \partial_{ms}^{e_i} f(\mathbf{x}))^2] &= \mathbb{E} [\partial_{ms}^{e_i} f(\mathbf{x} + \mathbf{h})^2] - 2\mathbb{E} [\partial_{ms}^{e_i} f(\mathbf{x} + \mathbf{h}) \partial_{ms}^{e_i} f(\mathbf{x})] + \mathbb{E} [\partial_{ms}^{e_i} f(\mathbf{x})^2] \\ &= \partial^{e_i, e_i} k(\mathbf{x} + \mathbf{h}, \mathbf{x} + \mathbf{h}) - 2\partial^{e_i, e_i} k(\mathbf{x} + \mathbf{h}, \mathbf{x}) + \partial^{e_i, e_i} k(\mathbf{x}, \mathbf{x}) \\ &\rightarrow 0 \end{aligned}$$

as $\mathbf{h} \rightarrow \mathbf{0}$, so $\partial_{ms}^{e_i} f$ is mean-square continuous.

(3) \Rightarrow (1): For $\mathbf{x}, \mathbf{y} \in O$ and $1 \leq i \leq d$ we have

$$\frac{k(\mathbf{x} + h\mathbf{e}_i, \mathbf{y}) - k(\mathbf{x}, \mathbf{y})}{h} = \mathbb{E} \left[\frac{f(\mathbf{x} + h\mathbf{e}_i) - f(\mathbf{x})}{h} f(\mathbf{y}) \right] \rightarrow \mathbb{E} [\partial_{ms}^{e_i} f(\mathbf{x}) f(\mathbf{y})]$$

as $h \rightarrow 0$. Hence $\partial^{e_i, \mathbf{0}} k(\mathbf{x}, \mathbf{y})$ exists. Furthermore, for $1 \leq j \leq d$,

$$\begin{aligned} \frac{\partial^{e_i, \mathbf{0}} k(\mathbf{x}, \mathbf{y} + h\mathbf{e}_j) - \partial^{e_i, \mathbf{0}} k(\mathbf{x}, \mathbf{y})}{h} &= \mathbb{E} \left[\frac{\partial_{ms}^{e_i} f(\mathbf{x}) (f(\mathbf{y} + h\mathbf{e}_j) - f(\mathbf{y}))}{h} \right] \\ &\rightarrow \mathbb{E} [\partial_{ms}^{e_i} f(\mathbf{x}) \partial_{ms}^{e_j} f(\mathbf{y})] \end{aligned} \quad (8)$$

as $h \rightarrow 0$. Hence $\partial^{e_i, e_j} k$ exists, and is continuous by Equation (8), since $\partial_{ms}^{e_i} f$ and $\partial_{ms}^{e_j} f$ are mean-square continuous.

Now it remains to show that in all cases $\partial_{ms}^{e_i} f \sim \mathcal{GP}(0, \partial^{e_i, e_i} k)$. Let $\{\mathbf{x}_1, \dots, \mathbf{x}_N\} \subset O$ a finite set of points. Then we need to show

$$(\partial_{ms}^{e_i} f(\mathbf{x}_1), \dots, \partial_{ms}^{e_i} f(\mathbf{x}_N)) \sim \mathcal{N} \left(\mathbf{0}, (\partial^{e_i, e_i} k(\mathbf{x}_p, \mathbf{x}_q))_{p, q=1}^N \right). \quad (9)$$

$(\partial_{ms}^{e_i} f(\mathbf{x}_1), \dots, \partial_{ms}^{e_i} f(\mathbf{x}_N))$ is the limit in $L^2(\Omega; \mathbb{R}^N)$ as $h \rightarrow 0$ of

$$\begin{aligned} &\left(\frac{f(\mathbf{x}_p + h\mathbf{e}_i) - f(\mathbf{x}_p)}{h} \right)_{p=1}^N \\ &\sim \mathcal{N} \left(\mathbf{0}, \left(\frac{k(\mathbf{x}_p + h\mathbf{e}_i, \mathbf{x}_q + h\mathbf{e}_i) - k(\mathbf{x}_p + h\mathbf{e}_i, \mathbf{x}_q) - k(\mathbf{x}_p, \mathbf{x}_q + h\mathbf{e}_i) + k(\mathbf{x}_p, \mathbf{x}_q)}{h^2} \right)_{p, q=1}^N \right). \end{aligned} \quad (10)$$

Now convergence in L^2 implies convergence in distribution, and we see that the multivariate normal distribution in Equation (10) converges to the multivariate normal distribution in Equation (9) as $h \rightarrow 0$. So this shows $\partial_{ms}^{e_i} f \sim \mathcal{GP}(0, \partial^{e_i, e_i} k)$. \blacksquare

We can now prove our main result, Theorem 7, which we restate here for convenience:

Theorem 7 (Sample path Hölder regularity) *Let $n \in \mathbb{N}_0$ and $\gamma \in (0, 1]$. The process $f \sim \mathcal{GP}(0, k)$ has samples in $C_{loc}^{(n+\gamma)^-}(O)$ if and only if,*

(1) for general k ,

- $k \in C^{n \otimes n}(O \times O)$,

- $|\partial^{\alpha,\beta}k(\mathbf{x}+\mathbf{h},\mathbf{x}+\mathbf{h})-\partial^{\alpha,\beta}k(\mathbf{x}+\mathbf{h},\mathbf{x})-\partial^{\alpha,\beta}k(\mathbf{x},\mathbf{x}+\mathbf{h})+\partial^{\alpha,\beta}k(\mathbf{x},\mathbf{x})| = \mathcal{O}(\|\mathbf{h}\|^{2\epsilon})$
as $\mathbf{h} \rightarrow \mathbf{0}$, locally uniformly in $\mathbf{x} \in O$, for all $\epsilon \in (0, \gamma)$ and $|\alpha| = |\beta| = n$.

(2) for stationary $k(\mathbf{x}, \mathbf{y}) = k_\delta(\mathbf{x} - \mathbf{y})$,

- $k_\delta \in C^{2n}(\mathbb{R}^d)$,
- $|\partial^\alpha k_\delta(\mathbf{h}) - \partial^\alpha k_\delta(\mathbf{0})| = \mathcal{O}(\|\mathbf{h}\|^{2\epsilon})$ as $\mathbf{h} \rightarrow \mathbf{0}$ for all $\epsilon \in (0, \gamma)$ and $|\alpha| = 2n$.

(3) for isotropic $k(\mathbf{x}, \mathbf{y}) = k_r(\|\mathbf{x} - \mathbf{y}\|)$,

- $k_r \in C^{2n}(\mathbb{R})$,
- $|k_r^{(2n)}(h) - k_r^{(2n)}(0)| = \mathcal{O}(|h|^{2\epsilon})$ as $h \rightarrow 0$ for all $\epsilon \in (0, \gamma)$.

In each case, differentiating sample-wise we have $\partial^\alpha f \sim \mathcal{GP}(0, \partial^{\alpha,\alpha}k)$ for all $|\alpha| \leq n$.

Proof We only prove the general case (1). Cases (2) and (3) then follow by taking intersections over $\epsilon \in (0, \gamma)$ in Lemma A.1.

\Leftarrow : For $n = 0$ the result follows from applying the Kolmogorov continuity theorem to GPs, see Azmoodeh et al. (2014, Theorem 1), or Potthoff (2009, Page 347) for the result in arbitrary dimension d .

For $n = 1$, the existence and continuity of $\partial^{e_i, e_i} k$ implies by Lemma A.3 (1) \Rightarrow (3) that f is mean-square differentiable in direction i . Moreover this mean-square partial derivative satisfies $\partial_{m.s}^{e_i} f \sim \mathcal{GP}(0, \partial^{e_i, e_i} k)$. Now applying the case $n = 0$ to $\partial^{e_i, e_i} k$ we deduce that $\partial_{m.s}^{e_i} f$ has samples in² $C_{loc}^{\gamma^-}(O)$. In particular it has continuous samples, and this implies by (Potthoff, 2010, Theorem 3.2) that f has differentiable samples in direction i , with $\partial^{e_i} f(\mathbf{x}) = \partial_{m.s}^{e_i} f(\mathbf{x})$ almost surely, for all $\mathbf{x} \in O$. Thus $\partial^{e_i} f \sim \mathcal{GP}(0, \partial^{e_i, e_i} k)$ and $\partial^{e_i} f$ has samples in $C_{loc}^{\gamma^-}(O)$. $1 \leq i \leq d$ was arbitrary, so f has samples in $C_{loc}^{(1+\gamma)^-}(O)$.

For $n > 1$, we apply the same argument inductively on the partial derivatives.

\Rightarrow : For $n = 0$, Azmoodeh et al. (2014, Theorem 1) gives the converse to Kolmogorov's theorem for GPs. To be precise, this theorem is only stated for dimension $d = 1$, but we can see that the proof generalises to arbitrary d .

Now suppose $n = 1$. Pick $1 \leq i \leq d$. For any finite set of points $\{\mathbf{x}_1, \dots, \mathbf{x}_N\} \subset O$, $(\partial^{e_i} f(\mathbf{x}_1), \dots, \partial^{e_i} f(\mathbf{x}_N))$ is the almost sure limit of centered multivariate Gaussians with distribution

$$\mathcal{N}\left(\mathbf{0}, \left(\frac{k(\mathbf{x}_p + h\mathbf{e}_i, \mathbf{x}_q + h\mathbf{e}_i) - k(\mathbf{x}_p + h\mathbf{e}_i, \mathbf{x}_q) - k(\mathbf{x}_p, \mathbf{x}_q + h\mathbf{e}_i) + k(\mathbf{x}_p, \mathbf{x}_q)}{h^2}\right)_{p,q=1}^N\right)$$

(see Equation (10)), so is itself a centered multivariate Gaussian distribution. Hence $\partial^{e_i} f$ is a GP.

Let $\mathbf{x}, \mathbf{x} + h\mathbf{e}_i \in O$. By the mean value theorem, for each $\omega \in \Omega$ there is $\xi_{\mathbf{x},\omega}(h)$ between \mathbf{x} and $\mathbf{x} + h\mathbf{e}_i$ such that

$$\frac{f(\mathbf{x} + h\mathbf{e}_i, \omega) - f(\mathbf{x}, \omega)}{h} = \partial^{e_i} f(\xi_{\mathbf{x},\omega}(h), \omega).$$

2. The mean-square partial derivative $\partial_{m.s}^{e_i} f(\mathbf{x})$ at a point $\mathbf{x} \in O$ is only well-defined almost surely on Ω . But we still talk of sample paths for the process $\partial_{m.s}^{e_i} f$, since we mean this up to modification, as described in Section 2.

Thus

$$\left| \frac{f(\mathbf{x} + h\mathbf{e}_i, \omega) - f(\mathbf{x}, \omega)}{h} - \partial^{e_i} f(\mathbf{x}, \omega) \right|^2 = |\partial^{e_i} f(\xi_{\mathbf{x}, \omega}(h), \omega) - \partial^{e_i} f(\mathbf{x}, \omega)|^2 \leq C_\omega^2 |h|^\gamma$$

since $\gamma \in (0, 2\gamma)$, for some constant C_ω depending on ω but not on h , assuming h is small enough. Thus we have

$$\mathbb{E} \left[\left(\frac{f(\mathbf{x} + h\mathbf{e}_i) - f(\mathbf{x})}{h} - \partial^{e_i} f(\mathbf{x}) \right)^2 \right] \leq \underbrace{\mathbb{E} [C_\omega^2]}_{(*)} \cdot \underbrace{|h|^\gamma}_{(**)}$$

for h small enough. $(*) < \infty$ by Azmoodeh et al. (2014, Theorem 1), which we can apply since we showed that $\partial^{e_i} f$ is a GP. Also $(**) \rightarrow 0$ as $h \rightarrow 0$. Hence

$$\frac{f(\mathbf{x} + h\mathbf{e}_i) - f(\mathbf{x})}{h} \xrightarrow{L^2} \partial^{e_i} f(\mathbf{x})$$

as $h \rightarrow 0$, i.e. f is mean-square differentiable in direction i . Hence by Lemma A.3 (3) \Rightarrow (1), $\partial^{e_i, e_j} k$ exists and is continuous for all $1 \leq i, j \leq d$, and $\partial^{e_i} f \sim \mathcal{GP}(0, \partial^{e_i, e_i} k)$. Now

$$\begin{aligned} & \partial^{e_i, e_j} k(\mathbf{x} + \mathbf{h}, \mathbf{x} + \mathbf{h}) - \partial^{e_i, e_j} k(\mathbf{x} + \mathbf{h}, \mathbf{x}) - \partial^{e_i, e_j} k(\mathbf{x}, \mathbf{x} + \mathbf{h}) + \partial^{e_i, e_j} k(\mathbf{x}, \mathbf{x}) \\ &= \mathbb{E} [(\partial^{e_i} f(\mathbf{x} + \mathbf{h}) - \partial^{e_i} f(\mathbf{x}))(\partial^{e_j} f(\mathbf{x} + \mathbf{h}) - \partial^{e_j} f(\mathbf{x}))] \\ &\leq \mathbb{E} [(\partial^{e_i} f(\mathbf{x} + \mathbf{h}) - \partial^{e_i} f(\mathbf{x}))^2]^{1/2} \mathbb{E} [(\partial^{e_j} f(\mathbf{x} + \mathbf{h}) - \partial^{e_j} f(\mathbf{x}))^2]^{1/2} \\ &\leq \underbrace{\mathbb{E} [A_\omega^2]^{1/2}}_{< \infty} \|\mathbf{h}\|^\epsilon \cdot \underbrace{\mathbb{E} [B_\omega^2]^{1/2}}_{< \infty} \|\mathbf{h}\|^\epsilon = \mathcal{O}(\|\mathbf{h}\|^{2\epsilon}) \end{aligned} \quad (11)$$

as $\mathbf{h} \rightarrow \mathbf{0}$, locally uniformly in $\mathbf{x} \in O$, for all $\epsilon \in (0, \gamma)$, where A_ω, B_ω are some constants depending on ω , by Azmoodeh et al. (2014, Theorem 1) applied to $\partial^{e_i} f$ and $\partial^{e_j} f$, for $1 \leq i, j \leq d$.

Finally, for $n > 1$ we apply the same argument inductively on the partial derivatives. \blacksquare

Appendix B. Proofs in Section 4

B.1 Proof of Proposition 10

Proof For $\rho > 0$ we can write for $\nu \notin \mathbb{N}$

$$K_\nu(\rho) = \frac{\pi I_{-\nu}(\rho) - I_\nu(\rho)}{2 \sin(\nu\pi)}$$

where

$$I_{\pm\nu}(\rho) = \left(\frac{\rho}{2}\right)^{\pm\nu} \sum_{j=0}^{\infty} \frac{1}{j! \Gamma(j \pm \nu + 1)} \left(\frac{\rho}{2}\right)^{2j}$$

are modified Bessel functions of the first kind (Abramowitz and Stegun, 1965, Equations 9.6.2 & 9.6.10). So we can write the Matérn kernel (3) as

$$k_r(x) = C_\nu \left(\underbrace{\sum_{j=0}^{\infty} \frac{1}{j! \Gamma(j - \nu + 1)} \left(\frac{x}{2}\right)^{2j}}_{(*)} - \frac{|x|^{2\nu}}{2^\nu} \sum_{j=0}^{\infty} \frac{1}{j! \Gamma(j + \nu + 1)} \left(\frac{x}{2}\right)^{2j} \right)_{(**)}$$

for $x \in \mathbb{R}$, where $C_\nu > 0$ is some constant. The power series (*) and (**) are smooth in x , and $|x|^{2\nu}$ has precisely $\lceil 2\nu - 1 \rceil$ continuous derivatives with $(2\nu - 2\lceil \nu - 1 \rceil)$ -Hölder continuous $2\lceil \nu - 1 \rceil^{\text{th}}$ derivative at 0. So, for $\nu \notin \mathbb{N}$, the result follows by Theorem 7 (3).

For $\nu = n \in \mathbb{N}$, we instead use the following formula for the modified Bessel function of the second kind (Abramowitz and Stegun, 1965, Equations 9.6.11): for $\rho > 0$

$$\begin{aligned} K_n(\rho) &= \frac{1}{2} \left(\frac{\rho}{2}\right)^{-n} \sum_{j=0}^{n-1} \frac{(n-j-1)!}{j!} \left(-\frac{\rho^2}{4}\right)^j + (-1)^{n+1} \log\left(\frac{\rho}{2}\right) I_n(\rho) \\ &\quad + (-1)^n \frac{1}{2} \left(\frac{\rho}{2}\right)^n \sum_{j=0}^{\infty} (\psi(j+1) + \psi(n+j+1)) \frac{\left(\frac{\rho^2}{4}\right)^j}{j!(n+j)!} \end{aligned}$$

where ψ is the digamma function. So we can write the Matérn kernel (3) as

$$\begin{aligned} k_r(x) &= A_n \underbrace{\sum_{j=0}^{n-1} \frac{(n-j-1)!}{j!} (-1)^j \left(-\frac{x^2}{4}\right)^j}_{(I)} + B_n \underbrace{x^{2n} \log\left(\frac{x}{2}\right)}_{(II)} \underbrace{\sum_{j=0}^{\infty} \frac{1}{j!(n+j)!} \left(\frac{x}{2}\right)^{2j}}_{(III)} \\ &\quad + C_n \underbrace{x^{2n} \sum_{j=0}^{\infty} (\psi(j+1) + \psi(n+j+1)) \frac{\left(\frac{x}{2}\right)^{2j}}{j!(n+j)!}}_{(IV)} \end{aligned}$$

for some constants $A_n, B_n, C_n \neq 0$. (I) and (IV) are smooth in x . (III) is smooth in x and is non-zero at $x = 0$. (II) has precisely $2n - 1$ continuous derivatives, with almost-2-Hölder continuous $(2n - 2)^{\text{th}}$ derivative at 0 (precisely, a Hölder decay of $\mathcal{O}(h^2 \log h)$). So the result follows by Theorem 7 (3). \blacksquare

B.2 Proof of Proposition 12

Proof It is shown in (Wendland, 2004, Theorem 9.12) that the Wendland kernels may be written as

$$k_r(x) = \sum_{j=0}^{\lfloor d/2 \rfloor + 3n + 1} d_{j,n}^{(\lfloor d/2 \rfloor + n + 1)} |x|^j$$

for $x \in \mathbb{R}$, where $d_{j,n}^{(\lfloor d/2 \rfloor + n + 1)} \in \mathbb{R}$ are coefficients. Furthermore the odd degree coefficients satisfy $d_{2j+1,n}^{(\lfloor d/2 \rfloor + n + 1)} = 0$ if and only if $0 \leq j \leq 2n - 1$. Now for all $j \in \mathbb{N}_0$, $|x|^{2j}$ is smooth

and $|x|^{2j+1}$ is precisely $2n$ times continuously differentiable with Lipschitz $2n^{\text{th}}$ derivative. So the result follows by Theorem 7 (3). \blacksquare

B.3 Proof of Proposition 15

Proof This follows from Theorem 7 (1), by noting that

$$\partial^{\alpha,\beta}k(\mathbf{x}, \mathbf{y}) = \partial^\alpha \phi(\mathbf{x})^T \partial^\beta \phi(\mathbf{y})$$

for all $\mathbf{x}, \mathbf{y} \in O$ and $\alpha, \beta \in \mathbb{N}_0^d$ with $|\alpha|, |\beta| \leq n$, and moreover when $|\alpha| = |\beta| = n$,

$$\begin{aligned} & |\partial^{\alpha,\beta}k(\mathbf{x} + \mathbf{h}, \mathbf{x} + \mathbf{h}) - \partial^{\alpha,\beta}k(\mathbf{x} + \mathbf{h}, \mathbf{x}) - \partial^{\alpha,\beta}k(\mathbf{x}, \mathbf{x} + \mathbf{h}) + \partial^{\alpha,\beta}k(\mathbf{x}, \mathbf{x})| \\ &= (\partial^\alpha \phi(\mathbf{x} + \mathbf{h}) - \partial^\alpha \phi(\mathbf{x}))^T (\partial^\beta \phi(\mathbf{x} + \mathbf{h}) - \partial^\beta \phi(\mathbf{x})) \\ &\leq \left(\sum_{i=1}^m |\partial^\alpha \phi_i(\mathbf{x} + \mathbf{h}) - \partial^\alpha \phi_i(\mathbf{x})|^2 \right)^{1/2} \left(\sum_{i=1}^m |\partial^\beta \phi_i(\mathbf{x} + \mathbf{h}) - \partial^\beta \phi_i(\mathbf{x})|^2 \right)^{1/2} \\ &= \mathcal{O}(\|\mathbf{h}\|^\epsilon) \cdot \mathcal{O}(\|\mathbf{h}\|^\epsilon) = \mathcal{O}(\|\mathbf{h}\|^{2\epsilon}) \end{aligned}$$

as $\mathbf{h} \rightarrow \mathbf{0}$, locally uniformly in $\mathbf{x} \in O$, for all $\epsilon \in (0, \gamma)$, where we used the Cauchy-Schwarz inequality. \blacksquare

B.4 Proof of Proposition 18

Proof We assume $m = 2$; the proof for general m can then be done inductively. Note that by the product rule

$$\partial^{\alpha,\beta}k(\mathbf{x}, \mathbf{y}) = \sum_{\gamma_1 + \gamma_2 = \alpha} \sum_{\delta_1 + \delta_2 = \beta} \binom{|\alpha|}{|\gamma_1|} \binom{|\beta|}{|\delta_1|} \partial^{\gamma_1, \delta_1} k_1(\mathbf{x}, \mathbf{y}) \partial^{\gamma_2, \delta_2} k_2(\mathbf{x}, \mathbf{y})$$

for all $\mathbf{x}, \mathbf{y} \in O$ and $\alpha, \beta \in \mathbb{N}_0^d$ with $|\alpha|, |\beta| \leq n$.

Now for $|\alpha| = |\beta| = n$,

$$\begin{aligned} & \partial^{\alpha,\beta}k(\mathbf{x} + \mathbf{h}, \mathbf{x} + \mathbf{h}) - \partial^{\alpha,\beta}k(\mathbf{x} + \mathbf{h}, \mathbf{x}) - \partial^{\alpha,\beta}k(\mathbf{x}, \mathbf{x} + \mathbf{h}) + \partial^{\alpha,\beta}k(\mathbf{x}, \mathbf{x}) \\ &= \sum_{\gamma_1 + \gamma_2 = \alpha} \sum_{\delta_1 + \delta_2 = \beta} \binom{|\alpha|}{|\gamma_1|} \binom{|\beta|}{|\delta_1|} \left(\partial^{\gamma_1, \delta_1} k_1(\mathbf{x} + \mathbf{h}, \mathbf{x} + \mathbf{h}) \partial^{\gamma_2, \delta_2} k_2(\mathbf{x} + \mathbf{h}, \mathbf{x} + \mathbf{h}) \right. \\ &\quad - \partial^{\gamma_1, \delta_1} k_1(\mathbf{x} + \mathbf{h}, \mathbf{x}) \partial^{\gamma_2, \delta_2} k_2(\mathbf{x} + \mathbf{h}, \mathbf{x}) - \partial^{\gamma_1, \delta_1} k_1(\mathbf{x}, \mathbf{x} + \mathbf{h}) \partial^{\gamma_2, \delta_2} k_2(\mathbf{x}, \mathbf{x} + \mathbf{h}) \\ &\quad \left. + \partial^{\gamma_1, \delta_1} k_1(\mathbf{x}, \mathbf{x}) \partial^{\gamma_2, \delta_2} k_2(\mathbf{x}, \mathbf{x}) \right). \end{aligned} \tag{12}$$

To show the expression in Equation (12) is $\mathcal{O}(\|\mathbf{h}\|^{2\epsilon})$, it is therefore sufficient to show that

$$\begin{aligned} & |\partial^{\gamma_1, \delta_1} k_1(\mathbf{x} + \mathbf{h}, \mathbf{x} + \mathbf{h}) \partial^{\gamma_2, \delta_2} k_2(\mathbf{x} + \mathbf{h}, \mathbf{x} + \mathbf{h}) - \partial^{\gamma_1, \delta_1} k_1(\mathbf{x} + \mathbf{h}, \mathbf{x}) \partial^{\gamma_2, \delta_2} k_2(\mathbf{x} + \mathbf{h}, \mathbf{x}) \\ &\quad - \partial^{\gamma_1, \delta_1} k_1(\mathbf{x}, \mathbf{x} + \mathbf{h}) \partial^{\gamma_2, \delta_2} k_2(\mathbf{x}, \mathbf{x} + \mathbf{h}) + \partial^{\gamma_1, \delta_1} k_1(\mathbf{x}, \mathbf{x}) \partial^{\gamma_2, \delta_2} k_2(\mathbf{x}, \mathbf{x})| = \mathcal{O}(\|\mathbf{h}\|^{2\epsilon}) \end{aligned} \tag{13}$$

as $\mathbf{h} \rightarrow \mathbf{0}$, locally uniformly in $\mathbf{x} \in O$, for all $\epsilon \in (0, \gamma)$, $\gamma_1 + \gamma_2 = \alpha$ and $\delta_1 + \delta_2 = \beta$. The inside of the absolute value on the left hand side of Equation (13) can be written as

$$\begin{aligned}
 & \partial^{\gamma_1, \delta_1} k_1(\mathbf{x} + \mathbf{h}, \mathbf{x} + \mathbf{h}) \partial^{\gamma_2, \delta_2} k_2(\mathbf{x} + \mathbf{h}, \mathbf{x} + \mathbf{h}) - \partial^{\gamma_1, \delta_1} k_1(\mathbf{x} + \mathbf{h}, \mathbf{x}) \partial^{\gamma_2, \delta_2} k_2(\mathbf{x} + \mathbf{h}, \mathbf{x}) \\
 & \quad - \partial^{\gamma_1, \delta_1} k_1(\mathbf{x}, \mathbf{x} + \mathbf{h}) \partial^{\gamma_2, \delta_2} k_2(\mathbf{x}, \mathbf{x} + \mathbf{h}) + \partial^{\gamma_1, \delta_1} k_1(\mathbf{x}, \mathbf{x}) \partial^{\gamma_2, \delta_2} k_2(\mathbf{x}, \mathbf{x}) \\
 & = \partial^{\gamma_2, \delta_2} k_2(\mathbf{x} + \mathbf{h}, \mathbf{x} + \mathbf{h}) \\
 & \quad \cdot \underbrace{(\partial^{\gamma_1, \delta_1} k_1(\mathbf{x} + \mathbf{h}, \mathbf{x} + \mathbf{h}) - \partial^{\gamma_1, \delta_1} k_1(\mathbf{x} + \mathbf{h}, \mathbf{x}) - \partial^{\gamma_1, \delta_1} k_1(\mathbf{x}, \mathbf{x} + \mathbf{h}) + \partial^{\gamma_1, \delta_1} k_1(\mathbf{x}, \mathbf{x}))}_{(*)} \\
 & \quad + \partial^{\gamma_1, \delta_1} k_1(\mathbf{x}, \mathbf{x}) \\
 & \quad \cdot \underbrace{(\partial^{\gamma_2, \delta_2} k_2(\mathbf{x} + \mathbf{h}, \mathbf{x} + \mathbf{h}) - \partial^{\gamma_2, \delta_2} k_2(\mathbf{x} + \mathbf{h}, \mathbf{x}) - \partial^{\gamma_2, \delta_2} k_2(\mathbf{x}, \mathbf{x} + \mathbf{h}) + \partial^{\gamma_2, \delta_2} k_2(\mathbf{x}, \mathbf{x}))}_{(**)} \\
 & \quad + \underbrace{(\partial^{\gamma_1, \delta_1} k_1(\mathbf{x} + \mathbf{h}, \mathbf{x}) - \partial^{\gamma_1, \delta_1} k_1(\mathbf{x}, \mathbf{x}))}_{(I)} \underbrace{(\partial^{\gamma_2, \delta_2} k_2(\mathbf{x} + \mathbf{h}, \mathbf{x} + \mathbf{h}) - \partial^{\gamma_2, \delta_2} k_2(\mathbf{x} + \mathbf{h}, \mathbf{x}))}_{(II)} \\
 & \quad + \underbrace{(\partial^{\gamma_1, \delta_1} k_1(\mathbf{x}, \mathbf{x} + \mathbf{h}) - \partial^{\gamma_1, \delta_1} k_1(\mathbf{x}, \mathbf{x}))}_{(III)} \underbrace{(\partial^{\gamma_2, \delta_2} k_2(\mathbf{x} + \mathbf{h}, \mathbf{x} + \mathbf{h}) - \partial^{\gamma_2, \delta_2} k_2(\mathbf{x}, \mathbf{x} + \mathbf{h}))}_{(IV)}
 \end{aligned} \tag{14}$$

By arguing as in the proof of Theorem 7, Equation (11) we have $(*), (**) = \mathcal{O}(\|\mathbf{h}\|^{2\epsilon})$. Moreover, by the Cauchy-Schwarz inequality,

$$\begin{aligned}
 |(I)| & = \left| \mathbb{E}[\partial^{\gamma_1} f_1(\mathbf{x} + \mathbf{h}) \partial^{\delta_1} f_1(\mathbf{x}) - \partial^{\gamma_1} f_1(\mathbf{x}) \partial^{\delta_1} f_1(\mathbf{x} + \mathbf{h})] \right| \\
 & \leq \mathbb{E}[\partial^{\delta_1} f_1(\mathbf{x})^2]^{1/2} \mathbb{E}[(\partial^{\gamma_1} f_1(\mathbf{x} + \mathbf{h}) - \partial^{\gamma_1} f_1(\mathbf{x}))^2]^{1/2} \\
 & = \partial^{\delta_1, \delta_1} k_1(\mathbf{x}, \mathbf{x})^{1/2} (\partial^{\gamma_1, \gamma_1} k_1(\mathbf{x} + \mathbf{h}, \mathbf{x} + \mathbf{h}) - \partial^{\gamma_1, \gamma_1} k_1(\mathbf{x} + \mathbf{h}, \mathbf{x}) \\
 & \quad - \partial^{\gamma_1, \gamma_1} k_1(\mathbf{x}, \mathbf{x} + \mathbf{h}) + \partial^{\gamma_1, \gamma_1} k_1(\mathbf{x}, \mathbf{x}))^{1/2} \\
 & = \mathcal{O}(\|\mathbf{h}\|^\epsilon)
 \end{aligned} \tag{15}$$

where $f_1 \sim \mathcal{GP}(0, k_1)$. Similarly $(II), (III), (IV) = \mathcal{O}(\|\mathbf{h}\|^\epsilon)$. By Equation (14) we therefore deduce Equation (13), which concludes the proof. \blacksquare

References

- Milton Abramowitz and Irene A. Stegun. *Handbook of Mathematical Functions: With Formulas, Graphs, and Mathematical Tables*. Courier Corporation, 1965. ISBN 978-0-486-61272-0.
- Robert Adler and Jonathan Taylor. *Random Fields and Geometry*. Springer Monographs in Mathematics. Springer, 2007. ISBN 978-0-387-48116-6. doi:10.1007/978-0-387-48116-6.
- Ehsan Azmoodeh, Tommi Sottinen, Lauri Viitasaari, and Adil Yazigi. Necessary and sufficient conditions for Hölder continuity of Gaussian processes. *Statistics & Probability Letters*, 94(C):230–235, 2014. doi:10.1016/j.spl.2014.07.030.

- Christian Berg, Jens Peter Reus Christensen, and Paul Ressel. *Harmonic Analysis on Semigroups*, volume 100 of *Graduate Texts in Mathematics*. Springer, 1984. ISBN 978-1-4612-1128-0. doi:10.1007/978-1-4612-1128-0.
- Viacheslav Borovitskiy, Alexander Terenin, Peter Mostowsky, and Marc Deisenroth. Matérn Gaussian processes on Riemannian manifolds. In *Advances in Neural Information Processing Systems*, volume 33, pages 12426–12437. Curran Associates, Inc., 2020.
- N. N. Chentsov. Weak convergence of stochastic processes whose trajectories have no discontinuities of the second kind and the “heuristic” approach to the Kolmogorov-Smirnov tests. *Theory of Probability & Its Applications*, 1(1):140–144, 1956. ISSN 0040-585X. doi:10.1137/1101013.
- Nathaël Da Costa, Cyrus Mostajeran, Juan-Pablo Ortega, and Salem Said. Invariant kernels on Riemannian symmetric spaces: a harmonic-analytic approach. 2023. doi:10.48550/arXiv.2310.19270.
- Xavier Fernique. Régularité des trajectoires des fonctions aleatoires gaussiennes. In *Ecole d’Eté de Probabilités de Saint-Flour IV—1974*, Lecture Notes in Mathematics, pages 1–96. Springer, 1975. ISBN 978-3-540-37600-2. doi:10.1007/BFb0080190.
- Karl Friston, Lancelot Da Costa, Noor Sajid, Conor Heins, Kai Ueltzhöffer, Grigorios A. Pavliotis, and Thomas Parr. The free energy principle made simpler but not too simple. *Physics Reports*, 1024:1–29, 2023. ISSN 0370-1573. doi:10.1016/j.physrep.2023.07.001.
- Tilmann Gneiting. On the derivatives of radial positive definite functions. *Journal of Mathematical Analysis and Applications*, 236(1):86–93, 1999. ISSN 0022-247X. doi:10.1006/jmaa.1999.6434.
- Conor Heins, Beren Millidge, Lancelot Da Costa, Richard Mann, Karl Friston, and Iain Couzin. Collective behavior from surprise minimization. 2023. doi:10.48550/arXiv.2307.14804.
- Iain Henderson. Sobolev regularity of Gaussian random fields. *Journal of Functional Analysis*, 286(3):110241, 2024. ISSN 0022-1236. doi:10.1016/j.jfa.2023.110241.
- Motonobu Kanagawa, Philipp Hennig, Dino Sejdinovic, and Bharath K. Sriperumbudur. Gaussian processes and kernel methods: a review on connections and equivalences. 2018. doi:10.48550/arXiv.1807.02582.
- Achim Klenke. *Probability Theory: A Comprehensive Course*. Universitext. Springer, 2014. ISBN 978-1-4471-5361-0. doi:10.1007/978-1-4471-5361-0.
- Marvin Pförtner, Ingo Steinwart, Philipp Hennig, and Jonathan Wenger. Physics-informed Gaussian process regression generalizes linear PDE solvers. 2022. doi:10.48550/arXiv.2212.12474.
- Albert Pierce. *Fundamentals of Nonparametric Statistics*. Dickenson Publishing Company, 1970. ISBN 978-0822143109.

- Jürgen Potthoff. Sample properties of random fields. II. Continuity. *Communications on Stochastic Analysis*, 3(3), 2009. ISSN 0973-9599. doi:10.31390/cosa.3.3.02.
- Jürgen Potthoff. Sample properties of random fields III: differentiability. *Communications on Stochastic Analysis*, 4(3), 2010. ISSN 2688-6669. doi:10.31390/cosa.4.3.03.
- Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2005. ISBN 978-0-262-25683-4. doi:10.7551/mitpress/3206.001.0001.
- Michael Scheuerer. *A Comparison of Models and Methods for Spatial Interpolation in Statistics and Numerical Analysis*. PhD thesis, 2010a.
- Michael Scheuerer. Regularity of the sample paths of a general second order random field. *Stochastic Processes and their Applications*, 120(10):1879–1897, 2010b. ISSN 0304-4149. doi:10.1016/j.spa.2010.05.009.
- Michael Schober, David K Duvenaud, and Philipp Hennig. Probabilistic ODE solvers with Runge-Kutta means. In *Advances in Neural Information Processing Systems*, volume 27, 2014.
- Michel Talagrand. Regularity of gaussian processes. *Acta Mathematica*, 159:99–149, 1987. ISSN 0001-5962, 1871-2509. doi:10.1007/BF02392556.
- Junyang Wang, Jon Cockayne, Oksana Chkrebtii, T. J. Sullivan, and Chris. J. Oates. Bayesian numerical methods for nonlinear partial differential equations. *Statistics and Computing*, 31(5):55, 2021. ISSN 1573-1375.
- Holger Wendland. *Scattered Data Approximation*. Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press, Cambridge, 2004. ISBN 978-0-521-84335-5. doi:10.1017/CBO9780511617539.