



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Convergence of Gaussian Process Regression with Estimated Hyper-parameters and Applications in Bayesian Inverse Problems

Citation for published version:

Teckentrup, AL 2020, 'Convergence of Gaussian Process Regression with Estimated Hyper-parameters and Applications in Bayesian Inverse Problems', *SIAM/ASA Journal on Uncertainty Quantification*, vol. 8, no. 4, pp. 1310–1337. <https://doi.org/10.1137/19M1284816>

Digital Object Identifier (DOI):

[10.1137/19M1284816](https://doi.org/10.1137/19M1284816)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

SIAM/ASA Journal on Uncertainty Quantification

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Convergence of Gaussian Process Regression with Estimated Hyper-parameters and Applications in Bayesian Inverse Problems

Aretha L. Teckentrup¹

¹ School of Mathematics, University of Edinburgh, James Clerk Maxwell Building, Edinburgh, EH9 3FD, UK. a.teckentrup@ed.ac.uk

Abstract

This work is concerned with the convergence of Gaussian process regression. A particular focus is on hierarchical Gaussian process regression, where hyper-parameters appearing in the mean and covariance structure of the Gaussian process emulator are a-priori unknown, and are learnt from the data, along with the posterior mean and covariance. We work in the framework of empirical Bayes, where a point estimate of the hyper-parameters is computed, using the data, and then used within the standard Gaussian process prior to posterior update. We provide a convergence analysis that (i) holds for a given, deterministic function f to be emulated; and (ii) shows that convergence of Gaussian process regression is unaffected by the additional learning of hyper-parameters from data, and is guaranteed in a wide range of scenarios. As the primary motivation for the work is the use of Gaussian process regression to approximate the data likelihood in Bayesian inverse problems, we provide a bound on the error introduced in the Bayesian posterior distribution in this context.

Keywords: inverse problem, Bayesian inference, surrogate model, Gaussian process regression, posterior consistency, hierarchical, empirical Bayes'

AMS 2020 subject classifications: 62G08, 62J07, 65D15, 65D40, 65J22

1 Introduction

Mathematical modelling and simulation are indispensable tools frequently used to inform decisions and assess risk. In practice, the parameters appearing in the models are often unknown, and have to be inferred from indirect observations. This leads to an *inverse problem*, where one infers the parameters of the model given incomplete, noisy observations of the model outputs. Adopting a Bayesian approach [19, 50], we incorporate our prior knowledge of the parameters into a probability distribution, referred to as the *prior distribution*, and obtain a more accurate representation of the parameters in the *posterior distribution*, which results from conditioning the prior distribution on the observations.

The goal of simulations is typically to (i) sample from the posterior distribution, using methods such as Markov chain Monte Carlo (MCMC), and/or (ii) compute a point estimate of the parameters, such as the most likely value under the posterior distribution (known as the maximum a-posteriori (MAP) estimate). Both of these tasks quickly become computationally infeasible when the mathematical model involved is complex. In many applications, for example when the forward

model is given by a partial differential equation, computing one instance of the forward model is computationally very expensive, and the sheer number of model evaluations required for the sampling and/or optimisation is prohibitively large.

This drawback of fully Bayesian inference for complex models was recognised several decades ago in the statistics literature, and resulted in key papers which had a profound influence on methodology [41, 21, 35]. These papers advocated the use of a Gaussian process surrogate model (also called emulator) to approximate the solution of the governing equations, and in particular the data likelihood, at a much lower computational cost.

The focus of this work is on the convergence analysis of Gaussian process surrogate models, in the case where the hyper-parameters in the distribution of the Gaussian process are a-priori unknown and inferred as part of the construction of the surrogate. This situation is of significant importance and interest, for, amongst others, the following reasons. Firstly, by correctly tuning the hyper-parameters, we will obtain a Gaussian process surrogate model that mimics closely the behaviour of the function we are approximating, resulting in a smaller error in the approximation. Secondly, the variance of the Gaussian process surrogate model is often used to represent the error in the approximation. However, for this interpretation to make sense, the hyper-parameters have to be chosen correctly. For example, the variance of the Gaussian process surrogate model can artificially be driven to zero by letting the marginal variance of the covariance kernel go to zero, but the error does not vanish in reality.

We adopt an empirical Bayes approach, also known as a plug-in approach, where we compute an estimate of the hyper-parameters, and plug this into the predictive equations for a Gaussian process surrogate model with known hyper-parameters. We present a convergence analysis of these hierarchical Gaussian process surrogate models, which shows that convergence of the mean and variance of the Gaussian process emulator is guaranteed under very mild assumptions on the estimated hyper-parameters. In particular, the convergence rates of the hierarchical Gaussian process emulator are the same as the convergence rates obtained for Gaussian process emulators with fixed, known values of the hyper-parameters, if the estimated hyper-parameters converge to the known values.

As particular examples of covariance kernels used to construct the emulators, we consider Matérn kernels and separable (or multiplicative/tensor-product) Matérn kernels. As we will see in section 3, the type of covariance kernel one should employ depends on the structure and smoothness of the function being emulated. The use of Matérn kernels corresponds to assuming a certain Sobolev smoothness, whereas the use of separable Matérn kernels assumes a tensor-product Sobolev structure (also known as mixed dominating smoothness).

The question of how the estimation of hyper-parameters influences the error in Gaussian process emulators is not new, and has been dealt with in the spatial statistics literature [47, 48, 37, 13, 44, 52, 8]. However, these results are of a different nature to our new results presented in section 3, and to the type of error bounds needed in section 5 to justify the use of Gaussian process emulators in Bayesian inverse problems (see also [49]). In particular, our results (i) give bounds for a fixed, deterministic function being emulated, rather than averaging over a certain distribution of functions, and (ii) do not require the hyper-parameters to be identifiable or the estimated hyper-parameters to converge.

A further distinction to previous studies, is that we do not require a notion of "true" values of the hyper-parameters. The customary (and often necessary) definition in spatial statistics (cf [47, 48, 37, 13]) is to choose the true parameter values such that the function being emulated

is a sample of the corresponding Gaussian process. In our analysis, we do not require any such assumption on the function being emulated. True parameter values in our context would simply represent a good choice of hyper-parameters, and can be defined in any way that the user finds suitable (including the customary definition above). Likewise, the estimated hyper-parameters can be defined in many suitable ways, e.g through maximum likelihood or maximum a-posteriori estimation (cf [17]) or cross-validation (cf [53]). Our results are independent of how the hyper-parameters are estimated.

1.1 Our Contributions

In this paper, we make the following contributions to the analysis of Gaussian process regression:

1. We provide a convergence analysis of Gaussian process regression with estimated hyper-parameters, which shows convergence of the emulators to the true function as the number of design points tends to infinity.
2. We justify the use of hierarchical Gaussian process emulators to approximate the data likelihood in Bayesian inverse problems, by bounding the error introduced in the posterior distribution. Previous results, well known in the spatial statistics literature, are not sufficient for this purpose.

1.2 Paper Structure

The paper is organised as follows. Section 2 introduces hierarchical Gaussian process regression, and summarises relevant results from the spatial statistics literature. Section 3 analyses the error in hierarchical Gaussian process regression in a wide range of scenarios. We set up the Bayesian inverse problem of interest in section 4, whereas Section 5 then considers the use of hierarchical Gaussian process emulators to approximate the posterior distribution in the Bayesian inverse problem. Section 6 provides a summary and discussion of the main results.

2 Hierarchical Gaussian Process Regression

We want to use Gaussian process regression (also known as Gaussian process *emulation* or *kriging*) to derive a computationally cheaper approximation to a given function $f : U \rightarrow \mathbb{R}$, where $U \subseteq \mathbb{R}^{d_u}$ is compact with Lipschitz boundary. We focus on the case where the hyper-parameters defining the Gaussian process emulator are unknown a-priori, and are inferred as part of the construction of the emulator. We denote these hyper-parameters by θ , and treat them using an empirical Bayes approach.

2.1 Set-up

Let $f : U \rightarrow \mathbb{R}$ be an arbitrary function. To derive the Gaussian process emulator of f , we use a Bayesian procedure and assign a Gaussian process prior distribution to f :

$$f_0 | \theta \sim \text{GP}(m(\theta; \cdot), k(\theta; \cdot, \cdot)). \quad (2.1)$$

To avoid confusion between the true function f and its prior distribution, we have added the subscript zero in the above prior. Here, $\theta \in R_\theta \subseteq \mathbb{R}^{d_\theta}$ are now hyper-parameters defining the mean

function $m(\theta; \cdot) : U \rightarrow \mathbb{R}$ and the two-point covariance function $k(\theta; \cdot, \cdot) : U \times U \rightarrow \mathbb{R}$, assumed to be positive-definite for all $\theta \in S$, for any compact subset $S \subseteq R_\theta$. Particular examples of covariance kernels $k(\theta)$ are the Matérn and separable Matérn families discussed in sections 2.2 and 2.3. For the mean function $m(\theta)$, we can for example use polynomials, in which case the hyper-parameters are typically the unknown polynomial coefficients. We will write $\theta = \{\theta_{\text{mean}}, \theta_{\text{cov}}\}$ when we want to explicitly distinguish between the hyper-parameters appearing in the mean and covariance function, respectively.

We further put a prior distribution $\mathbb{P}(\theta)$ on θ , with Lebesgue density $p(\theta)$. The joint prior distribution on (f, θ) is then given by

$$\mathbb{P}(f_0, \theta) = \mathbb{P}(f_0 | \theta) \mathbb{P}(\theta).$$

Then, given data in the form of a set of distinct design points $D_N := \{u^n\}_{n=1}^N \subseteq U$, together with corresponding function values

$$f(D_N) := [f(u^1), \dots, f(u^N)] \in \mathbb{R}^N,$$

we condition the prior distribution $\mathbb{P}(f_0, \theta)$ on the observed data $f(D_N)$ to obtain the posterior distribution

$$\mathbb{P}(f_0, \theta | f(D_N)) = \mathbb{P}(f_0 | \theta, f(D_N)) \mathbb{P}(\theta | f(D_N)).$$

The distribution $\mathbb{P}(f_0 | \theta, f(D_N))$ is again a Gaussian process, with explicitly known mean function $m_N^f(\theta; \cdot)$ and covariance kernel $k_N(\theta; \cdot, \cdot)$:

$$m_N^f(\theta; u) = m(\theta; u) + k(\theta; u, D_N)^T K(\theta; D_N)^{-1} (f(D_N) - m(\theta; D_N)), \quad (2.2)$$

$$k_N(\theta; u, u') = k(\theta; u, u') - k(\theta; u, D_N)^T K(\theta; D_N)^{-1} k(\theta; D_N, u'), \quad (2.3)$$

where $k(\theta; u, D_N) = [k(\theta; u, u^1), \dots, k(\theta; u, u^N)] \in \mathbb{R}^N$, $K(\theta; D_N) \in \mathbb{R}^{N \times N}$ is the matrix with ij^{th} entry equal to $k(\theta; u^i, u^j)$ and $m(\theta; D_N) := [m(\theta; u^1), \dots, m(\theta; u^N)] \in \mathbb{R}^N$. These are the well-known formulae for Gaussian process emulation [38], here adopting notation that will enable us to make use of the analysis of such emulators in [49]. When we wish to make explicit the dependence on the prior mean m , we will denote the predictive mean in (2.2) by $m_N^{f,m}(\theta)$.

The marginal distribution

$$\mathbb{P}(f_0 | f(D_N)) = \int_{\theta} \mathbb{P}(f_0, \theta | f(D_N)) \, d\theta = \int_{\theta} \mathbb{P}(f_0 | \theta, f(D_N)) \mathbb{P}(\theta | f(D_N)) \, d\theta$$

is typically not available in closed form, since the integrals involved are intractable. In practice one therefore often uses a *plug-in* approach, also known as *empirical Bayes*. This consists of calculating an estimate $\hat{\theta}_N$ of θ using the data $f(D_N)$, and then approximating

$$\mathbb{P}(f_0 | f(D_N)) \approx \mathbb{P}(f_0 | \hat{\theta}_N, f(D_N)) = \text{GP}(m_N^f(\hat{\theta}_N; \cdot), k_N(\hat{\theta}_N; \cdot, \cdot)).$$

This corresponds to approximating the distribution $\mathbb{P}(\theta | f(D_N))$ by a Dirac measure at $\theta = \hat{\theta}_N$. For the remainder of this work, we will use

$$f_N(\hat{\theta}_N) \sim \text{GP}(m_N^f(\hat{\theta}_N; \cdot), k_N(\hat{\theta}_N; \cdot, \cdot)) \quad (2.4)$$

as a Gaussian process emulator of f . The process in (2.4) is also referred to as the *predictive process*, and we shall refer to $m_N^f(\hat{\theta}_N; \cdot)$ and $k_N(\hat{\theta}_N; \cdot, \cdot)$ as the predictive mean and the predictive covariance, respectively.

In this work, we will focus on the convergence of the emulator $f_N(\hat{\theta}_N)$ to the true function f , and how this is affected by the learning of the hyper-parameters θ . Computing a good estimate $\hat{\theta}_N$ of the hyper-parameters θ from the data $f(D_N)$ is an important and difficult question in practice. However, our results are independent of how this estimate is computed, and our results are also independent of whether the hyper-parameters are *identifiable*.

Following [57], the random field model $g \sim \text{GP}(m(\theta; \cdot), k(\theta; \cdot, \cdot))$ is identifiable if it is theoretically possible to learn the true value of θ after obtaining an infinite number of observations of g on U . In other words, the model $g \sim \text{GP}(m(\theta; \cdot), k(\theta; \cdot, \cdot))$ is identifiable if different values of θ give rise to orthogonal Gaussian measures.

By the Cameron-Martin Theorem ([7], see also [36, Proposition 2.24]) it follows in particular that models with polynomial mean functions $m(\theta_{\text{mean}})$, where the parameters θ_{mean} represent the coefficients or the degree of the polynomial, are in most cases not identifiable, since polynomials are typically contained in the reproducing kernel Hilbert space (a.k.a. Cameron Martin space) associated to k . In particular, this is the case for the Matérn and separable Matérn kernels presented below.

2.2 Matérn Covariance Kernels

Covariance functions $k(\theta; \cdot, \cdot)$ frequently used in applications are the Matérn covariance functions

$$k_{\text{Mat}}(\theta; u, u') = \frac{\sigma^2}{\Gamma(\nu)2^{\nu-1}} \left(\frac{\|u - u'\|_2}{\lambda} \right)^\nu B_\nu \left(\frac{\|u - u'\|_2}{\lambda} \right), \quad (2.5)$$

with hyper-parameters $\theta_{\text{cov}} = \{\sigma^2, \lambda, \nu\} \in (0, \infty)^3$. Here, Γ denotes the Gamma function, and B_ν denotes the modified Bessel function of the second kind [27]. The parameter σ^2 is usually referred to as the (marginal) variance, λ as the correlation length and ν as the smoothness parameter. The expression for the Matérn covariance kernel simplifies for particular choices of ν . Notable examples include the exponential covariance kernel $\sigma^2 \exp(-\|u - u'\|_2/\lambda)$ with $\nu = 1/2$, and the Gaussian covariance kernel $\sigma^2 \exp(-\|u - u'\|_2^2/\lambda^2)$ in the limit $\nu \rightarrow \infty$.

The identifiability of the Matérn model has been studied in [58, 57, 3]. While all parameters θ_{cov} are identifiable for $d_u \geq 5$, only the quantities ν and $\sigma^2 \lambda^{-2\nu}$ are identifiable when $d_u \leq 3$. The case $d_u = 4$ remains open. To alleviate problems with identifiability for $d_u \leq 3$, the recent paper [17] discusses choices for the prior distribution on the hyper-parameters $\{\sigma^2, \lambda\}$, such that the MAP estimate $\hat{\theta}_N^{\text{MAP}}$ gives a good estimate of the true value of θ .

We would briefly like to point out here that the identifiability issues mentioned above are related to the fact that our parameter space U is bounded, which means that we are dealing with *in-fill* asymptotics. If U were unbounded, we would be dealing with *increasing domain* asymptotics, where all parameters are identifiable also when $d_u \leq 3$ [58].

2.3 Separable Matérn Covariance Kernels

As an alternative to the classical Matérn covariance functions in the previous section, one can consider using their separable versions (also called multiplicative or tensor-product versions). These

are obtained by taking the product of one-dimensional Matérn covariance functions:

$$k_{\text{sepMat}}(\theta; u, u') = \prod_{j=1}^{d_u} k_{\text{Mat}}(\theta_j; u_j, u'_j). \quad (2.6)$$

Since the marginal variances σ_j^2 only enter as multiplicative pre-factors, the hyper-parameters in this case are $\{\nu_j, \lambda_j\}_{j=1}^{d_u}$ and $\sigma^2 := \prod_{j=1}^{d_u} \sigma_j^2$, leading to $\theta_{\text{cov}} \in (0, \infty)^{2d_u+1}$. A particular example is the separable exponential covariance kernel, which corresponds to $\nu_j \equiv 1/2$ and hence takes the form

$$k_{\text{sepExp}}(\theta; u, u') = \sigma^2 \exp \left(- \sum_{j=1}^{d_u} \frac{|u_j - u'_j|}{\lambda_j} \right). \quad (2.7)$$

The separable versions of Matérn kernels can have better properties than the classical Matérn kernels in terms of identifiability. For example, [56, Theorem 1] shows that, provided $d_u > 1$, the model $g \sim \text{GP}(0, k_{\text{sepExp}}(\theta; \cdot, \cdot))$ on U is identifiable. The case of general separable Matérn covariance kernels appears to be open, but related results in this direction can be found in [25, 26, 12]. Note that for $d_u = 1$, the classical and separable Matérn kernels coincide. In particular, for the stationary Ornstein-Uhlenbeck process given by $d_u = 1$ and $\nu = 1/2$, only the quantity $\sigma^2 \lambda^{-1}$ is identifiable on the bounded domain U .

3 Error Analysis of Hierarchical Gaussian Process Regression

In this section, we are concerned with the convergence of the hierarchical Gaussian process emulator f_N to the function f . Although the main idea behind the error estimates in this section is related to those in [47, 48, 37], we are here interested in error bounds which (i) do not assume that the function being emulated is a sample of a particular Gaussian process, (ii) bound the error for a given, deterministic function f , and (iii) are flexible with respect to the definition of the estimated hyper-parameters, so do not require any assumptions on identifiability of the hyper-parameters. Furthermore, the error analysis here will be performed in norms amenable to the use of the hierarchical Gaussian process emulators as surrogate models in Bayesian inverse problems, see section 5 for details. For a short discussion of the prediction error typically studied in the spatial statistics literature, see section 3.3.

Since the error analysis depends on various properties of the covariance kernel, such as the corresponding reproducing kernel Hilbert space (also known as the native space or Cameron-Martin space), we will consider two particular examples, namely the classical and the separable Matérn covariance kernels already considered in sections 2.2 and 2.3.

The definition of the estimated parameter values $\hat{\theta}_N$ is open, and our analysis does not require any assumptions on how these estimates are computed. We do not require that the sequence $\{\hat{\theta}_N\}_{N \in \mathbb{N}}$ converges, neither do we require the parameters θ to be identifiable. We could for example use maximum likelihood or maximum a-posteriori estimators, choose $\hat{\theta}_N$ to minimise the error $\|f - m_N^f(\hat{\theta}_N)\|_{L^2(U)}$, or use a combination of different approaches for different hyper-parameters. Note, however, that we do not want to minimise the predictive variance $\|k_N^{1/2}(\hat{\theta}_N)\|_{L^2(U)}^1$, since

¹By slight abuse of notation, we denote by $k_N(\hat{\theta}_N)$ the function of one variable that gives the predictive variance at a point u , $k_N(\hat{\theta}_N; u, u)$. The quantity $\|k_N^{1/2}(\hat{\theta}_N)\|_{L^2(U)} = \sqrt{\int_U k_N(\hat{\theta}_N; u, u) du}$ is hence an averaged predictive variance.

this can be made arbitrarily small by letting $\hat{\sigma}_N^2 \rightarrow 0$. We want to choose $\hat{\sigma}_N^2$ such that $k_N^{1/2}(\hat{\theta}_N)$ is a good representation of our remaining uncertainty about the function f , after observing the $f(D_N)$.

We would like to quantify the performance of the predictive mean and covariance functions $m_N^f(\hat{\theta}_N)$ and $k_N(\hat{\theta}_N)$. In particular, in light of the error bounds required for the Bayesian posterior distribution in section 5, we are interested in the quantities $\|f - m_N^f(\hat{\theta}_N)\|_{L^2(U)}$ and $\|k_N(\hat{\theta}_N)\|_{L^2(U)}$. We recall the following fundamental results, which hold for any kernel $k(\theta)$.

Proposition 3.1. [38, Section 6.2], [45, Theorem 1] *The function $m_N^{f,0}(\theta)$, given by (2.2) with $m(\theta) \equiv 0$, is the minimal norm interpolant of f in the native space corresponding to $k(\theta)$, with interpolation points D_N :*

$$m_N^{f,0}(\theta) = \arg \min_{g \in H_{k(\theta)}(U) : g(D_N) = f(D_N)} \|g\|_{H_{k(\theta)}(U)}.$$

In particular, $\|m_N^{f,0}(\theta)\|_{H_{k(\theta)}(U)} \leq \|f\|_{H_{k(\theta)}(U)}$.

Proposition 3.2. [49, Proposition 3.5] *Suppose $k_N(\theta)$ is given by (2.3). Then*

$$k_N(\theta; u, u)^{\frac{1}{2}} = \sup_{\|g\|_{H_{k(\theta)}(U)}=1} |g(u) - m_N^{g,0}(\theta; u)|,$$

where $m_N^{g,0}(\theta)$ is given by (2.2) with $m(\theta) \equiv 0$.

3.1 Matérn Covariance Kernels

Suppose we use a Matérn covariance kernel $k_{\text{Mat}}(\hat{\theta}_N)$, defined in (2.5), to construct the hierarchical Gaussian process emulator f_N , defined in (2.4).

Given the set of design points $D_N = \{u^n\}_{n=1}^N \subseteq U$, we define the fill distance $h_{D_N,U}$, separation radius $q_{D_N,U}$ and mesh ratio $\rho_{D_N,U}$ by

$$h_{D_N,U} := \sup_{u \in U} \inf_{u^n \in U} \|u - u^n\|, \quad q_{D_N,U} := \frac{1}{2} \min_{i \neq j} \|u^i - u^j\|, \quad \rho_{D_N,U} := \frac{h_{D_N,U}}{q_{D_N,U}} \geq 1. \quad (3.1)$$

The fill distance (also known as the maximin distance [18] or dispersion [30]) is the maximum distance any point in U can be from a design point in D_N , and the separation radius is half the smallest distance between any two distinct points in D_N . The three quantities above provide measures of how uniformly the design points D_N are distributed in U .

The fill distance $h_{D_N,U}$ and the separation radius $q_{D_N,U}$ are decreasing functions of N , and these quantities will tend to zero as N tends to infinity for space-filling designs. The best possible rate of convergence for the fill distance for any choice of D_N is $h_{D_N,U} \leq CN^{-1/d_u}$ (see e.g. [30, 39]). The separation radius can decrease at an arbitrarily fast rate. The mesh ratio $\rho_{D_N,U}$, on the other hand, is a non-decreasing function of N . Point sets D_N for which $\rho_{D_N,U}$ can be bounded uniformly in N , i.e. sets for which the fill distance and the separation radius decrease at the same rate with N , are called quasi-uniform. In general, however, the mesh ratio can be strictly increasing in N .

3.1.1 Predictive Mean

We first consider the predictive mean $m_N^f(\widehat{\theta}_N)$. The main result is given in Theorem 3.5. To prove explicit error bounds, recall the following characterisation of the native space (also known as reproducing kernel Hilbert space) of the Matérn kernel.

Proposition 3.3. [54, Corollary 10.48] *Let U be a bounded Lipschitz domain, and let $k(\theta) = k_{\text{Mat}}(\theta)$, with $\theta_{\text{cov}} = \{\nu, \lambda, \sigma^2\} \subseteq S$, for some compact set $S \subseteq (0, \infty)^3$. Then the native space $H_{k_{\text{Mat}}(\theta)}(U)$ is equal to the Sobolev space $H^{\nu+d_u/2}(U)$ as a vector space, and the native space norm and the Sobolev norm are equivalent.*

There hence exist constants $C_{\text{low}}(\theta_{\text{cov}})$ and $C_{\text{up}}(\theta_{\text{cov}})$ such that for all $g \in H^{\nu+d_u/2}(U)$

$$C_{\text{low}}(\theta_{\text{cov}})\|g\|_{H_{k_{\text{Mat}}(\theta)}(U)} \leq \|g\|_{H^{\nu+d_u/2}(U)} \leq C_{\text{up}}(\theta_{\text{cov}})\|g\|_{H_{k_{\text{Mat}}(\theta)}(U)}. \quad (3.2)$$

Lemma 3.4. *For any compact set $S \subseteq (0, \infty)^3$, we have*

$$\max_{\theta_{\text{cov}} \in S} C_{\text{up}}(\theta_{\text{cov}})C_{\text{low}}(\theta_{\text{cov}})^{-1} \leq \max_{\theta_{\text{cov}} \in S} \max\{\lambda, \lambda^{-1}\} =: C_{3.4}.$$

Proof. First note that the conclusion of Proposition 3.3 holds also on the domain \mathbb{R}^{d_u} [54, Corollary 10.13]. By [54, Theorem 10.12], for any $g \in H_{k_{\text{Mat}}(\theta)}(\mathbb{R}^{d_u})$, we can express the native space norm as

$$\|g\|_{H_{k_{\text{Mat}}(\theta)}(\mathbb{R}^{d_u})}^2 = (2\pi)^{-d_u/2} \int_{\mathbb{R}^{d_u}} \frac{|\widehat{g}(\omega)|^2}{\widehat{k}(\theta; \omega)} d\omega,$$

where $\widehat{\cdot}$ denotes the Fourier transform. Furthermore, the Matérn covariance kernel has the Fourier transform [27, Example 7.17]

$$\widehat{k}(\theta; \omega) = \sigma^2 \frac{\Gamma(\nu + d_u/2)}{\Gamma(\nu) \pi^{d_u/2}} \frac{\lambda^{d_u}}{(1 + \lambda^2 \|\omega\|^2)^{\nu+d_u/2}}.$$

With

$$\|g\|_{H^{\nu+d_u/2}(\mathbb{R}^{d_u})}^2 = (2\pi)^{-d_u/2} \int_{\mathbb{R}^{d_u}} |\widehat{g}(\omega)|^2 (1 + \|\omega\|^2)^{\nu+d_u/2} d\omega,$$

it then follows that

$$\begin{aligned} \|g\|_{H^{\nu+d_u/2}(\mathbb{R}^{d_u})}^2 &\leq \frac{\sigma^2 \Gamma(\nu + d_u/2) \lambda^{d_u}}{\pi^{d_u/2} \Gamma(\nu)} \max\{1, \lambda^{-2}\} \|g\|_{H_{k_{\text{Mat}}(\theta)}(\mathbb{R}^{d_u})}^2 \\ &:= C_{\text{up}}(\theta_{\text{cov}})^2 \|g\|_{H_{k_{\text{Mat}}(\theta)}(\mathbb{R}^{d_u})}^2, \\ \|g\|_{H_{k_{\text{Mat}}(\theta)}(\mathbb{R}^{d_u})}^2 &\leq \frac{\pi^{d_u/2} \Gamma(\nu)}{\sigma^2 \Gamma(\nu + d_u/2) \lambda^{d_u}} \max\{1, \lambda^2\} \|g\|_{H^{\nu+d_u/2}(\mathbb{R}^{d_u})}^2 \\ &:= C_{\text{low}}(\theta_{\text{cov}})^{-2} \|g\|_{H_{k_{\text{Mat}}(\theta)}(\mathbb{R}^{d_u})}^2. \end{aligned}$$

On the bounded Lipschitz domain U , the same inequalities then hold for the norms

$\|g\|_{H^{\nu+d_u/2}(U)} = \inf_{\mathcal{E}g \in H^{\nu+d_u/2}(\mathbb{R}^{d_u})} \|g\|_{H^{\nu+d_u/2}(\mathbb{R}^{d_u})}$ and $\|g\|_{H_{k_{\text{Mat}}(\theta)}(U)} = \inf_{\mathcal{E}g \in H_{k_{\text{Mat}}(\theta)}(\mathbb{R}^{d_u})} \|g\|_{H_{k_{\text{Mat}}(\theta)}(\mathbb{R}^{d_u})}$, where $\mathcal{E}g$ denotes an extension of g . The claim of the Lemma then follows, with $C_{3.4} = \max_{\theta_{\text{cov}} \in S} \max\{\lambda, \lambda^{-1}\}$. \square

We then have the following result on the convergence of $m_N^f(\widehat{\theta}_N)$ to f as $N \rightarrow \infty$. In particular, it shows that we obtain convergence in a wide range of scenarios, under very mild assumptions on the estimated hyper-parameters. If the estimated hyper-parameters converge, we obtain the same convergence rate as in the case where all the hyper-parameters are fixed at the limiting value, cf [49, Proposition 3.4]. Note that Theorem 3.5 trivially also applies to the special case $\widehat{\theta}_N = \widehat{\theta}$, where a fixed value of the hyper-parameter is used.

Theorem 3.5. (Convergence in N of $m_N^f(\widehat{\theta}_N)$) *Suppose we have a sequence of estimates $\{\widehat{\theta}_N\}_{N=1}^\infty \subseteq S$, for some compact set $S \subseteq R_\theta$. Assume*

- (a) $U \subseteq \mathbb{R}^{d_u}$ is compact, with Lipschitz boundary, and satisfies an interior cone condition,
- (b) the native space $H_{k(\theta)}(U)$ is isomorphic to the Sobolev space $H^{\tau(\theta)}(U)$,
- (c) $f \in H^{\tilde{\tau}}(U)$, for some $\tilde{\tau} = n + r$, with $n \in \mathbb{N}$, $n > d_u/2$ and $0 \leq r < 1$,
- (d) $m(\theta) \in H^{\tilde{\tau}}(U)$ for all $\theta \in S$,
- (e) for some $N^* \in \mathbb{N}$, the quantities $\tau^- := \inf_{N \geq N^*} \tau(\widehat{\theta}_N)$ and $\tau^+ := \sup_{N \geq N^*} \tau(\widehat{\theta}_N)$ satisfy $\tau_- = n' + r'$, with $n' \in \mathbb{N}$, $n' > d_u/2$ and $0 \leq r' < 1$.

Then there exists a constant C , which is independent of f , m and N , such that for any $\beta \leq \tilde{\tau}$

$$\|f - m_N^f(\widehat{\theta}_N)\|_{H^\beta(U)} \leq C h_{D_N,U}^{\min\{\tilde{\tau}, \tau^-\} - \beta} \rho_{D_N,U}^{\max\{\tau^+ - \tilde{\tau}, 0\}} \left(\|f\|_{H^{\tilde{\tau}}(U)} + \sup_{N \geq N^*} \|m(\widehat{\theta}_N)\|_{H^{\tilde{\tau}}(U)} \right),$$

provided $N \geq N^*$ and $h_{D_N,U} \leq h_0$.

Proof. First, we note that it follows from (2.2) that $m_N^f(\widehat{\theta}_N) = m_N^{f,0}(\widehat{\theta}_N) + m(\widehat{\theta}_N) - m_N^{m,0}(\widehat{\theta}_N)$. An application of the triangle inequality hence gives

$$\begin{aligned} \|f - m_N^f(\widehat{\theta}_N)\|_{H^\beta(U)} &= \|f - m_N^{f,0}(\widehat{\theta}_N) - m(\widehat{\theta}_N) + m_N^{m,0}(\widehat{\theta}_N)\|_{H^\beta(U)} \\ &\leq \|f - m_N^{f,0}(\widehat{\theta}_N)\|_{H^\beta(U)} + \|m(\widehat{\theta}_N) - m_N^{m,0}(\widehat{\theta}_N)\|_{H^\beta(U)}. \end{aligned}$$

By assumption d), it follows from [29, Lemma 4.1] that for fixed $N \in \mathbb{N}$,

$$\|m(\widehat{\theta}_N) - m_N^{m,0}(\widehat{\theta}_N)\|_{H^\beta(U)} \leq C_1(\tilde{\tau}, \widehat{\theta}_N) h_{D_N,U}^{\tilde{\tau} - \beta} \|m(\widehat{\theta}_N)\|_{H^{\tilde{\tau}}(U)},$$

for some constant $C_1(\tilde{\tau}, \widehat{\theta}_N)$ independent of $h_{D_N,U}$ and m , provided $h_{D_N,U} \leq C_h(U)n^{-2}$. If $\tau(\widehat{\theta}_N) \leq \tilde{\tau}$, a similar estimate holds for f :

$$\|f - m_N^{f,0}(\widehat{\theta}_N)\|_{H^\beta(U)} \leq C_1(\tau(\widehat{\theta}_N), \widehat{\theta}_N) h_{D_N,U}^{\tau(\widehat{\theta}_N) - \beta} \|f\|_{H^{\tilde{\tau}}(U)},$$

provided $h_{D_N,U} \leq C_h(U)[\tau(\widehat{\theta}_N)]^{-2}$, where we have used $\|f\|_{H^{\tau(\widehat{\theta}_N)}(U)} \leq \|f\|_{H^{\tilde{\tau}}(U)}$ (see e.g. the proof of Lemma 3.4).

If $\tau(\widehat{\theta}_N) > \tilde{\tau}$, [29, Theorem 4.2] gives

$$\|f - m_N^{f,0}(\widehat{\theta}_N)\|_{H^\beta(U)} \leq C_2(\tilde{\tau}, \widehat{\theta}_N) h_{D_N,U}^{\tilde{\tau} - \beta} \rho_{D_N,U}^{\tau(\widehat{\theta}_N) - \tilde{\tau}} \|f\|_{H^{\tilde{\tau}}(U)},$$

for some constant $C_2(\tilde{\tau}, \widehat{\theta}_N)$ independent of f , $h_{D_N,U}$ and $\rho_{D_N,U}$, provided $h_{D_N,U} \leq C_h(U)n^{-2}$.

An inspection of the proofs of [29, Lemma 4.1 and Theorem 4.2] further gives the following. The constant $C_1(\tau', \theta)$ is of the form $C_1(\tau', \theta) = 2C'(\tau')C''(\theta)$, where

- $C'(\tau')$ is the constant appearing in the sampling inequality [28, Theorem 2.12]. This constant depends only on the integer part of τ' , and can hence only take a finite set of values for τ' in a compact set.
- $C''(\theta)$ is such that $\|m_N^f(\theta)\|_{H^{\tau(\theta)}(U)} \leq C''(\theta)\|f\|_{H^{\tau(\theta)}(U)}$. Using Propositions 3.1 and 3.3, an appropriate choice for $C''(\theta)$ is hence $C_{\text{up}}(\theta_{\text{cov}})C_{\text{low}}(\theta_{\text{cov}})^{-1}$, which is uniformly bounded on compact sets by Lemma 3.4.

Similarly, we have $C_2(\tau', \theta) = C'(\tau')(C'''(\tau') + 2C''(\theta))$, where C' and C'' are as above, and $C'''(\tau')$ is the constant appearing in the Bernstein inequality [29, Corollary 3.5].

The conclusion then follows, with $h_0 := C_h(U) \min_{\hat{\theta}_N \in S'} \min\{\lfloor \tau_+ \rfloor^{-2}, n^{-2}\}$ and

$$C = \max_{\hat{\theta}_N \in S'} \max\{C_1(\tilde{\tau}, \hat{\theta}_N), C_1(\tau(\hat{\theta}_N), \hat{\theta}_N), C_2(\tilde{\tau}, \hat{\theta}_N)\},$$

where $S' = S \cap \{\theta : \lfloor \tau \rfloor > d_u/2\}$. □

Assumption (a) in Theorem 3.5 is an assumption on the domain U being sufficiently regular, containing no sharp corners or cusps, and is satisfied, for example, for the unit cube $U = [0, 1]^{d_u}$. Assumption (b) reiterates that the native space of Matérn kernels is a Sobolev space. Theorem 3.5 applies in fact not just to Matérn kernels, but to any kernel which has a Sobolev space as native space, including the compactly supported Wendland functions [54].

Assumption (c) is an assumption on the regularity of the function f being emulated. We point out here that this assumption is rather mild, and modulo some technicalities (cf Remark 3.7), this assumption simply means that f should be an element of a Sobolev space that is compactly embedded into the space of continuous functions. We also point out here that Theorem 3.5 does not require the function f to be in the native space of any of the kernels $k(\hat{\theta}_N)$. The smoothness of f , denoted by $\tilde{\tau}$, can be both greater or smaller than the estimated smoothness $\tau(\hat{\theta}_N)$. The best possible convergence rates are obtained when the estimated smoothness matches the true smoothness of f (cf Remark 3.6). Recall that $h_{D_N, U}$ is decreasing in N , whereas $\rho_{D_N, U}$ is either constant or increasing in N . If we are underestimating the smoothness of f , then $\tau_- < \tilde{\tau}$, and we do not achieve the best possible exponent in $h_{D_N, U}$. If we are overestimating the smoothness of f , then $\tau_+ > \tilde{\tau}$ and we obtain a positive power of $\rho_{D_N, U}$. See section 6 for a further discussion on the optimality of the rates.

Assumption (d) ensures that the chosen mean m has at least the same regularity as f . This can be relaxed, but less regularity in m would lead to lower convergence rates in the error, so in practice, one should ensure that m is sufficiently smooth.

The quantities τ^- and τ^+ in assumption (e) can be thought of as $\liminf_{N \rightarrow \infty} \tau(\hat{\theta}_N)$ and $\limsup_{N \rightarrow \infty} \tau(\hat{\theta}_N)$, respectively. If $\lim_{N \rightarrow \infty} \tau(\hat{\theta}_N)$ exists, then this can be substituted for both quantities. Assumption (e) is the only assumption we make on the estimated hyper-parameters, other than that $\{\hat{\theta}_N\}_{N=1}^\infty \subseteq S$, for some compact set $S \subseteq R_\theta$. In particular, this means that the only assumptions required on $\hat{\lambda}_N$ and $\hat{\sigma}_N^2$ are that they are bounded away from zero and infinity. For the estimated smoothness $\hat{\nu}_N$, we again essentially require $0 < \hat{\nu}_N < \infty$, however, due some technical issues in the proof (cf Remark 3.7), we require a slightly larger lower bound on $\hat{\nu}_N$.

The error bounds in Theorem 3.5 can be translated into error bounds in terms of the number of design points N for specific choices of point sets. For example, the uniform grid $\tilde{D}_N = \{\frac{i}{N}\}_{i=1}^N$ with N points in $U = [0, 1]$ has fill distance and separation radius equal to $h_{\tilde{D}_N, U} = q_{\tilde{D}_N, U} = N^{-1}$.

In higher dimensions, the Cartesian product of one-dimensional uniform grids \tilde{D}_N with N points in $U = [0, 1]^{d_u}$ has fill distance $h_{\tilde{D}_N, U} = \sqrt{d_u} N^{-\frac{1}{d_u}}$ and separation radius $q_{\tilde{D}_N, U} = N^{-\frac{1}{d_u}}$. Hence uniform grids are quasi-uniform, with constant mesh ratio $\rho_{D_N, U} = \sqrt{d_u}$, and Theorem 3.5 gives a convergence rate of $N^{-\frac{\min\{\tilde{\tau}, \tau_-\}}{d_u}}$ in the L^2 -norm (corresponding to $\beta = 0$).

Low-discrepancy point sets, such as the Halton sequence, Sobol nets and lattice rules (see e.g. [30]), also have a small fill distance. The fill distance $h_{D_N, U}$ can be bounded in terms of the discrepancy $d_{D_N, U}$ as $h_{D_N, U} \leq d_{D_N, U}^{\frac{1}{d_u}}$ (see e.g. [30, Theorem 6.6]). Sequences such as the Halton sequence, for which $d_{D_N, U} \leq CN^{-1}(\log N)^{d_u}$, then have a fill distance $h_{D_N, U} \leq CN^{-\frac{1}{d_u}} \log N$, which up to the log factor decays at optimal rate. However, it is unclear whether these point sets are quasi-uniform. For further discussion on specific point sets and their fill distances, we refer the reader to [55] and the references therein.

For a given $f \in H^{\tilde{\tau}}(U)$, the fastest rate obtainable for $\|f - m_N^f(\hat{\theta}_N)\|_{L^2(U)}$ is $N^{-\frac{\tilde{\tau}}{d_u}}$. Given $\|f - m_N^f(\hat{\theta}_N)\|_{L^2(U)} = CN^{-\frac{\tilde{\tau}}{d_u}}$, the number of points needed to obtain an error $\|f - m_N^f(\hat{\theta}_N)\|_{L^2(U)} = \varepsilon$ is $N = C^{d_u/\tilde{\tau}} \varepsilon^{-d_u/\tilde{\tau}}$, and the number of points required to achieve a given accuracy hence grows exponentially in the dimension.

Remark 3.6. (Choice of ν) *Under the assumptions of Theorem 3.5, we have*

$$\|f - m_N^f(\theta_0)\|_{H^\beta(U)} \leq Ch_{D_N, U}^{\min\{\tilde{\tau}, \tau(\theta_0)\} - \beta} \rho_{D_N, U}^{\max\{\tau(\theta_0) - \tilde{\tau}, 0\}} \|f\|_{H^{\tilde{\tau}}(U)},$$

for any $\theta_0 \in S$. The best possible convergence rate in N of this bound is obtained when $\tilde{\tau} = \tau(\theta_0)$, i.e. when f is in the reproducing kernel Hilbert space corresponding to $k(\theta_0)$: $f \in H^{\tau(\theta_0)}(U) = H_{k(\theta_0)}(U)$. This is different to defining θ_0 such that f is a sample of the Gaussian process $GP(m(\theta_0; \cdot), k(\theta_0; \cdot, \cdot))$, since samples of a Gaussian process are almost surely not in the corresponding reproducing kernel Hilbert space. This point has also already been noted in [44].

Remark 3.7. (Valid choice of $\tilde{\tau}$) *The Sobolev space $H^{\tilde{\tau}}(U)$ is a reproducing kernel Hilbert space for any $\tilde{\tau} > d_u/2$. The restriction on $\tilde{\tau}$ in assumption (c) of Theorem 3.5 is hence slightly stronger than expected, requiring that the integer part of $\tilde{\tau}$ is greater than $d_u/2$. This is due to a technical detail in the proofs of [29, Lemma 4.1 and Theorem 4.2]. As noted in [55], one can use [5, Theorem 3.2] instead of [29], and assumption (c) in Theorem 3.5 can then be relaxed to the expected $f \in H^{\tilde{\tau}}(U)$, for some $\tilde{\tau} > d_u/2$. The rest of the proof remains identical. The same comment applies to τ_- in assumption (e) in Theorem 3.5, and assumptions (d) and (f) in Theorem 3.11.*

3.1.2 Predictive Variance

Next, we investigate the predictive variance $k_N(\hat{\theta}_N)$. An application of Proposition 3.2 gives the following result on the convergence of $k_N(\hat{\theta}_N)$ to 0.

Theorem 3.8. (Convergence in N of $k_N(\hat{\theta}_N)$) *Let the assumptions of Theorem 3.5 hold. Then there exists a constant C , independent of N , such that*

$$\|k_N^{1/2}(\hat{\theta}_N)\|_{L^2(U)} \leq Ch_{D_N, U}^{\min\{\tilde{\tau}, \tau_-\} - d_u/2 - \varepsilon} \rho_{D_N, U}^{\max\{\tau_+ - \tilde{\tau}, 0\}},$$

for any $N \geq N^*$, $h_{D_N, U} \leq h_0$ and $\varepsilon > 0$.

Proof. An application of Proposition 3.2, gives

$$\begin{aligned}
\|k_N^{1/2}(\hat{\theta}_N)\|_{L^2(U)} &= \left(\int_U k_N(\hat{\theta}_N; u, u) du \right)^{1/2} \\
&\leq C \sup_{u \in U} k_N(\hat{\theta}_N; u, u)^{1/2} \\
&= C \sup_{u \in U} \sup_{\|g\|_{H_{k(\hat{\theta}_N)}(U)=1}} |g(u) - m_N^g(\hat{\theta}_N; u)| \\
&= C \sup_{\|g\|_{H_{k(\hat{\theta}_N)}(U)=1}} \sup_{u \in U} |g(u) - m_N^g(\hat{\theta}_N; u)|.
\end{aligned}$$

The Sobolev embedding theorem gives the compact embedding of $H^{d_u/2+\varepsilon}(U)$ into the space of bounded continuous functions (see e.g. [2, Theorem 4.12, Part II]). Together with Theorem 3.5, this gives

$$\begin{aligned}
\|k_N^{1/2}(\hat{\theta}_N)\|_{L^2(U)} &\leq C' \sup_{\|g\|_{H_{k(\hat{\theta}_N)}(U)=1}} \|g(u) - m_N^g(\hat{\theta}_N; u)\|_{H^{d_u/2+\varepsilon}(U)} \\
&\leq C'' h_{D_N, U}^{\min\{\bar{\tau}, \tau_-\} - d_u/2 - \varepsilon} \rho_{D_N, U}^{\max\{\tau_+ - \bar{\tau}, 0\}} \sup_{\|g\|_{H_{k(\hat{\theta}_N)}(U)=1}} \|g\|_{H^{\bar{\tau}}(U)}.
\end{aligned}$$

Finally, using Proposition 3.3 gives $\|g\|_{H^{\bar{\tau}}(U)} \leq C_{\text{up}}(\hat{\theta}_N) \|g\|_{H_{k(\hat{\theta}_N)}(U)}$. The expression for $C_{\text{up}}(\hat{\theta}_N)$ derived in the proof of Lemma 3.4, and the compactness of S , then finish the proof. \square

Remark 3.9. (Dependency on $\hat{\lambda}_N, \hat{\sigma}_N^2$ and $\hat{\nu}_N$) *A careful inspection of the proofs of Theorems 3.5 and 3.8 reveals more details about the dependency on the different hyper-parameters. The correlation length $\hat{\lambda}_N$ enters only through the norm-equivalence constants $C_{\text{low}}(\hat{\theta}_N)$ and $C_{\text{up}}(\hat{\theta}_N)$, and the constants in Theorems 3.5 and 3.8 are larger for extreme (i.e. very large or very small) values of $\hat{\lambda}_N$. The constant in Theorem 3.5 is in fact independent of $\hat{\sigma}_N^2$, since it cancels out in the product $C_{\text{low}}(\hat{\theta}_N)^{-1} C_{\text{up}}(\hat{\theta}_N)$. This makes sense intuitively since the predictive mean $m_N^f(\hat{\theta}_N)$ is independent of $\hat{\sigma}_N^2$. In Theorem 3.8, $\hat{\sigma}_N^2$ enters linearly in the constant C through $C_{\text{up}}(\hat{\theta}_N)$. Again, this makes sense intuitively since $\hat{\sigma}_N^2$ enters as a multiplicative constant in the kernel $k(\hat{\theta}_N)$. The dependency on $\hat{\nu}_N$ is much more intricate, and influences the constants, as well as the convergence rates and valid choices for h_0 .*

3.2 Separable Matérn Covariance Kernels

Rather than the Matérn kernels employed in the previous section, suppose now that we use a separable Matérn covariance kernel $k_{\text{sepMat}}(\hat{\theta}_N)$, as defined in (2.6), to define the Gaussian process emulator (2.4). Due to the tensor product structure of the kernel $k_{\text{sepMat}}(\hat{\theta}_N)$, we will assume that our parameter domain U also has a tensor product structure $U = \prod_{j=1}^{d_u} U_j$, with $U_j \subset \mathbb{R}$ compact.

3.2.1 Predictive Mean

We again start with the predictive mean $m_N^f(\hat{\theta}_N)$. The main result in this section is Theorem 3.11. We have the following equivalent of Proposition 3.3, characterising the native space of separable Matérn covariance kernels on the tensor-product domain U .

Proposition 3.10. [32, 40] Suppose $U = \prod_{j=1}^{d_u} U_j$, with $U_j \subset \mathbb{R}$ bounded. Let $k(\theta) = k_{\text{sepMat}}(\theta)$ be a separable Matérn covariance kernel with $\theta_{\text{cov}} \in S$, for some compact set $S \subseteq (0, \infty)^{2d_u+1}$. Then the native space $H_{k(\theta)}(U)$ is equal to the tensor product Sobolev space $H_{\otimes d_u}^{\{\nu_j+1/2\}}(U) := \otimes_{j=1}^{d_u} H^{\nu_j+1/2}(U_j)$ as a vector space, and the native space norm and the Sobolev norm are equivalent.

There hence exist constants $C'_{\text{low}}(\theta_{\text{cov}})$ and $C'_{\text{up}}(\theta_{\text{cov}})$ such that for all $g \in H_{\otimes d_u}^{\{\nu_j+1/2\}}(U)$

$$C'_{\text{low}}(\theta_{\text{cov}}) \|g\|_{H_{k_{\text{sepMat}}(\theta)}(U)} \leq \|g\|_{H_{\otimes d_u}^{\{\nu_j+1/2\}}(U)} \leq C'_{\text{up}}(\theta; U) \|g\|_{H_{k_{\text{sepMat}}(\theta_{\text{cov}})}(U)}.$$

We will write $\{\beta_j\} \leq \{\alpha_j\}$ if $\beta_j \leq \alpha_j$ for all $1 \leq j \leq d_u$.

For our further analysis, we now want to make use of the convergence results from [32], related results are also found in [40] and the references in [32]. For the design points D_N , we will use Smolyak sparse grids [6]. For $1 \leq j \leq d_u$, we choose a sequence $X_j^{(i)} := \{x_{j,1}^{(i)}, \dots, x_{j,m_i}^{(i)}\}$, $i \in \mathbb{N}$, of nested sets of points in U_j . We then define the sparse grid $H(q, d_u) \subseteq U$ as the set of points

$$H(q, d_u) := \bigcup_{|\mathbf{i}|=q} X_1^{(i_1)} \times \dots \times X_{d_u}^{(i_{d_u})}, \quad (3.3)$$

where $|\mathbf{i}| = i_1 + \dots + i_{d_u}$ for a multi-index $\mathbf{i} \in \mathbb{N}^{d_u}$, and $q \geq d_u$. We denote by $N = N(q, d_u)$ the number of points in the sparse grid $H(q, d_u)$.

We then have the following equivalent of Theorem 3.5, which is again concerned with the convergence as $N \rightarrow \infty$.

Theorem 3.11. (Convergence in N of $m_N^f(\hat{\theta}_N)$) Suppose we have a sequence of estimates $\{\hat{\theta}_N\}_{N=1}^\infty \subseteq S$, for some compact set $S \subseteq R_\theta$. Assume

(a) $U = \prod_{j=1}^{d_u} U_j$, with $U_j \subset \mathbb{R}$ compact,

(b) D_N is chosen as the Smolyak sparse grid $H(q, d_u)$, for some $q \geq d_u$, with

$$h_{X_j^{(i)}, U_j} \leq C_1 m_i^{-r_h}, \quad \text{and} \quad \rho_{X_j^{(i)}, U_j} \leq C_2 m_i^{r_\rho},$$

for positive constants C_1, C_2, r_h and r_ρ independent of m_i and j ,

(c) the native space $H_{k(\theta)}(U)$ is isomorphic to a tensor product Sobolev space $H_{\otimes d_u}^{\{r_j(\theta)\}}(U)$,

(d) $f \in H_{\otimes d_u}^{\{\tilde{r}_j\}}(U)$, for some $\{\tilde{r}_j\}$ such that $\min_{1 \leq j \leq d_u} \tilde{r}_j \geq 1$,

(e) $m(\theta) \in H_{\otimes d_u}^{\{\tilde{r}_j\}}(U)$ for all $\theta \in S$,

(f) for some $N^* \in \mathbb{N}$, the quantities $r_{j,-} := \inf_{N \geq N^*} r_j(\hat{\theta}_N)$ and $r_{j,+} := \sup_{N \geq N^*} r_j(\hat{\theta}_N)$ satisfy $\min_{1 \leq j \leq d_u} r_{j,-} \geq 1$.

Then there exists a constant C , which is independent of f and N , such that

$$\|f - m_N^f(\hat{\theta}_N)\|_{H_{\otimes d_u}^{\{\beta_j\}}(U)} \leq C N^{-\alpha} (\log N)^{(1+\alpha')(d_u-1)} \left(\|f\|_{H_{\otimes d_u}^{\{\tilde{r}_j\}}(U)} + \sup_{N \geq N^*} \|m(\hat{\theta}_N)\|_{H_{\otimes d_u}^{\{\tilde{r}_j\}}(U)} \right),$$

for any $\{\beta_j\} \leq \{\tilde{r}_j\}$ and $N \geq N^*$, where

$$\alpha = \min_{1 \leq j \leq d_u} r_h(\min\{\tilde{r}_j, r_{j,-}\} - \beta_j) - r_\rho \max\{r_{j,+} - \tilde{r}_j, 0\},$$

$$\text{and } \alpha' = \min_{1 \leq j \leq d_u} r_h(\min\{\tilde{r}_j, r_{j,+}\} - \beta_j) - r_\rho \max\{r_{j,-} - \tilde{r}_j, 0\},$$

Proof. Theorem 3.11 is based on a generalisation of [32, Theorem 3], to the case where the function f is not necessarily in the native space of the kernel used to construct $m_N^f(\hat{\theta}_N)$. The structure of the proof remains identical, and we only need to replace [32, Proposition 4] with a corresponding result. All other assumptions required for [32, Theorem 3] remain valid.

So let $U' \subseteq \mathbb{R}$ be bounded, and let $X_n := \{x_1, \dots, x_n\} \subseteq U'$ be a set of n points in U' . For any $r \geq 1$ and $\beta \leq r$, let us denote by $\text{Id} : H^r(U') \rightarrow H^\beta(U')$ the identity operator, and by $S_{X_n}^\theta : H^r(U') \rightarrow H^\beta(U')$ the interpolation operator defined by $S_{X_n}^\theta(g) = m_n^{g,0}(\theta)$ (defined as in (2.2) with $m \equiv 0$). As in the proof of Theorem 3.5, we have

$$\begin{aligned} \|\text{Id} - S_{X_n}^\theta\|_{H^r(U') \rightarrow H^\beta(U')} &:= \sup_{\|g\|_{H^r(U')}=1} \|g - m_n^{g,0}(\theta)\|_{H^\beta(U')} \\ &\leq C_1(\theta) h_{X_n, U'}^{\min\{r, r(\theta)\} - \beta} \rho_{X_n, U'}^{\max\{r(\theta) - r, 0\}}, \end{aligned} \quad (3.4)$$

for any $\beta \leq r$ and $h_{X_n, U'} \leq h_0(\theta)$, where $H^{r(\theta)}(U')$ is the reproducing kernel Hilbert space corresponding to $k(\theta)$ used to construct $m_N^{g,0}(\theta)$. The fill distance $h_{X_n, U'}$ and mesh ratio $\rho_{X_n, U'}$ are as defined in (3.1), and the constants $C_1(\theta)$ and $h_0(\theta)$ are as in the proof of Theorem 3.5.

Following the proof of [32, Theorem 3] and replacing [32, Proposition 4] with (3.4), gives

$$\|g' - m_N^{g',0}(\theta)\|_{H_{\otimes d_u}^{\{\beta_j\}}(U)} \leq C_2(\theta) N^{-\alpha(\theta)} (\log N)^{(1+\alpha(\theta))(d_u-1)} \|g'\|_{H_{\otimes d_u}^{\{\tilde{r}_j\}}(U)}, \quad (3.5)$$

for any $\{\beta_j\} \leq \{\tilde{r}_j\}$ and $g' \in H_{\otimes d_u}^{\{\tilde{r}_j\}}(U)$, where

$$\alpha(\theta) = \min_{1 \leq j \leq d_u} r_h(\min\{\tilde{r}_j, r_j(\theta)\} - \beta_j) - r_\rho \max\{r_j(\theta) - \tilde{r}_j, 0\}.$$

By [32, Remark 4], the constant $C_2(\theta)$ depends on θ only through $\prod_{j=1}^{d_u} C_1(\theta)^{\alpha(\theta)}$.

To finish the proof, we use the triangle inequality and the equality $m_N^f(\hat{\theta}_N) = m_N^{f,0}(\hat{\theta}_N) + m(\hat{\theta}_N) - m_N^{m,0}(\hat{\theta}_N)$, and apply (3.5):

$$\begin{aligned} \|f - m_N^f(\hat{\theta}_N)\|_{H_{\otimes d_u}^{\{\beta_j\}}(U)} &\leq \|f - m_N^{f,0}(\hat{\theta}_N)\|_{H_{\otimes d_u}^{\{\beta_j\}}(U)} + \|m(\hat{\theta}_N) - m_N^{m,0}(\hat{\theta}_N)\|_{H_{\otimes d_u}^{\{\beta_j\}}(U)} \\ &\leq C_2(\hat{\theta}_N) N^{-\alpha(\hat{\theta}_N)} (\log N)^{(1+\alpha(\hat{\theta}_N))(d_u-1)} \left(\|f\|_{H_{\otimes d_u}^{\{\tilde{r}_j\}}(U)} + \|m(\hat{\theta}_N)\|_{H_{\otimes d_u}^{\{\tilde{r}_j\}}(U)} \right). \end{aligned}$$

The claim then follows as in Theorem 3.5. \square

Many of the same comments apply as to Theorem 3.5. Assumption (a) means that the domain U is of tensor-product structure, which is natural when using tensor-product kernels, and is satisfied, for example, for the unit cube $U = [0, 1]^{d_u}$.

Assumption (b) is a specific choice of design points D_N , and in contrast to Theorem 3.5, the choice of design points D_N as a sparse grid is explicitly used in the proof and is crucial for obtaining

the error bound. The values of r_h and r_ρ will depend on the particular choice of one-dimensional point sets. A particular choice of one-dimensional nested point sets often used in sparse grids are the Clenshaw-Curtis point sets $X_{CC}^{(i)}$, defined on $[-1, 1]$ by $m_i = 2^{(i-1)} + 1$, $X_{CC}^{(1)} = \{0\}$ and

$$x_{j,n}^{(i)} = -\cos\left(\frac{\pi(n-1)}{m_i-1}\right), \quad 1 \leq n \leq m_i.$$

General intervals $[a, b]$ can be dealt with through a linear transformation. For this particular point set, we can use the Lipschitz continuity of the cosine function, together with Kober's inequality $1 - \frac{x^2}{2} < \cos(x) < 1 - \frac{4x^2}{\pi^2}$ for $x \in [0, \frac{\pi}{2}]$ [42], to show that $h_{X_{CC}^{(i)}, U_j} \leq C_1 m_i^{-1}$ and $\rho_{X_{CC}^{(i)}, U_j} \leq C_2 m_i$. This shows that assumption (b) is satisfied for Clenshaw-Curtis point sets, with $r_h = r_\rho = 1$. Note that Clenshaw-Curtis points are known to cluster around the boundaries, so are not quasi-uniform (as evidenced by $r_\rho = 1$, which is sharp due Kober's inequality.) Alternatively, the one-dimensional point sets can be chosen as the uniform grids shown to be quasi-uniform in section 3.1, for which $r_h = 1$ and $r_\rho = 0$.

Assumption (c) reiterates that the native space of the separable Matérn kernel is a tensor-product Sobolev space, and Theorem 3.11 applies to any kernel with such a native space.

Assumption (d) is a regularity assumption on f , and again roughly corresponds to the function f being in a Sobolev space of mixed dominating smoothness that is compactly embedded into the space of continuous functions. We would ideally have the restriction $\min_{1 \leq j \leq d_u} \tilde{r}_j > 1/2$; however, we need a slightly stronger restriction due to some technicalities in the proof (cf Remark 3.7). Theorem 3.11 does not require the function f to be in the native space of any of the kernels $k(\hat{\theta}_N)$. The fastest convergence rates are again obtained when the estimated smoothness matches the true smoothness of f , see section 6 for a discussion on the optimality of the results.

Assumption (e) ensures that the mean m is at least as smooth as the function f , but this can again be relaxed. The assumptions on the estimated hyper-parameters $\{\nu_j, \lambda_j\}_{j=1}^{d_u}$ and σ^2 are again very mild. $\{\lambda_j\}_{j=1}^{d_u}$ and σ^2 are simply required to be bounded away from zero and infinity, and we require only a slightly larger lower bound on $\{\nu_j\}_{j=1}^{d_u}$ (cf assumption (f)).

Explicit convergence rates can again be obtained for specific choices of the design points. For sparse grids based on nested one-dimensional uniform grids, we obtain for $\beta_j \equiv 0$ the values $\alpha = \min_{1 \leq j \leq d_u} \min\{\tilde{r}_j, r_{j,-}\}$ and $\alpha' = \min_{1 \leq j \leq d_u} \min\{\tilde{r}_j, r_{j,+}\}$, giving the error estimate $\|f - m_N^f(\hat{\theta}_N)\|_{L^2(U)} \leq C N^{-\min_{1 \leq j \leq d_u} \min\{\tilde{r}_j, r_{j,-}\}} (\log N)^{(1+\alpha')(d_u-1)}$. Note in particular that the dimension d_u enters only in the log factor.

The set-up in Theorem 3.11 is much more restrictive than that of Theorem 3.5. Firstly, the design points need to be chosen as a sparse grid based on nested one-dimensional point sets. This limits the admissible choices of number of points N . The choice $m_1 = 1$ ensures that the growth of N , as a function of level q and dimension d_u , is as slow as possible. Furthermore, it can be shown that N grows at most polynomially in d_u (although it grows exponentially in the level q), see e.g. [31, Lemma 3.9].

Secondly, the set of functions which satisfy the regularity assumptions in Theorem 3.11 is a strict subset of those which satisfy the regularity assumptions in Theorem 3.5. The so-called mixed regularity of f assumed here is crucial to obtaining the error bound in Theorem 3.11. Although mixed regularity is quite a strong assumption, it is fulfilled in many important applications such as parametric partial differential equations [10, 31]. Also note that this function class is very different to those usually considered in the non-parametric regression literature (see e.g. [52, 22]), which correspond to the ones in Theorem 3.5.

The convergence rate in N is (up to logarithmic factors) independent of the dimension d_u , and can in high dimensions be much larger than the convergence rate obtained in Theorem 3.5. For illustrative purposes, let us look at the example where the smoothness of f is the same in every dimension, i.e. $\tilde{r}_j = \tilde{r}$. The fastest rate of convergence is obtained for example with sparse grids based on uniform one-dimensional grids and correctly estimated smoothness $r(\hat{\theta}_N) = \tilde{r}$, in which case the convergence rate in the $L^2(U)$ -norm is $\|f - m_N^f(\hat{\theta}_N)\|_{L^2(U)} \leq CN^{-\tilde{r}}(\log N)^{(1+\tilde{r})(d_u-1)}$. This means that to get an error $\|f - m_N^f(\hat{\theta}_N)\|_{L^2(U)} = \varepsilon$, it suffices to choose $N = c\varepsilon^{-1/\tilde{r}}|\log \varepsilon|^{(1+1/\tilde{r})(d_u-1)}$ (cf [32, Theorem 3]). In contrast with the setting of Theorem 3.5, the number of function evaluations required to achieve a given accuracy no longer grows exponentially in ε with dimension d_u .

3.2.2 Predictive Variance

Next, we investigate the predictive variance $k_N(\hat{\theta}_N)$, proving the convergence of $k_N(\hat{\theta}_N)$ to 0 as $N \rightarrow \infty$.

Theorem 3.12. (Convergence in N of $k_N(\hat{\theta}_N)$) *Let the assumptions of Theorem 3.11 hold. Then there exists a constant C , independent of N , such that*

$$\|k_N^{1/2}(\hat{\theta}_N)\|_{L^2(U)} \leq CN^{-\alpha}|\log N|^{(1+\alpha')(d_u-1)}$$

for any $N \geq N^*$ and $\varepsilon > 0$, where

$$\begin{aligned} \alpha &= \min_{1 \leq j \leq d_u} r_h(\min\{\tilde{r}_j, r_{j,-}\} - 1/2 - \varepsilon) - r_\rho \max\{r_{j,+} - \tilde{r}_j, 0\}, \\ \text{and } \alpha' &= \min_{1 \leq j \leq d_u} r_h(\min\{\tilde{r}_j, r_{j,+}\} - 1/2 - \varepsilon) - r_\rho \max\{r_{j,-} - \tilde{r}_j, 0\}. \end{aligned}$$

Proof. The proof is identical to Theorem 3.8, using Theorem 3.11 instead of Theorem 3.5. \square

3.3 Point-wise prediction error

We now briefly discuss the point-wise prediction error, i.e. the error in using $m_N^f(\hat{\theta}_N)$ or $f_N(\hat{\theta}_N)$ to predict $f(u)$, for an observed location $u \in U \setminus D_N$. This error is often considered in the spatial statistics literature, see e.g. [47, 48, 37].

For prediction using the mean $m_N^f(\hat{\theta}_N)$, we immediately obtain an error bound using Theorem 3.5 or 3.11, together with the Sobolev embedding theorem, which gives the compact embedding of $H^{d_u/2+\varepsilon}(U)$ into the space of bounded continuous functions (see e.g. [2, Theorem 4.12, Part II])

$$\begin{aligned} |f(u) - m_N^f(\hat{\theta}_N; u)| &\leq \sup_{u \in U} |f(u) - m_N^f(\hat{\theta}_N; u)| \\ &\leq \begin{cases} C\|f - m_N^f(\hat{\theta}_N)\|_{H^{d_u/2+\varepsilon}(U)}, & \text{in the set-up of Theorem 3.5,} \\ C\|f - m_N^f(\hat{\theta}_N)\|_{H_{\otimes d_u}^{1/2+\varepsilon}(U)}, & \text{in the set-up of Theorem 3.11.} \end{cases} \end{aligned}$$

Convergence of this prediction error to zero as $N \rightarrow \infty$ then follows, under very mild assumptions on the estimated hyper-parameters (as in Theorem 3.5 or 3.11).

For prediction using the predictive process $f_N(\hat{\theta}_N)$, we obtain a bound on the error

$$\mathbb{E}((f(u) - f_N(\hat{\theta}_N; u))^2)^{1/2} = k_N(\hat{\theta}_N; u, u)^{1/2},$$

where the expectation is over the distribution of $f_N(\widehat{\theta}_N)$. As in the proof of Theorem 3.8, we use Theorem 3.5 or 3.11, together with Proposition 3.2 and the Sobolev embedding theorem, which gives the compact embedding of $H^{d_u/2+\varepsilon}(U)$ into the space of bounded continuous functions (see e.g. [2, Theorem 4.12, Part II]):

$$\begin{aligned} k_N(\widehat{\theta}_N; u, u)^{1/2} &\leq \sup_{u \in U} k_N(\widehat{\theta}_N; u, u)^{1/2} \\ &\leq \begin{cases} C \sup_{\|g\|_{H_{k(\theta)}=1}} \|g - m_N^g(\widehat{\theta}_N)\|_{H^{d_u/2+\varepsilon}(U)}, & \text{in the set-up of Theorem 3.5,} \\ C \sup_{\|g\|_{H_{k(\theta)}=1}} \|g - m_N^g(\widehat{\theta}_N)\|_{H^{1/2+\varepsilon}_{\otimes^{d_u}}(U)}, & \text{in the set-up of Theorem 3.11.} \end{cases} \end{aligned}$$

Convergence of this prediction error to zero as $N \rightarrow \infty$ then follows, under very mild assumptions on the estimated hyper-parameters (as in Theorem 3.5 or 3.11).

We can also obtain convergence rates for the prediction error

$$\mathbb{E}((f(u) - m_N^f(\widehat{\theta}_N; u)^2))^{1/2},$$

where the expected value is now over some probability distribution over f . In particular, consider the setting $f \sim \text{GP}(0, k(\theta_0; \cdot, \cdot))$, for some true value θ_0 of the hyper-parameters, as is used for example [47, 48, 37]. Assume for simplicity that we are using Matérn kernels; similar arguments apply in the case of separable Matérn kernels.

Every sample of the Gaussian process $\text{GP}(0, k(\theta_0; \cdot, \cdot))$ belongs to the Sobolev space $H^{\nu_0}(U)$ (see e.g. [43]). Hence, we can apply Theorem 3.5 sample-wise, with $\tilde{\tau} = \nu_0$. The error bounds on $|f(u) - m_N^f(\theta; u)|$ coming from Theorem 3.5 depend on f only through $\|f\|_{H^{\tilde{\tau}}(U)}$, from which it follows that

$$\mathbb{E}_{\theta_0}((f(u) - m_N^f(\widehat{\theta}_N; u)^2))^{1/2} \leq C h_{D_N, U}^{\min\{\nu_0, \tau_-\} - d_u/2 - \varepsilon} \rho_{D_N, U}^{\max\{\tau_+ - \nu_0, 0\}} \mathbb{E}_{\theta_0}(\|f\|_{H^{\nu_0}(U)}^2)^{1/2}.$$

Since $\mathbb{E}_{\theta_0}(\|f\|_{H^{\nu_0}(U)}^2)^{1/2} < \infty$ (see e.g. [51, Propositions A.2.1 and A.2.3]), the above gives convergence to zero of the prediction error as N tends to ∞ , under the assumptions of Theorem 3.5. Note that this does not require any particular relation between the true hyper-parameters θ_0 and the employed hyper-parameters θ .

If the observed values of f at the design points D_N are used to form the estimate $\widehat{\theta}_N$, then

$$\mathbb{E}_{\theta_0}((f(u) - m_N^f(\widehat{\theta}_N; u)^2))^{1/2} \leq C \sup_{f \in H^{\nu_0}(U)} \left(h_{D_N, U}^{\min\{\nu_0, \tau_-\} - d_u/2 - \varepsilon} \rho_{D_N, U}^{\max\{\tau_+ - \nu_0, 0\}} \right) \mathbb{E}_{\theta_0}(\|f\|_{H^{\nu_0}(U)}^2)^{1/2},$$

and convergence to zero as $N \rightarrow \infty$ is again guaranteed. The restrictions on $\widehat{\sigma}_N^2$ and $\widehat{\lambda}_N$ are unchanged from Theorem 3.5. The restrictions on $\widehat{\nu}_N$ become slightly stronger, and require the quantities τ_- and τ_+ to be uniformly bounded in f . This can easily be achieved by imposing a fixed upper and lower bound on $\widehat{\nu}_N$, which is independent of the observed function values $f(u^1), \dots, f(u^N)$. Alternatively Hölder's inequality can be used to weaken the supremum in the above bound to an L^p -norm, with $p < \infty$.

4 Bayesian Inverse Problems

Our motivation for studying Gaussian process emulators was their use to approximate posterior distributions in Bayesian inverse problems. The inverse problem of interest is to determine the

unknown parameters $u \in U$ from noisy data $y \in \mathbb{R}^{d_y}$ given by

$$y = \mathcal{G}(u) + \eta. \quad (4.1)$$

We assume that the noise η is a realisation of the \mathbb{R}^{d_y} -valued Gaussian random variable $\mathcal{N}(0, \Gamma)$, for some known, positive-definite covariance matrix Γ , and that the parameter space U is a compact subset of \mathbb{R}^{d_u} , for some finite $d_u \in \mathbb{N}$. The map \mathcal{G} will be referred to as the *parameter-to-observation map* or *forward model*. For $x \in \mathbb{R}^m$, $m \in \mathbb{N}$, we denote by $\|x\|_2 = x^T x$ the Euclidean norm, and by $\|x\|_A = x^T A^{-1} x$ the norm weighted by (the inverse of) a positive-definite matrix $A \in \mathbb{R}^{m \times m}$.

We adopt a Bayesian perspective in which, in the absence of data, u is distributed according to a prior measure μ_0 . We are interested in the posterior distribution μ^y on the conditioned random variable $u|y$, which can be characterised as follows.

Proposition 4.1. (*[19, 50]*) *Suppose $\mathcal{G} : U \rightarrow \mathbb{R}^{d_y}$ is continuous and $\mu_0(U) = 1$. Then the posterior distribution μ^y on the conditioned random variable $u|y$ is absolutely continuous with respect to μ_0 and given by Bayes' Theorem:*

$$\frac{d\mu^y}{d\mu_0}(u) = \frac{1}{Z} \exp(-\Phi(u)),$$

where

$$\Phi(u) = \frac{1}{2} \|y - \mathcal{G}(u)\|_\Gamma^2 \quad \text{and} \quad Z = \mathbb{E}_{\mu_0}(\exp(-\Phi(u))).$$

Common to many of practical applications is that the evaluation of the parameter-to-observation map \mathcal{G} is analytically impossible and computationally very expensive, and, in simulations, it is therefore often necessary to approximate \mathcal{G} (or directly Φ) by a surrogate model. In this work, we are interested in Gaussian process emulators as surrogate models, as already discussed in [49].

Remark 4.2. (Distribution of the noise η) *The assumption that the distribution of the observational noise η is Gaussian with zero mean is not essential, and is for ease of presentation only. Inclusion of a non-zero mean, representing for example model discrepancy [21], is straightforward, and leads only to a shift in the misfit functional Φ . Other distributions, leading to other forms of the log-likelihood Φ , are also possible, and it is only the smoothness of Φ as a function of u that is important for the analysis presented in this paper. See for example [24] for a more general formulation.*

5 Approximation of the Bayesian Posterior Distribution

We now use the hierarchical Gaussian process emulator to define computationally cheaper approximations to the Bayesian posterior distribution μ^y . We will consider emulation of either the parameter-to-observation map $\mathcal{G} : U \rightarrow \mathbb{R}^{d_y}$ or the negative log-likelihood $\Phi : U \rightarrow \mathbb{R}$. An emulator of \mathcal{G} in the case $d_y > 1$ is constructed by emulating each entry independently.

The analysis presented in this section is for the most part independent of the specific covariance kernel used to construct the Gaussian process emulator. When the analysis does depend on the covariance kernel, we again consider the classical and separable Matérn families.

5.1 Approximation Based on the Predictive Mean

Using simply the predictive mean of a Gaussian process emulator of the parameter-to-observation map \mathcal{G} or the negative log-likelihood Φ , we can define the approximations $\mu_{\text{mean}}^{y, N, \mathcal{G}, \theta}$ and $\mu_{\text{mean}}^{y, N, \Phi, \theta}$,

given by

$$\begin{aligned}\frac{d\mu_{\text{mean}}^{y,N,\mathcal{G},\theta}}{d\mu_0}(u) &= \frac{1}{Z_{\text{mean}}^{N,\mathcal{G},\theta}} \exp\left(-\frac{1}{2} \|y - m_N^{\mathcal{G}}(\theta; u)\|_{\Gamma}^2\right), \\ Z_{\text{mean}}^{N,\mathcal{G},\theta} &= \mathbb{E}_{\mu_0}\left(\exp\left(-\frac{1}{2} \|y - m_N^{\mathcal{G}}(\theta)\|_{\Gamma}^2\right)\right), \\ \frac{d\mu_{\text{mean}}^{y,N,\Phi,\theta}}{d\mu_0}(u) &= \frac{1}{Z_{\text{mean}}^{N,\Phi,\theta}} \exp\left(-m_N^{\Phi}(\theta; u)\right), \\ Z_{\text{mean}}^{N,\Phi,\theta} &= \mathbb{E}_{\mu_0}\left(\exp\left(-m_N^{\Phi}(\theta)\right)\right),\end{aligned}$$

where $m_N^{\mathcal{G}}(\theta; u) = [m_N^{\mathcal{G}^1}(\theta; u), \dots, m_N^{\mathcal{G}^{d_y}}(\theta; u)] \in \mathbb{R}^{d_y}$.

We have the following result on the convergence of the approximate posterior distributions. A combination of Theorem 5.1 with Theorem 3.5 or Theorem 3.11 allows us to obtain convergence rates in N for the error in approximate posterior distributions.

Theorem 5.1. *Suppose we have a sequence of estimates $\{\hat{\theta}_N\}_{N=1}^{\infty} \subseteq S$, for some compact set $S \subseteq R_{\theta}$. Assume*

- (a) $U \subseteq \mathbb{R}^{d_u}$ is compact,
- (b) $\sup_{u \in U} \|\mathcal{G}(u) - m_N^{\mathcal{G}}(\hat{\theta}_N; u)\|$ and $\sup_{u \in U} |\Phi(u) - m_N^{\Phi}(\hat{\theta}_N; u)|$ can be bounded uniformly in N ,
- (c) $\sup_{u \in U} \|\mathcal{G}(u)\| \leq C_{\mathcal{G}} < \infty$.

Then there exist constants C_1 and C_2 , independent of N , such that

$$\begin{aligned}d_{\text{Hell}}(\mu^y, \mu_{\text{mean}}^{y,N,\mathcal{G},\hat{\theta}_N}) &\leq C_1 \left\| \mathcal{G} - m_N^{\mathcal{G}}(\hat{\theta}_N) \right\|_{L_{\mu_0}^2(U; \mathbb{R}^{d_y})}, \\ \text{and } d_{\text{Hell}}(\mu^y, \mu_{\text{mean}}^{y,N,\Phi,\hat{\theta}_N}) &\leq C_2 \left\| \Phi - m_N^{\Phi}(\hat{\theta}_N) \right\|_{L_{\mu_0}^2(U)}.\end{aligned}$$

Proof. This is essentially [49, Theorem 4.2]. The change from the distribution $\eta \sim \mathcal{N}(0, \sigma_{\eta}^2 \mathbf{I})$, considered in [49], to the more general distribution $\eta \sim \mathcal{N}(0, \Gamma)$ considered here in (4.1), only influences the values of the constants C_1 and C_2 , since all norms on \mathbb{R}^{d_y} are equivalent. The constants C_1 and C_2 involve taking the supremum over $\hat{\theta}_N$ over the corresponding constants in [49, Theorem 4.2], and we use the compactness of S to make sure this can be bounded independently of N . Furthermore, it is sufficient for the quantities in (b) to be bounded uniformly in N rather than converging to 0 as N tends to infinity (cf [49, Proof of Lemma 4.1]). \square

Since Theorems 3.5 and 3.11 hold only on bounded domains U , we have for simplicity assumed that U is bounded in assumption (a). This assumption can be relaxed in general (cf [24]). Assumption (b) is required to ensure the constants C_1 and C_2 are independent of N . Assumption (c) is satisfied for example when \mathcal{G} is continuous on U .

5.2 Approximation Based on the Predictive Process

We now consider approximations to the posterior distribution μ^y obtained using the full predictive processes \mathcal{G}_N and Φ_N . In contrast to the mean, the full Gaussian process also carries information about the uncertainty in the emulator due to only using a finite number of function evaluations to construct it. Randomising the approximations to \mathcal{G} and Φ , with the randomness tuned to represent the surrogate modelling error, can be crucial to obtaining statistically efficient sampling algorithms for the approximate posterior distributions [11, 9].

For the remainder of this section, we denote by $\nu_N^{\mathcal{G},\theta}$ the distribution of $\mathcal{G}_N(\theta)$ and by $\nu_N^{\Phi,\theta}$ the distribution of $\Phi_N(\theta)$, for $N \in \mathbb{N}$. The process \mathcal{G}_N consists of d_y independent Gaussian processes \mathcal{G}_N^j , so the measure $\nu_N^{\mathcal{G},\theta}$ is a product measure, $\nu_N^{\mathcal{G},\theta} = \prod_{j=1}^{d_y} \nu_N^{\mathcal{G}^j,\theta}$. Φ_N is a Gaussian process with mean m_N^Φ and covariance kernel k_N , and \mathcal{G}_N^j , for $j = 1, \dots, d_y$, is a Gaussian process with mean $m_N^{\mathcal{G}^j}$ and covariance kernel k_N . Replacing \mathcal{G} by \mathcal{G}_N in (4.1), we obtain the approximation $\mu_{\text{sample}}^{y,N,\mathcal{G},\theta}$ given by

$$\frac{d\mu_{\text{sample}}^{y,N,\mathcal{G},\theta}}{d\mu_0}(u) = \frac{1}{Z_{\text{sample}}^{N,\mathcal{G},\theta}} \exp\left(-\frac{1}{2} \|y - \mathcal{G}_N(\theta; u)\|_\Gamma^2\right),$$

where

$$Z_{\text{sample}}^{N,\mathcal{G},\theta} = \mathbb{E}_{\mu_0}\left(\exp\left(-\frac{1}{2} \|y - \mathcal{G}_N(\theta)\|_\Gamma^2\right)\right).$$

Similarly, we define for the predictive process Φ_N the approximation $\mu_{\text{sample}}^{y,N,\Phi,\theta}$ by

$$\frac{d\mu_{\text{sample}}^{y,N,\Phi,\theta}}{d\mu_0}(u) = \frac{1}{Z_{\text{sample}}^{N,\Phi,\theta}} \exp\left(-\Phi_N(\theta; u)\right), \quad Z_{\text{sample}}^{N,\Phi,\theta} = \mathbb{E}_{\mu_0}\left(\exp\left(-\Phi_N(\theta)\right)\right).$$

The measures $\mu_{\text{sample}}^{y,N,\mathcal{G},\theta}$ and $\mu_{\text{sample}}^{y,N,\Phi,\theta}$ are random approximations of the deterministic measure μ^y . The uncertainty in the posterior distribution introduced in this way can be thought of representing the uncertainty in the emulator, which in applications can be large (or comparable) to the uncertainty present in the observations. A user may want to take this into account to "inflate" the variance of the posterior distribution and avoid over-confident inference.

Deterministic approximations of the posterior distribution μ^y can now be obtained by fixing a sample of \mathcal{G}_N or Φ_N , or by taking the expected value with respect to the distribution of the Gaussian processes. The latter results in the marginal approximations

$$\begin{aligned} \frac{d\mu_{\text{marginal}}^{y,N,\mathcal{G},\theta}}{d\mu_0}(u) &= \frac{1}{\mathbb{E}_{\nu_N^{\mathcal{G},\theta}}(Z_{\text{sample}}^{N,\mathcal{G},\theta})} \mathbb{E}_{\nu_N^{\mathcal{G},\theta}}\left(\exp\left(-\frac{1}{2\sigma_\eta^2} \|y - \mathcal{G}_N(\theta; u)\|^2\right)\right), \\ \frac{d\mu_{\text{marginal}}^{y,N,\Phi,\theta}}{d\mu_0}(u) &= \frac{1}{\mathbb{E}_{\nu_N^{\Phi,\theta}}(Z_{\text{sample}}^{N,\Phi,\theta})} \mathbb{E}_{\nu_N^{\Phi,\theta}}\left(\exp\left(-\Phi_N(\theta; u)\right)\right). \end{aligned}$$

It can be shown that the above marginal approximation of the likelihood is optimal in the sense that it minimises a certain L^2 -error to the true likelihood [46]. The likelihood in the marginal approximations involves computing an expectation, and methods from the pseudo-marginal MCMC literature can be used within an MCMC method in this context [4, 9].

We have the following result on the convergence of the approximate posterior distributions, which can then be combined with Theorems 3.5 and 3.8 or Theorems 3.11 and 3.12 to obtain convergence rates in N for the error in approximate posterior distributions (cf [49, Corollary 4.10 and 4.12]). This requires the parameter-to-observation map \mathcal{G} to be sufficiently smooth.

Theorem 5.2. *Suppose we have a sequence of estimates $\{\hat{\theta}_N\}_{N=1}^\infty \subseteq S$, for some compact set $S \subseteq R_\theta$. Assume*

- (a) $U \subseteq \mathbb{R}^{d_u}$ is bounded,
- (b) $\sup_{u \in U} \left\| \mathcal{G}(u) - m_N^{\mathcal{G}}(\hat{\theta}_N; u) \right\|$ and $\sup_{u \in U} \left| \Phi(u) - m_N^{\Phi}(\hat{\theta}_N; u) \right|$ can be bounded uniformly in N , and $\sup_{u \in U} k_N(\hat{\theta}_N; u, u)$ converges to 0 as N tends to infinity,
- (c) $\sup_{u \in U} \|\mathcal{G}(u)\| \leq C_{\mathcal{G}} < \infty$,
- (d) $\mathbb{E} \left(\sup_{u \in U} \left(\Phi_N(u) - m_N^{\Phi}(\hat{\theta}_N; u) \right) \right)$ and $\mathbb{E} \left(\sup_{u \in U} \left(\mathcal{G}_N^j(u) - m_N^{\mathcal{G}}(\hat{\theta}_N; u) \right) \right)$, for $1 \leq j \leq d_y$, can be bounded uniformly in N .

Then there exist constants C_1, C_2, C_3 and C_4 , independent of N , such that for any $\delta > 0$,

$$d_{\text{Hell}}(\mu^y, \mu_{\text{marginal}}^{y, N, \mathcal{G}}) \leq C_1 \left\| \left(\mathbb{E}_{\nu_N^{\mathcal{G}}} \left(\|\mathcal{G} - \mathcal{G}_N(\hat{\theta}_N)\|^{1+\delta} \right) \right)^{1/(1+\delta)} \right\|_{L_{\mu_0}^2(U)},$$

$$d_{\text{Hell}}(\mu^y, \mu_{\text{marginal}}^{y, N, \Phi}) \leq C_2 \left\| \mathbb{E}_{\nu_N^{\Phi}} \left(|\Phi - \Phi_N(\hat{\theta}_N)|^{1+\delta} \right)^{1/(1+\delta)} \right\|_{L_{\mu_0}^2(U)}.$$

and

$$\left(\mathbb{E}_{\nu_N^{\mathcal{G}}} \left(d_{\text{Hell}}(\mu^y, \mu_{\text{sample}}^{y, N, \mathcal{G}})^2 \right) \right)^{1/2} \leq C_3 \left\| \left(\mathbb{E}_{\nu_N^{\mathcal{G}}} \left(\|\mathcal{G} - \mathcal{G}_N(\hat{\theta}_N)\|^{2+\delta} \right) \right)^{1/(2+\delta)} \right\|_{L_{\mu_0}^2(U)},$$

$$\left(\mathbb{E}_{\nu_N^{\Phi}} \left(d_{\text{Hell}}(\mu^y, \mu_{\text{sample}}^{y, N, \Phi})^2 \right) \right)^{1/2} \leq C_4 \left\| \left(\mathbb{E}_{\nu_N^{\Phi}} \left(|\Phi - \Phi_N(\hat{\theta}_N)|^{2+\delta} \right) \right)^{1/(2+\delta)} \right\|_{L_{\mu_0}^2(U)}.$$

Proof. This is essentially [49, Theorems 4.9 and 4.11]. As in Theorem 5.1, it is sufficient for the first two quantities in (b) to be bounded uniformly in N rather than converging to 0 as N tends to infinity (cf [49, Proof of Lemma 4.7]), and the change in the distribution of η in (4.1) only influences the constants. In [49], assumption (d) is replaced by an assumption involving the Sudakov-Fernique inequality (see Proposition 5.3 below), which is a sufficient condition for (d) to hold. However, that assumption is not satisfied in the case of the hierarchical Gaussian process emulators considered here, so we have introduced the more general assumption (d). \square

We have for simplicity again assumed that U is bounded in assumption (a). This assumption can be relaxed in general; see [24] for a more general statement of Theorem 5.2. Assumptions (b) and (d) are required to ensure the constants C_1, C_2, C_3 and C_4 are independent of N . Assumption (c) is satisfied for example when \mathcal{G} is continuous on U .

To verify assumption (d) in Theorem 5.2, we make use of the following two results.

Proposition 5.3. (Sudakov-Fernique Inequality, [36]) *Let g and h be scalar, Gaussian fields on the compact domain $U \subseteq \mathbb{R}^{d_u}$, and suppose g and h are almost surely bounded, i.e $\mathbb{P}[\sup_{u \in U} g(u) < \infty] = \mathbb{P}[\sup_{u \in U} h(u) < \infty] = 1$. Suppose $\mathbb{E}((g(u) - g(u'))^2) \leq \mathbb{E}((h(u) - h(u'))^2)$ and $\mathbb{E}(g(u)) = \mathbb{E}(h(u))$, for all $u, u' \in U$. Then*

$$\mathbb{E}(\sup_{u \in U} g(u)) \leq \mathbb{E}(\sup_{u \in U} h(u)).$$

Proposition 5.4. (Dudley's Inequality, [14, 23]) *Let g be a scalar Gaussian field on the compact domain $U \subseteq \mathbb{R}^{d_u}$, with zero mean $\mathbb{E}(g(u)) \equiv 0$, and define on U the pseudo-metric $d_g(u, u') = \mathbb{E}((g(u) - g(u'))^2)^{1/2}$. For $\epsilon > 0$, denote by $M(U, d_g, \epsilon)$ the minimal number of open d_g -balls of radius ϵ required to cover U . Then*

$$\mathbb{E}(\sup_{u \in U} g(u)) \leq C_D \int_0^\infty \sqrt{\log M(U, d_g, \epsilon)} d\epsilon,$$

for a constant C_D independent of g .

The Sudakov-Fernique inequality is a comparison inequality between Gaussian processes, whereas Dudley's inequality relates extreme values of a Gaussian process to its metric entropy. These results can be used to verify assumption (d) in Theorem 5.2 for general covariance functions $k(\theta)$, but we will in the following lemma concentrate on the particular case of covariance kernels chosen from the Matérn family or the separable Matérn family.

Lemma 5.5. *Suppose $U \subseteq \mathbb{R}^{d_u}$ is compact, and $k(\theta)$ is chosen as either the Matérn kernel in (2.5) with $\theta = \{\nu, \lambda, \sigma^2\}$ and $\nu > 1$, or the separable Matérn kernel in (2.6), with $\theta = \{\{\nu_j\}_{j=1}^{d_u}, \{\lambda_j\}_{j=1}^{d_u}, \sigma^2\}$ and $\nu_j > 1$, for $1 \leq j \leq d_u$. Assume $\{\hat{\theta}_N\}_{N \in \mathbb{N}} \subseteq S$, for some bounded set $S \subseteq (0, \infty)^{d_\theta}$. Then there exists a constant C , independent of N , such that*

$$\mathbb{E}(\sup_{u \in U} \Phi_N(u) - m_N^\Phi(\hat{\theta}_N; u)) \leq C, \quad \text{and} \quad \mathbb{E}(\sup_{u \in U} \mathcal{G}_N^j(u) - m_N^{\mathcal{G}^j}(\hat{\theta}_N; u)) \leq C, \quad j = 1, \dots, d_y.$$

Proof. We will give the proof for Φ , the proof for \mathcal{G}^j is similar. By [49, Lemma 4.8], it follows that the assumptions of Proposition 5.3 are satisfied with $g = \Phi_N - m_N^\Phi(\hat{\theta}_N)$ and $h = \tilde{\Phi}_N$, where $\tilde{\Phi}_N$ is the Gaussian process with mean zero and covariance kernel $k(\hat{\theta}_N)$. We hence have

$$\mathbb{E} \left(\sup_{u \in U} \left(\Phi_N(u) - m_N^\Phi(\hat{\theta}_N; u) \right) \right) \leq \mathbb{E}(\sup_{u \in U} \tilde{\Phi}_N(u)).$$

We now use Proposition 5.4, and consider separately the two types of covariance functions.

The covariance kernel $k_{\text{Mat}}(\hat{\theta}_N)$ is continuously differentiable, and hence Lipschitz continuous, jointly in u and u' for $\nu > 1$ (see e.g. [33, Lemma C.1]), and so

$$|k(\hat{\theta}_N; u, u') - k(\hat{\theta}_N; u, \tilde{u})| \leq L(\hat{\theta}_N) \|u' - \tilde{u}\|_2.$$

Thus, for any $u, u' \in U$,

$$\begin{aligned} d_{\tilde{\Phi}_N}(u, u')^2 &= \mathbb{E}((\tilde{\Phi}_N(u) - \tilde{\Phi}_N(u'))^2) \\ &= k_{\text{Mat}}(\hat{\theta}_N; u, u) - k_{\text{Mat}}(\hat{\theta}_N; u, u') - k_{\text{Mat}}(\hat{\theta}_N; u', u) + k_{\text{Mat}}(\hat{\theta}_N; u', u') \\ &\leq 2L(\hat{\theta}_N) \|u - u'\|_2, \end{aligned}$$

with Lipschitz constant $L(\widehat{\theta}_N) := \sup_{u, u' \in U} \|\nabla_u k_{\text{Mat}}(\widehat{\theta}_N; u, u')\|$. Using the formulas $\frac{d}{dr} r^\nu B_\nu(r) = -r^\nu B_{\nu-1}(r)$ [1] and $\Gamma(\nu) = \nu\Gamma(\nu-1)$, as well as the chain rule, then gives

$$\begin{aligned} \frac{d}{du_i} k_{\text{Mat}}(\{\sigma^2, \lambda, \nu\}; u, u') &= \frac{d}{du_i} \frac{\sigma^2}{\Gamma(\nu) 2^{\nu-1}} \left(\frac{\|u - u'\|_2}{\lambda} \right)^\nu B_\nu \left(\frac{\|u - u'\|_2}{\lambda} \right) \\ &= -2(u_i - u'_i) \frac{\|u - u'\|_2}{2\nu\lambda^2} k_{\text{Mat}}(\{\sigma^2, \lambda, \nu-1\}; u, u') \end{aligned}$$

Since $0 \leq k_{\text{Mat}}(\{\sigma^2, \lambda, \nu-1\}; u, u') \leq \sigma^2$, it then follows from the compactness of U and S that $L(\widehat{\theta}_N)$ can be bounded independently of N : $L(\widehat{\theta}_N) \leq L := \sup_{\theta \in S} L(\theta)$.

It follows that $M(U, d_{\widehat{\Phi}_N}, \epsilon) \sim \epsilon^{-2d_u}$ can be chosen independently of N , which together with Proposition 5.4 gives that $\mathbb{E}(\sup_{u \in U} \widetilde{\Phi}_N(u))$ can be bounded independently of N .

The proof for $k_{\text{sepMat}}(\widehat{\theta}_N)$ is similar. Iterating the inequality $|ab - cd| \leq a|b - d| + d|a - c|$, for real, positive numbers a, b, c, d , and using the Lipschitz continuity of the Matérn kernel for $\nu > 1$, as well as the bound $k_{\text{Mat}}(\widehat{\theta}_N; u_j, u'_j) \leq \widehat{\sigma}_N^2$, we have for any $u, u' \in U$,

$$\begin{aligned} d_{\widehat{\Phi}_N}(u, u')^2 &= \mathbb{E}((\widetilde{\Phi}_N(u) - \widetilde{\Phi}_N(u'))^2) \\ &= \prod_{j=1}^{d_u} k_{\text{Mat}}(\widehat{\theta}_N; u_j, u_j) - \prod_{j=1}^{d_u} k_{\text{Mat}}(\widehat{\theta}_N; u_j, u'_j) - \prod_{j=1}^{d_u} k_{\text{Mat}}(\widehat{\theta}_N; u'_j, u_j) + \prod_{j=1}^{d_u} k_{\text{Mat}}(\widehat{\theta}_N; u'_j, u'_j) \\ &\leq 2d_u(\widehat{\sigma}_N^2)^{d_u-1} L(\widehat{\theta}_N) \sum_{j=1}^{d_u} |u_j - u'_j| \\ &\leq \widetilde{L} \|u - u'\|_1, \end{aligned}$$

where $\widetilde{L} = 2d_u \sup_{\widehat{\theta}_N \in S} (\widehat{\sigma}_N^2)^{d_u-1} L(\widehat{\theta}_N)$. It follows that $M(U, d_{\widehat{\Phi}_N}, \epsilon) \sim \epsilon^{-2d_u}$, as in the case of Matérn kernels. This finishes the proof. \square

6 Conclusions and Discussion

Gaussian process regression is frequently used to approximate complex models. In this work, we looked at how the accuracy of the approximation depends on the number of model evaluations used to construct the Gaussian process emulator, in the setting where the hyper-parameters in the Gaussian process emulator are a-priori unknown and inferred as part of the emulation. The main results here are Theorems 3.5, 3.8, 3.11 and 3.12. These results show how fast we can expect the error to decay as a function of the number of model evaluations, and relate the decay rate of the error to the smoothness of both the function we are approximating and the employed kernel.

Generally speaking, we obtain error estimates of the form

$$\|f - m_N^f\|_{L^2(U)} \leq C_1 N^{-r_1} (\|f\| + \|m\|), \quad \text{and} \quad \|k_N^{\frac{1}{2}}\|_{L^2(U)} \leq C_2 N^{-r_2},$$

for the predictive mean m_N^f (as in (2.2)) and predictive variance k_N (as in (2.3)). The constants C_1 and C_2 depend on all hyper-parameters, whereas the rates r_1 and r_2 depend only on the estimated smoothness parameter(s) and the true smoothness parameter(s) (i.e. the smoothness of the given f). For a given function f , convergence of the Gaussian process emulator f_N is guaranteed under

very mild assumptions on the values of the estimated hyper-parameters, cf Theorems 3.5 and 3.11 and the discussions thereafter.

Let us briefly examine the optimality of our results. By [34, Theorem 23], we have the following bound for the best approximation of $f \in H^{\tilde{\tau}}(U)$ based on N function values $f(u^1), \dots, f(u^N)$:

$$c_1 N^{-\frac{\tilde{\tau}}{d_u}} \leq \inf_{\substack{\{u^1, \dots, u^N\} \subseteq U \\ \phi_1, \dots, \phi_N \in L^2(U)}} \sup_{\|f\|_{H^{\tilde{\tau}}(U)} \leq 1} \left\| f - \sum_{n=1}^N f(u^n) \phi_n \right\|_{L^2(U)} \leq c_2 N^{-\frac{\tilde{\tau}}{d_u}}.$$

We can then draw the following conclusions about the rates in Theorem 3.5:

- We obtain optimal convergence rates when the estimated smoothness matches the true smoothness, i.e. $\tau(\hat{\theta}_N) = \tilde{\tau}$, for any choice of design points D_N with optimal decay of the fill distance $h_{D_N, U} \leq CN^{-\frac{1}{d_u}}$.
- We obtain optimal convergence rates when the estimated smoothness is greater than or equal to the true smoothness, i.e. $\tau_- \geq \tilde{\tau}$, for any choice of *quasi-uniform* design points D_N with optimal decay of the fill distance $h_{D_N, U} \leq CN^{-\frac{1}{d_u}}$.
- We obtain suboptimal convergence rates when the estimated smoothness is greater than the true smoothness, i.e. $\tau_- \geq \tilde{\tau}$, and the design points D_N are not quasi-uniform. These issues arise due to the bound depending on $\|m_N^f(\hat{\theta}_N)\|_{H^{\tilde{\tau}}(U)}$, which can generally blow up as $N \rightarrow \infty$. If the mesh ratio grows with $\rho_{D_N, U} \leq CN^r$, Theorem 3.5 still gives $\|f - m_N^f\|_{L^2(U)} \rightarrow 0$ as $N \rightarrow \infty$, provided $\tau_+ \leq \tilde{\tau}(1 + (rd_u)^{-1})$. If τ_+ is too large, convergence is no longer guaranteed.
- We obtain suboptimal convergence rates when the estimated smoothness is less than the true smoothness, i.e. $\tau_+ \leq \tilde{\tau}$. Theorem 3.5 still gives $\|f - m_N^f\|_{L^2(U)} \rightarrow 0$ as $N \rightarrow \infty$, under very mild conditions on τ_- . We note that there are some results that allow to recover a faster convergence rate in this setting, but these typically require a particular relation between $\tau(\hat{\theta}_N)$ and $\tilde{\tau}$, and are hence difficult to apply in a general setting. For example, the results in [54, Section 11.5] require $\tilde{\tau} \geq 2\tau(\hat{\theta}_N)$.

A similar discussion applies to $\|k_N^{\frac{1}{2}}\|_{L^2(U)}$. There are no optimal rates for comparison, but we note that Theorem 3.8 in some settings gives almost the optimal rate $N^{-\frac{\tilde{\tau}}{d_u} + \frac{1}{2}}$ for $\|f - \sum_{n=1}^N f(u^n) \phi_n\|_{L^\infty(U)}$ (defined as above, see [34, Theorem 23]), which is crucially used as an upper bound in the proof.

A similar discussion also applies to Theorems 3.11 and 3.12. By e.g. [15, Theorem 4.5.1], we have the following bound for the best approximation of $f \in H_{\otimes d_u}^{\{\tilde{\tau}\}}(U)$ based on N function values $f(u^1), \dots, f(u^N)$, in the case of the d -dimensional torus $U = \mathbb{T}^d$:

$$c_1 N^{-\tilde{\tau}} \log(N)^{\tilde{\tau}(d_u-1)} \leq \inf_{\substack{\{u^1, \dots, u^N\} \subseteq U \\ \phi_1, \dots, \phi_N \in L^2(U)}} \sup_{\|f\|_{H_{\otimes d_u}^{\{\tilde{\tau}\}}(U)} \leq 1} \left\| f - \sum_{n=1}^N f(u^n) \phi_n \right\|_{L^2(U)}.$$

An algorithm that achieves this lower bound is not yet known. Up to the logarithmic factors, Theorem 3.11 again gives optimal convergence rates when the estimated smoothness matches the true smoothness and the fill distance of the one-dimensional point sets decays at the optimal rate

$h_{X_j^{(i)}, U_j} \leq C_1 m_i^{-1}$. For quasi-uniform one-dimensional point sets, we also get optimal convergence rates when the smoothness is overestimated. Underestimated smoothness leads to suboptimal convergence rates, and overestimated smoothness leads to convergence in Theorem 3.11 only if $r_+ < \tilde{r}(1 + r_\rho^{-1})$.

We note here that the estimation of hyper-parameters in an empirical Bayes' framework can in general have severe effects on issues such as consistency of MAP estimators; see the recent work [16] for a discussion. Gaussian process regression, viewed as an inverse problem to recover the function f from the function values $f(u^1), \dots, f(u^N)$, does however not fit into the framework considered in [16], and the results in this paper show that we do get consistency of the MAP estimate (i.e. the convergence of m_N^f to f) also with estimated hyper-parameters.

In section 3.3, we briefly examine the point-wise prediction error, and bound the error of using the predictive mean m_N^f or the predictive process f_N (as in (2.4)) to predict $f(u)$ at some unobserved location $u \in U \setminus D_N$. Again, we obtain convergence to zero as N tends to infinity under very mild assumptions on the estimated hyper-parameters.

Furthermore, we looked at the effect of approximating the parameter-to-observation map, or directly the log-likelihood, in a Bayesian inference problem by a Gaussian process emulator in section 5. This results in a computationally cheaper approximation to the Bayesian posterior distribution, which is crucial in large scale applications. The main results in this context are Theorems 5.1 and 5.2, which bound the error between the true posterior and the approximate posterior in terms of the accuracy of the Gaussian process emulator. These results give a justification for using Gaussian process emulators to approximate the Bayesian posterior, as they show that the approximate Bayesian posterior is close to the true posterior as long as the Gaussian process emulator approximates the data likelihood sufficiently well.

As a next step, it would be interesting to combine the results in this paper with results on the convergence of the estimated hyper-parameters $\hat{\theta}_N$. For example, the recent work [20] studies the asymptotics of the maximum likelihood estimator of the marginal variance σ^2 in the Matérn model, under assumptions similar to this work. It would also be useful to include the Gaussian covariance kernel, corresponding to the limit $\nu = \infty$ in the Matérn model, in our results.

Acknowledgements

The author would like to thank Andrew Stuart, Finn Lindgren, Peter Challenor, David Ginsbourger and Sören Wolfers for helpful discussions, and Toni Karvonen for pointing out the missed dependency of C' and h_0 on $\tau(\hat{\theta}_N)$ in the proof of Theorem 3.5. The author was partially supported by The Alan Turing Institute under the EPSRC grant EP/N510129/. The author would also like to thank the Isaac Newton Institute for Mathematical Sciences, Cambridge, for support and hospitality during the Uncertainty Quantification programme where work on this paper was partially undertaken. This programme was supported by EPSRC grant EP/K032208/1.

References

- [1] <http://functions.wolfram.com/bessel-typefunctions/besselk/20/showall.html>, Accessed June 02, 2020.
- [2] R. A. ADAMS AND J. J. FOURNIER, *Sobolev spaces*, Academic press, 2003.

- [3] E. ANDERES, *On the consistent separation of scale and variance for Gaussian random fields*, The Annals of Statistics, (2010), pp. 870–893.
- [4] C. ANDRIEU AND G. O. ROBERTS, *The pseudo-marginal approach for efficient Monte Carlo computations*, The Annals of Statistics, (2009), pp. 697–725.
- [5] R. ARCANGÉLI, M. C. L. DE SILANES, AND J. J. TORRENS, *Extension of sampling inequalities to sobolev semi-norms of fractional order and derivative data*, Numerische Mathematik, 121 (2012), pp. 587–608.
- [6] H.-J. BUNGARTZ AND M. GRIEBEL, *Sparse grids*, Acta numerica, 13 (2004), pp. 147–269.
- [7] R. H. CAMERON AND W. T. MARTIN, *Transformations of Wiener integrals under translations*, Annals of Mathematics, (1944), pp. 386–396.
- [8] T. CHOI AND M. J. SCHERVISH, *On posterior consistency in nonparametric regression problems*, Journal of Multivariate Analysis, 98 (2007), pp. 1969–1987.
- [9] P. R. CONRAD, M. GIROLAMI, S. SÄRKKÄ, A. STUART, AND K. ZYGALAKIS, *Statistical analysis of differential equations: introducing probability measures on numerical solutions*, Statistics and Computing, 27 (2017), pp. 1065–1082.
- [10] S. COTTER, M. DASHTI, AND A. STUART., *Variational data assimilation using targetted random walks*, Int. J. Numer. Meth. Fluids., 68 (2012), pp. 403–421.
- [11] T. CUI, C. FOX, AND M. J. O’SULLIVAN, *A posteriori stochastic correction of reduced models in delayed-acceptance mcmc, with application to multiphase subsurface inverse problems*, International Journal for Numerical Methods in Engineering, 118 (2019), pp. 578–605.
- [12] W. DAQING, *Fixed Domain Asymptotics and Consistent Estimation for Gaussian Random Field Models in Spatial Statistics and Computer Experiments*, PhD thesis, National University of Singapore, 2010.
- [13] P. DIAMOND AND M. ARMSTRONG, *Robustness of variograms and conditioning of kriging matrices*, Mathematical Geology, 16 (1984), pp. 809–822.
- [14] R. M. DUDLEY, *The sizes of compact subsets of Hilbert space and continuity of Gaussian processes*, Journal of Functional Analysis, 1 (1967), pp. 290–330.
- [15] D. DŨNG, V. TEMLYAKOV, AND T. ULLRICH, *Hyperbolic cross approximation*, Springer, 2018.
- [16] M. M. DUNLOP, T. HELIN, AND A. M. STUART, *Hyperparameter Estimation in Bayesian MAP Estimation: Parameterizations and Consistency*, arXiv preprint arXiv:1905.04365, (2019).
- [17] G.-A. FUGLSTAD, D. SIMPSON, F. LINDGREN, AND H. RUE, *Constructing priors that penalize the complexity of Gaussian random fields*, Journal of the American Statistical Association, 114 (2019), pp. 445–452.
- [18] M. E. JOHNSON, L. M. MOORE, AND D. YLVISAKER, *Minimax and maximin distance designs*, Journal of statistical planning and inference, 26 (1990), pp. 131–148.

- [19] J. P. KAIPIO AND E. SOMERSALO, *Statistical and Computational Inverse Problems*, Springer, 2005.
- [20] T. KARVONEN, G. WYNNE, F. TRONARP, C. J. OATES, AND S. SÄRKKÄ, *Maximum likelihood estimation and uncertainty quantification for gaussian process approximation of deterministic functions*, arXiv preprint arXiv:2001.10965, (2020).
- [21] M. C. KENNEDY AND A. O’HAGAN, *Bayesian calibration of computer models*, Journal of the Royal Statistical Society: Series B (Statistical Methodology), 63 (2001), pp. 425–464.
- [22] B. T. KNAPIK, B. T. SZABÓ, A. W. VAN DER VAART, AND J. H. VAN ZANTEN, *Bayes procedures for adaptive inference in inverse problems for the white noise model*, Probability Theory and Related Fields, 164 (2016), pp. 771–813.
- [23] M. LEDOUX AND M. TALAGRAND, *Probability in Banach Spaces: isoperimetry and processes*, Springer, 1991.
- [24] H. C. LIE, T. SULLIVAN, AND A. L. TECKENTRUP, *Random forward models and log-likelihoods in Bayesian inverse problems*, SIAM/ASA Journal on Uncertainty Quantification, 6 (2018), pp. 1600–1629.
- [25] W.-L. LOH, *Fixed-domain asymptotics for a subclass of Matérn-type Gaussian random fields*, The Annals of Statistics, 33 (2005), pp. 2344–2394.
- [26] W.-L. LOH AND T.-K. LAM, *Estimating structured correlation matrices in smooth Gaussian random field models*, The Annals of Statistics, 28 (2000), pp. 880–904.
- [27] G. J. LORD, C. E. POWELL, AND T. SHARDLOW, *An introduction to computational stochastic PDEs*, Cambridge University Press, 2014.
- [28] F. NARCOWICH, J. WARD, AND H. WENDLAND, *Sobolev bounds on functions with scattered zeros, with applications to radial basis function surface fitting*, Mathematics of Computation, 74 (2005), pp. 743–763.
- [29] F. J. NARCOWICH, J. D. WARD, AND H. WENDLAND, *Sobolev error estimates and a Bernstein inequality for scattered data interpolation via radial basis functions*, Constructive Approximation, 24 (2006), pp. 175–186.
- [30] H. NIEDERREITER, *Random Number Generation and quasi-Monte Carlo methods*, SIAM, 1994.
- [31] F. NOBILE, R. TEMPONE, AND C. G. WEBSTER, *A sparse grid stochastic collocation method for partial differential equations with random input data*, SIAM Journal on Numerical Analysis, 46 (2008), pp. 2309–2345.
- [32] F. NOBILE, R. TEMPONE, AND S. WOLFERS, *Sparse approximation of multilinear problems with applications to kernel-based methods in UQ*, Numerische Mathematik, (2017), pp. 1–34.
- [33] F. NOBILE AND F. TESEI, *A Multi Level Monte Carlo method with control variate for elliptic PDEs with log-normal coefficients*, Stochastic Partial Differential Equations: Analysis and Computations, 3 (2015), pp. 398–444.

- [34] E. NOVAK AND H. TRIEBEL, *Function spaces in Lipschitz domains and optimal rates of convergence for sampling*, Constructive approximation, 23 (2006), pp. 325–350.
- [35] A. O’HAGAN, *Bayesian analysis of computer code outputs: a tutorial*, Reliability Engineering & System Safety, 91 (2006), pp. 1290–1300.
- [36] G. D. PRATO AND J. ZABCZYK., *Stochastic Equations in Infinite Dimensions*, vol. 44 of Encyclopedia Math. Appl., Cambridge University Press, Cambridge, 1992.
- [37] H. PUTTER AND G. A. YOUNG, *On the effect of covariance function estimation on the accuracy of kriging predictors*, Bernoulli, 7 (2001), pp. 421–438.
- [38] C. E. RASMUSSEN AND C. K. WILLIAMS, *Gaussian processes for machine learning*, MIT Press, 2006.
- [39] A. REZNIKOV AND E. SAFF, *The covering radius of randomly distributed points on a manifold*, International Mathematics Research Notices, 2016 (2015), pp. 6065–6094.
- [40] C. RIEGER AND H. WENDLAND, *Sampling inequalities for sparse grids*, Numerische Mathematik, 136 (2017), pp. 439–466.
- [41] J. SACKS, W. J. WELCH, T. J. MITCHELL, AND H. P. WYNN, *Design and analysis of computer experiments*, Statistical science, (1989), pp. 409–423.
- [42] J. SÁNDOR, *On new refinements of kobers and jordans trigonometric inequalities*, Notes Number Theory Discrete Math, 19 (2013), pp. 73–83.
- [43] M. SCHEUERER, *Regularity of the sample paths of a general second order random field*, Stochastic Processes and their Applications, 120 (2010), pp. 1879–1897.
- [44] M. SCHEUERER, R. SCHABACK, AND M. SCHLATHER, *Interpolation of spatial data—A stochastic or a deterministic problem?*, European Journal of Applied Mathematics, 24 (2013), pp. 601–629.
- [45] B. SCHÖLKOPF, R. HERBRICH, AND A. J. SMOLA, *A generalized representer theorem*, in Computational learning theory, Springer, 2001, pp. 416–426.
- [46] M. SINSBECK AND W. NOWAK, *Sequential Design of Computer Experiments for the Solution of Bayesian Inverse Problems*, SIAM/ASA Journal on Uncertainty Quantification, 5 (2017), pp. 640–664.
- [47] M. L. STEIN, *Asymptotically efficient prediction of a random field with a misspecified covariance function*, The Annals of Statistics, 16 (1988), pp. 55–63.
- [48] ———, *A simple condition for asymptotic optimality of linear predictions of random fields*, Statistics & Probability Letters, 17 (1993), pp. 399–404.
- [49] A. STUART AND A. TECKENTRUP, *Posterior consistency for Gaussian process approximations of Bayesian posterior distributions*, Mathematics of Computation, 87 (2018), pp. 721–753.
- [50] A. M. STUART, *Inverse problems*, vol. 19 of Acta Numerica, Cambridge University Press, 2010, pp. 451–559.

- [51] A. W. VAART AND J. A. WELLNER, *Weak convergence and empirical processes: with applications to statistics*, Springer, 1996.
- [52] A. VAN DER VAART AND H. VAN ZANTEN, *Information rates of nonparametric Gaussian process methods*, Journal of Machine Learning Research, 12 (2011), pp. 2095–2119.
- [53] G. WAHBA, *Spline models for observational data*, Siam, 1990.
- [54] H. WENDLAND, *Scattered Data Approximation*, Cambridge University Press, 2005.
- [55] G. WYNNE, F.-X. BRIOL, AND M. GIROLAMI, *Convergence guarantees for gaussian process approximations under several observation models*, arXiv preprint arXiv:2001.10818, (2020).
- [56] Z. YING, *Maximum likelihood estimation of parameters under a spatial sampling scheme*, The Annals of Statistics, (1993), pp. 1567–1590.
- [57] H. ZHANG, *Inconsistent Estimation and Asymptotically Equal Interpolations in Model-Based Geostatistics*, Journal of the American Statistical Association, 99 (2004), pp. 250–261.
- [58] H. ZHANG AND D. L. ZIMMERMAN, *Towards reconciling two asymptotic frameworks in spatial statistics*, Biometrika, 92 (2005), pp. 921–936.