



UNIVERSITY OF HELSINKI

<https://helda.helsinki.fi>

Maximum likelihood estimation in Gaussian process regression is ill-posed

Karvonen, Toni; Oates, Chris

2023

Microtome Publishing

<http://hdl.handle.net/10138/567837>

Karvonen, T & Oates, C 2023, 'Maximum likelihood estimation in Gaussian process regression is ill-posed', *Journal of Machine Learning Research*, vol. 24, no. 120, 120, pp. 1-47.

Downloaded from Helda, University of Helsinki institutional repository. <https://helda.helsinki.fi>
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.
Please cite the original version.

Maximum Likelihood Estimation in Gaussian Process Regression is Ill-Posed

Toni Karvonen

*Department of Mathematics and Statistics
University of Helsinki
PL 56 (Pietari Kalmin katu 5)
00014 Helsingin yliopisto, Finland*

TONI.KARVONEN@HELSINKI.FI

Chris J. Oates

*School of Mathematics, Statistics and Physics
Newcastle University
Newcastle upon Tyne, NE1 7RU, United Kingdom*

CHRIS.OATES@NCL.AC.UK

Editor: Marc Peter Deisenroth

Abstract

Gaussian process regression underpins countless academic and industrial applications of machine learning and statistics, with maximum likelihood estimation routinely used to select appropriate parameters for the covariance kernel. However, it remains an open problem to establish the circumstances in which maximum likelihood estimation is well-posed, that is, when the predictions of the regression model are insensitive to small perturbations of the data. This article identifies scenarios where the maximum likelihood estimator *fails* to be well-posed, in that the predictive distributions are not Lipschitz in the data with respect to the Hellinger distance. These failure cases occur in the noiseless data setting, for any Gaussian process with a stationary covariance function whose lengthscale parameter is estimated using maximum likelihood. Although the failure of maximum likelihood estimation is part of Gaussian process folklore, these rigorous theoretical results appear to be the first of their kind. The implication of these negative results is that well-posedness may need to be assessed *post-hoc*, on a case-by-case basis, when maximum likelihood estimation is used to train a Gaussian process model.

Keywords: Gaussian processes, maximum likelihood estimation, ill-posedness, stationary kernels

1. Introduction

Gaussian process regression is a popular tool used to construct a predictive model for a response variable as a function of one or more covariates of interest. As a strict generalisation of classical linear regression, and with the support of production-level software, Gaussian process regression has found myriad applications in both the academic and industrial contexts. The success of Gaussian process regression, both in terms of predictive performance and quality of uncertainty quantification, is contingent on the use of a suitable covariance kernel K for the Gaussian process model. This is often achieved by choosing K from a parametric set $\{K_\theta\}_{\theta \in \Theta}$ of candidate covariance kernels, with the parameter $\theta \in \Theta$ being selected based on the training data set. The predictive performance of Gaussian processes is well-understood

in a variety of asymptotic settings (e.g., Stein, 1999; Anderes, 2010; van der Vaart and van Zanten, 2011; Bachoc et al., 2017; Karvonen et al., 2020). However, not much is known about the non-asymptotic setting when θ is estimated from a training data set. In particular, it is an open problem to understand when the predictions from a Gaussian process model are *well-posed*, in the sense that the predictions of the model are either continuous in the training data set or insensitive to small perturbations of the training data set. Theoretical understanding of well-posedness is urgently needed to support the use of Gaussian process regression in sensitive applications, such as in mine gas safety monitoring (Dong, 2012), malicious maritime activity detection (Kowalska and Peel, 2012), and climate modelling (Revell et al., 2018), where the reliability and robustness of predictions is critical.

Several methods exist to estimate θ , including maximum likelihood estimation (Mardia and Marshall, 1984), maximum *a posteriori* estimation (e.g., Cunningham et al., 2008), cross-validation (Geisser and Eddy, 1979), Bayesian inference (MacKay, 1992), kernel flows (Chen et al., 2021), and various bespoke approaches, for example when θ are the parameters of a neural network in a deep kernel (Wilson et al., 2016). Among these, maximum likelihood estimators θ_{ML} are arguably most widely used, for example being the default approach in Gaussian process software (Rasmussen and Nickisch, 2010; Pedregosa et al., 2011; Roustant et al., 2012; GPy, since 2012; Matthews et al., 2017). In addition to statistics and machine learning, maximum likelihood is also occasionally used in the applied mathematical literature to construct a kernel interpolant; see Fasshauer and McCourt (2015, Section 14.3) and Cavoretto (2021) for recent examples. Compared to other approaches, maximum likelihood is attractive due to the absence of any additional degrees of freedom (e.g., cross-validation requires a choice for how data are partitioned) and the possibility for automatic gradient-based optimisation.

Suppose that the data are modelled as being generated by a Gaussian process defined by a mean function m and a positive-definite covariance kernel K_θ . For a noiseless training data set $Y = (y_1, \dots, y_n) \in \mathbb{R}^n$, associated to a set X of distinct covariates $x_1, \dots, x_n \in \mathbb{R}^d$, a maximum likelihood estimator θ_{ML} of θ satisfies

$$\theta_{\text{ML}} \in \arg \min_{\theta \in \Theta} \ell(\theta | Y) \quad \text{with} \quad \ell(\theta | Y) = Y_m^\top K_\theta(X, X)^{-1} Y_m + \log \det K_\theta(X, X), \quad (1.1)$$

where $K_\theta(X, X) = (K_\theta(x_i, x_j))_{i,j=1}^n \in \mathbb{R}^{n \times n}$ is the positive-definite covariance matrix and $Y_m = (y_i - m(x_i))_{i=1}^n \in \mathbb{R}^n$. See Stein (1999, Section 6.4) or Rasmussen and Williams (2006, Section 5.4.1). Despite the simple form of the optimisation problem (1.1), there is only limited understanding of the behaviour of θ_{ML} in the *deterministic interpolation* regime, where the data $y_i = f(x_i)$ are, in truth, generated from a fixed but unknown function $f: \mathbb{R}^d \rightarrow \mathbb{R}$. An important open problem is to understand the behaviour of θ_{ML} in terms of the data-generating function f and the set of covariates, and the implications of this behaviour for predictions produced by the Gaussian process model.

The deterministic interpolation regime represents a simple but important instance of Gaussian process regression widely used in, for example, emulation of computer experiments (Sacks et al., 1989; Kennedy and O’Hagan, 2002), probabilistic numerical computation (Diaconis, 1988; Cockayne et al., 2019; Hennig et al., 2022), and Bayesian optimisation (Snoek et al., 2012). However, this regime is challenging to analyse, and results concerning maximum likelihood estimation appear limited to the asymptotic analyses. Xu and Stein (2017); Karvonen

et al. (2020); and Wang (2021) exploited a closed form for the maximum likelihood estimator of a *scale* parameter (i.e., $\theta = \{\sigma\}$ and $K_\theta = \sigma^2 K$) to analyse its behaviour as $n \rightarrow \infty$ in the fixed domain setting where the covariates $\{x_i\}_{i=1}^\infty$ are dense in a compact subset of \mathbb{R}^d . In a similar manner, Karvonen et al. (2019) analysed maximum likelihood estimation of the scale and *lengthscale* parameters (i.e., $\theta = \{\sigma, \lambda\}$ and $K_\theta(x, x') = \sigma^2 K(x/\lambda, x'/\lambda)$) for a particular non-stationary Ornstein–Uhlenbeck process, while Karvonen (2023) obtained asymptotic lower bounds on estimates of the smoothness parameter of the Matérn model. The non-asymptotic behaviour of maximum likelihood estimation in the deterministic interpolation regime has yet to be studied.

1.1 Contributions

The principal contributions of this work are to demonstrate how the concept of well-posedness can be applied to Gaussian process interpolation and how rigorous theoretical analysis of well-posedness can be performed. We prove that, in the deterministic interpolation regime, maximum likelihood estimation of a lengthscale parameter can fail to be well-posed. We emphasise that this is a *non-asymptotic* (i.e., the number n of observations is kept fixed) result, in contrast to earlier work, and is based on the observation that the maximum likelihood estimate λ_{ML} of a lengthscale parameter λ is infinite if the observations differ from the prior mean function by a constant vertical shift.

To be more precise, consider a Gaussian process with prior mean function m and stationary covariance function of the form

$$K_\lambda(x, y) = \Phi\left(\frac{x - y}{\lambda}\right) \quad \text{for } x, y \in \mathbb{R}^d, \quad (1.2)$$

where the lengthscale parameter $\lambda > 0$ determines the spatial correlation distance of the resulting Gaussian random field. If the function $\Phi: \mathbb{R}^d \rightarrow \mathbb{R}$ is continuous, the covariance function (1.2) tends pointwise to the constant $\Phi(0)$ as $\lambda \rightarrow \infty$. Thus, if the data are approximately shifted from the mean function m by a constant, it is intuitive that a large value will be taken by the maximum likelihood estimate λ_{ML} . It is rigorously proven in this article (see Theorem 2.3) that if (a) $n \geq 2$ is fixed, (b) the function Φ in (1.2) satisfies certain mild regularity conditions, and (c) there is a constant $c \in \mathbb{R}$ such that the data are *m-constant* in that

$$y_i = m(x_i) + c \quad \text{for } i = 1, \dots, n, \quad (1.3)$$

then $\lambda_{\text{ML}} = \infty$. This result can be viewed as a generalisation of the simple fact that maximum likelihood estimation fails if the data are fully explained by the prior mean: if $y_i = m(x_i)$ for every $i = 1, \dots, n$, the first term of $\ell(\lambda | Y)$ in (1.1) is zero and thus the estimate of λ must be infinite because $K_\lambda(X, X)$ tends to a singular matrix if and only if $\lambda \rightarrow \infty$. From $\lambda_{\text{ML}} = \infty$ it follows that the pointwise predictions produced by the fitted Gaussian process model assign all probability to a single point (see Theorem 2.6). This phenomenon is undesirable, as it is not reasonable to claim infinite precision from a finite data set (which could, in this case, involve as few as $n = 2$ values being observed). Secondly, we prove that if the data are *not m-constant*, then $\lambda_{\text{ML}} < \infty$ so that the predictive distributions are non-degenerate. Using these results we show in Section 2.4 that maximum likelihood estimation is not well-posed in general, in the sense that the resulting predictive distributions are not Lipschitz in the

data with respect to the Hellinger distance, which means that predictive inference can be sensitive to small perturbations of the data set.¹

A constant mean shift in (1.3) is *not* a pathological case in the deterministic interpolation context, though it does highlight one sense in which mathematical analysis may be easier when the data are assumed to come from a stochastic process, since a constant mean shift may be neglected as a measure zero event in that setting. For example, Gaussian process regression has been used to explore *discrepancy* between computer models (Brevault et al., 2020), where a constant mean shift between the output of two computer models for the same phenomenon could reasonably be expected. Similarly, in probabilistic numerical computation one could encounter a constant mean shift (e.g., when modelling an integrand in Bayesian cubature, if that integrand is in fact constant; Briol et al., 2019), or in applications of Bayesian optimisation, where the data are obtained in a region where the objective function is constant. The role of this article is therefore to highlight an important failure mode of maximum likelihood estimation in Gaussian process interpolation and, in doing so, to underscore the need for an improved theoretical understanding of parameter estimation in general.

It would be tempting to attribute these failings to the simplicity of the maximum likelihood estimator and the modelling choices that we consider, for surely something more sophisticated ought to render the problem well-posed. Section 3 demonstrates that there may not exist an easy solution in the deterministic interpolation regime, at least if the tractability of a Gaussian process model is to be retained. Namely, we prove the following extensions:

- (i) A certain cross-validation estimator of the scale parameter shares the undesirable property of producing infinite lengthscale estimates when the data are m -constant.
- (ii) Inclusion of a parametric prior mean function, which too is estimated from the data (i.e., as in universal kriging), does not prevent ill-posedness.
- (iii) Simultaneous maximum likelihood estimation of the scale and lengthscale parameters does not prevent ill-posedness.

What *does* guarantee well-posedness is the inclusion of a regularisation or a nugget term, which corresponds to an assumption by the user that the data are corrupted by additive Gaussian noise. It is intuitive that the maximum likelihood estimator should be more well-behaved if the data-generating process is noisy as it may be merely by chance that the data set is m -constant. Less intuitive is the numerical evidence in Section 4.1, which indicates that λ_{ML} is infinite even in the regularised setting as long as the regularisation term is sufficiently large. A possible interpretation of this observation is that for sufficiently large assumed noise level *any* data could have been plausibly generated by a constant mean shift of the prior mean. The use of regularisation can hardly be considered a proper solution to ill-posedness because the interpolation property of the conditional process, desirable in many applications, is lost and an influential degree of freedom is introduced. As demonstrated in Section 4.2, another option for guaranteeing that the lengthscale estimates are always

1. Note that this notion of well-posedness of a parameter estimation method is stronger than that of its *robustness*, defined by Gu et al. (2018, Section 3) essentially as the impossibility of obtaining a singular covariance matrix.

finite is to place a hyperprior on the lengthscale and use *maximum a posteriori* estimation. However, because the hyperprior determines the estimator, this approach is rather arbitrary and does not lend itself well to automation in software.

Most of our results apply to Matérn-type covariance functions whose Fourier transforms decay polynomially (see Assumption 2.2). Section 5 discusses generalisations of this Fourier assumption, lengthscale estimation for product kernels, and the use of general linear information, such as derivative data. Our proofs are predominantly based on reproducing kernel Hilbert space (RKHS) techniques and approximation theory in Sobolev spaces. Complete proofs are relegated to Section 7. No familiarity with RKHSs or Sobolev spaces is required to understand the statements of our main results and, outside the proofs, it is only in Section 5 that these concepts are used. Practical and theoretical implications of our results are discussed in Section 6.

1.2 Related Literature

For fixed kernel parameters θ , the mathematical properties of Gaussian process interpolation as $n \rightarrow \infty$ can be deduced from the equivalent perspective of optimal interpolation in an RKHS. See Fasshauer (2011); Scheuerer et al. (2013); and Kanagawa et al. (2018) for reviews on this equivalence. The setting where the data-generating function f is randomised has received considerable attention and a large number of results have been obtained. Consistency and asymptotic normality results for maximum likelihood estimation of scale and lengthscale parameters can be found in, for example, Ying (1991); Zhang (2004); Loh (2005); Du et al. (2009); Anderes (2010); Kaufman and Shaby (2013); Bachoc (2013); and Bevilacqua et al. (2019). The use of maximum likelihood to estimate *smoothness* parameters (i.e., controlling the differentiability of the Gaussian process sample paths) has also been considered, notably by Szabó et al. (2015); Knapik et al. (2016); Chen et al. (2021); and Karvonen (2023). A well known issue with maximum likelihood estimation occurs when θ is non-identifiable, as can happen when one attempts to simultaneously estimate scale, lengthscale, and smoothness parameters in the Matérn covariance model (Zhang, 2004). However, a lack of identifiability does not affect the predictions produced by the model or the combinations of parameters that can be identified (the so-called *microergodic* parameters). These related works, whilst providing useful insight and powerful theoretical tools and techniques, do not apply in the non-asymptotic deterministic interpolation regime where θ is estimated using maximum likelihood.

2. Maximum Likelihood Estimation Is Not Well-Posed

Given any function f and a set X of points x_1, \dots, x_n in the domain of f , we use $f(X)$ to denote the n -vector $(f(x_1), \dots, f(x_n))$. All vectors in this article are to be understood to be in column format.

2.1 Gaussian Process Interpolation

In Gaussian process interpolation, the data are modelled as discrete and noiseless observations from a Gaussian process sample path. A Gaussian process $f_{\text{GP}} \sim \text{GP}(m, K)$ on \mathbb{R}^d is a stochastic process characterised by a mean function $m: \mathbb{R}^d \rightarrow \mathbb{R}$ and a covariance kernel

$K: \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$:

$$\mathbb{E}[f_{\text{GP}}(x)] = m(x) \quad \text{and} \quad \text{Cov}[f_{\text{GP}}(x), f_{\text{GP}}(y)] = K(x, y)$$

for all $x, y \in \mathbb{R}^d$. Let $K(X, X) = (K(x_i, x_j))_{i,j=1}^n \in \mathbb{R}^{n \times n}$. That f_{GP} is Gaussian means that all its finite-dimensional distributions are normally distributed, which is to say that

$$f_{\text{GP}}(X) \sim \text{N}(m(X), K(X, X))$$

for any finite set of points $X \subset \mathbb{R}^d$. Throughout this article the covariance kernel is assumed to be (strictly) *positive-definite*, which means that

$$\sum_{i=1}^n \sum_{j=1}^n a_i a_j K(x_i, x_j) > 0 \tag{2.1}$$

for any $n \in \mathbb{N}$, any non-zero vector $a = (a_1, \dots, a_n)$, and any distinct points $x_i \in \mathbb{R}^d$. This implies that the covariance matrix $K(X, X)$ is positive-definite and non-singular if X consists of distinct points. A covariance kernel is *stationary* if there is a function $\Phi: \mathbb{R}^d \rightarrow \mathbb{R}$ such that

$$K(x, y) = \Phi(x - y) \quad \text{for all} \quad x, y \in \mathbb{R}^d. \tag{2.2}$$

If Φ is to yield a positive-definite covariance kernel, it is necessary that $\Phi(0) > 0$ because otherwise the condition (2.1) fails for $n = 1$. For the purposes of this article, the Matérn kernels constitute the most important class of stationary covariance functions. Let σ , λ , and ν be positive. A Matérn kernel with scale σ , lengthscale λ , and smoothness ν is given by

$$K(x, y) = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu} \|x - y\|}{\lambda} \right)^\nu \mathcal{K}_\nu \left(\frac{\sqrt{2\nu} \|x - y\|}{\lambda} \right), \tag{2.3}$$

where Γ is the Gamma function and \mathcal{K}_ν the modified Bessel function of the second kind.

Suppose that a noiseless training data set $Y = (y_1, \dots, y_n) \in \mathbb{R}^n$ associated to a set X of distinct covariates $x_1, \dots, x_n \in \mathbb{R}^d$ has been obtained and let $Y_m = Y - m(X)$. Define

$$K(x, X) = (K(x, x_i))_{i=1}^n \in \mathbb{R}^n \quad \text{and} \quad K(X, y) = (K(x_i, y))_{i=1}^n \in \mathbb{R}^n$$

for any $x, y \in \mathbb{R}^d$. The mean and covariance functions of the conditional Gaussian process $f_{\text{GP}} | Y$ are obtained from the well known expressions

$$\mu(x) = \mathbb{E}[f_{\text{GP}}(x) | Y] = m(x) + K(x, X)^\top K(X, X)^{-1} Y_m \tag{2.4}$$

and

$$P(x, y)^2 = \text{Cov}[f_{\text{GP}}(x), f_{\text{GP}}(y) | Y] = K(x, y) - K(x, X)^\top K(X, X)^{-1} K(X, y). \tag{2.5}$$

We use the simplified notation $P(x, x)^2 = P(x)^2$ for the conditional variance. It is often convenient, especially in our proofs, to use the function

$$s(x) = K(x, X)^\top K(X, X)^{-1} Y_m \tag{2.6}$$

and write the conditional mean as $\mu = m + s$. The assumption that there is no noise means that we are in the deterministic interpolation regime where μ is an interpolant to the data, which is to say that $\mu(x_i) = y_i$ and $P(x_i) = 0$ for every $i = 1, \dots, n$. Crucially, this assumption allows us to leverage well known equivalences, reviewed in Section 7.1, between Gaussian process interpolation and kernel-based minimum-norm interpolation (e.g., Scheuerer et al., 2013; Kanagawa et al., 2018).




Figure 1: Translates of Matérn kernels in (2.3), as well as sample paths from the corresponding Gaussian processes on the domain $[0, 1]$, with $\sigma = 1$, $\nu = 3/2$, and three different λ .

2.2 Maximum Likelihood Estimation

Maximum likelihood estimation is the most common method used to select parameters $\theta \in \Theta$ of a parametrised covariance kernel K_θ . Under the Gaussian process model $\text{GP}(m, K_\theta)$, the probability density function of the data Y given θ is (e.g., Rasmussen and Williams, 2006, Section 5.4.1)

$$\frac{1}{\det(2\pi K_\theta(X, X))^{1/2}} \exp\left(-\frac{1}{2} Y_m^\top K_\theta(X, X)^{-1} Y_m\right). \quad (2.7)$$

Maximising (2.7) over Θ is equivalent to minimising

$$\ell(\theta | Y) = Y_m^\top K_\theta(X, X)^{-1} Y_m + \log \det K_\theta(X, X), \quad (2.8)$$

which we call the *modified log-likelihood function*, being the log-likelihood function up to subtraction and multiplication by negative constants. That is, any maximum likelihood estimate (which may not be unique if ℓ is multimodal) of θ satisfies

$$\theta_{\text{ML}} \in \arg \min_{\theta \in \Theta} \ell(\theta | Y).$$

The two terms that comprise the modified log-likelihood function (2.8) are usually called the (negative) *data-fit* and *model complexity* terms, respectively. Throughout the article we use subscripts to denote that various quantities depend on the kernel parameters. For example, the parameter-dependent conditional mean and covariance are

$$\begin{aligned} \mu_\theta(x) &= m(x) + K_\theta(x, X)^\top K_\theta(X, X)^{-1} Y_m, \\ P_\theta(x, y)^2 &= K_\theta(x, y) - K_\theta(x, X)^\top K_\theta(X, X)^{-1} K_\theta(X, y). \end{aligned}$$

We are interested in estimation of the lengthscale parameter $\lambda > 0$ which parametrises any stationary kernel of the form (2.2) as

$$K_\lambda(x, y) = \Phi\left(\frac{x - y}{\lambda}\right)$$

and determines the spatial correlation distance of the resulting Gaussian process (see Figure 1). A maximum likelihood estimate of λ therefore satisfies

$$\lambda_{\text{ML}} \in \arg \min_{\lambda > 0} \ell(\lambda | Y) = \arg \min_{\lambda > 0} \{Y_m^\top K_\lambda(X, X)^{-1} Y_m + \log \det K_\lambda(X, X)\}.$$

Specifically, we are interested in rigorously proving that $\lambda_{\text{ML}} = \infty$ in certain cases. As λ increases, the kernel K_λ tends pointwise to $\Phi(0)$, which is not positive-definite, and the covariance matrices in the predictive mean (2.4) and covariance (2.5) become singular if $n > 1$. However, the limits $\lim_{\lambda \rightarrow \infty} \mu_\lambda$ and $\lim_{\lambda \rightarrow \infty} P_\lambda(x, y)^2$ do exist for most commonly used kernels, such as Matérns and the Gaussian, and it is in this limiting sense that one should interpret the conditional moments when $\lambda = \infty$. This is discussed in more detail in Section 2.3.

Our results on the behaviour of λ_{ML} (and, later, other estimators of λ) show that Gaussian process interpolation fails, producing degenerate predictive distributions, when the data are m -constant and the function Φ has a polynomially decaying Fourier transform.

Definition 2.1 (Constant data) *Given a mean function m , we say that the data Y are m -constant if there is a constant $c \in \mathbb{R}$ such that*

$$Y_m = Y - m(X) = (c, \dots, c) \in \mathbb{R}^n.$$

Let $\widehat{f}(\xi) = \int_{\mathbb{R}^d} g(x) e^{-i\xi^\top x} dx$ denote the Fourier transform of an integrable function $f: \mathbb{R}^d \rightarrow \mathbb{R}$. We use the following assumption on the rate of decay of the Fourier transform of a stationary kernel.

Assumption 2.2 (Stationary Sobolev kernel) *There are a continuous and integrable function $\Phi: \mathbb{R}^d \rightarrow \mathbb{R}$ and constants $C_1, C_2 > 0$ and $\alpha > d/2$ such that $K(x, y) = \Phi(x - y)$ for all $x, y \in \mathbb{R}^d$ and*

$$C_1(1 + \|\xi\|^2)^{-\alpha} \leq \widehat{\Phi}(\xi) \leq C_2(1 + \|\xi\|^2)^{-\alpha} \tag{2.9}$$

for all $\xi \in \mathbb{R}^d$.

If $d = 1$ and Assumption 2.2 holds for $\alpha = p + 1 \in \mathbb{N}$, the kernel is p times differentiable in that the derivative

$$\left. \frac{\partial^{2p}}{\partial x^p \partial y^p} K(x, y) \right|_{x=0} = (-1)^p \Phi^{2p}(0)$$

exists. As a consequence, the process $f_{\text{GP}} \sim \text{GP}(m, K)$ is p times mean-square differentiable (Stein, 1999, Section 2.4). That a kernel satisfying (2.9) is called a *Sobolev kernel* is because its RKHS is norm-equivalent to the Sobolev space $W_2^\alpha(\mathbb{R}^d)$ of order α . The norm-equivalence is a crucial ingredient in several of our proofs and is reviewed, together with Sobolev spaces, in more detail in Section 7.3. One can also prove that the sample paths of f_{GP} are elements of certain Sobolev spaces (Scheuerer, 2011; Steinwart, 2019; Henderson, 2022). The Fourier transform of the function

$$\Phi(z) = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} (\sqrt{2\nu} \|z\|)^\nu \mathcal{K}_\nu(\sqrt{2\nu} \|z\|), \quad z \in \mathbb{R}^d,$$

which defines a Matérn kernel in (2.3), is (e.g., Stein, 1999, p. 49)

$$\widehat{\Phi}(\xi) = \sigma^2 \frac{\Gamma(\nu + d/2)}{\pi^{d/2} \Gamma(\nu)} (2\nu)^\nu (2\nu + \|\xi\|^2)^{-(\nu + d/2)}. \tag{2.10}$$

Therefore a Matérn kernel with smoothness $\nu > 0$ satisfies Assumption 2.2 with $\alpha = \nu + d/2$.

With these preliminaries we are ready to state the main result of this article on the behaviour of maximum likelihood estimates of λ . The result is illustrated in Figure 2.




Figure 2: The modified log-likelihood function for the Matérn kernel (2.3) with parameters $\sigma = 1$ and $\nu = 5/2$ given three different data vectors Y obtained at the points $X = \{1, 1.2, 2\} \subset \mathbb{R}$. Note the non-linear y -axis.

Theorem 2.3 (Maximum likelihood estimation) *Suppose that the kernel K satisfies Assumption 2.2 and $n \geq 2$. If the data Y are m -constant, then*

$$\lim_{\lambda \rightarrow \infty} \ell(\lambda | Y) = -\infty \quad \text{and} \quad \lambda_{\text{ML}} = \infty. \quad (2.11)$$

If the data Y are not m -constant, then

$$\lim_{\lambda \rightarrow \infty} \ell(\lambda | Y) = \infty \quad \text{and} \quad \lambda_{\text{ML}} < \infty. \quad (2.12)$$

Proof See Section 7.4. The proof uses RKHS techniques to show that under Assumption 2.2 the data-fit term, as a function of λ , (a) is upper bounded if the data are m -constant and (b) grows polynomially if the data are not m -constant, while the covariance matrix tends to the matrix consisting of $\Phi(0)$'s as $\lambda \rightarrow \infty$ and it can be shown that its log-determinant (i.e., the model complexity) tends to negative infinity with at most rate $-\log \lambda$. ■

Remark 2.4 *It would be very interesting and useful to obtain a more quantitative version of Theorem 2.3 which would, for example, state that*

$$\lambda_{\text{ML}} \geq g(\text{const}_m(Y)) \quad (2.13)$$

for some measure $\text{const}_m(Y)$ of how far Y are from being m -constant (i.e., $\text{const}_m(Y) = 0$ if and only if Y are m -constant) and some decreasing function $g: (0, \infty) \rightarrow \mathbb{R}$ such that $g(r) \rightarrow \infty$ as $r \rightarrow 0$. Unfortunately, the techniques we use to prove Theorem 2.3 are not precise enough to prove any form of (2.13).

Remark 2.5 *For simplicity, suppose that $m \equiv 0$ so that the data being m -constant means that $Y = (c, \dots, c)$ for some $c \in \mathbb{R}$. Because the case $\lambda = \infty$ formally corresponds to a model*

with a constant kernel, one might be tempted to interpret Theorem 2.3 as a special case of a general theorem which would state that θ_{ML} takes the value $\bar{\theta}$ for which the data are fully explained by a single translate of $K_{\bar{\theta}}$. However, there can be no such theorem. For suppose that there are $\bar{\theta} \in \Theta$ and $i \in \{1, \dots, n\}$ such that the data could have been generated by the translate of $K_{\bar{\theta}}$ at x_i . That is, $Y = aK_{\bar{\theta}}(x_i, X)$ for some $a \in \mathbb{R}$. Then

$$\ell(\theta | Y) = Y^\top K_\theta(X, X)^{-1}Y + \log \det K_\theta(X, X)$$

would have to attain its minimum at $\theta = \bar{\theta}$. But because the data-fit term is non-negative and the model complexity term does not depend on the data, it is clear that a minimum can be attained at $\bar{\theta}$ only by “chance” or if $\log \det K_{\bar{\theta}}(X, X) = -\infty$, which happens only if $K_{\bar{\theta}}(X, X)$ is singular or, in other words, if $K_{\bar{\theta}}$ is not a valid positive-definite kernel. That is, there can be no general theorem that $\theta_{\text{ML}} = \bar{\theta}$ if it is required that $K_{\bar{\theta}}$ be a well-defined positive-definite kernel.

Next we discuss the behaviour of the conditional mean and covariance and give a precise meaning to ill-posedness of Gaussian process interpolation that we have repeatedly alluded to.

2.3 Conditional Mean and Variance in the Flat Limit

Given Theorem 2.3, the question that arises is how the conditional mean (2.4) and covariance (2.5) behave if the data are m -constant. As the kernel becomes constant for $\lambda = \infty$, the linear systems in the equations which define the conditional moments are singular if $n \geq 2$. The sensible approach is therefore to consider the limits of $\mu_\lambda(x)$ and $P_\lambda(x, y)$ as $\lambda \rightarrow \infty$. This *flat limit* has been extensively studied during the past twenty years in the literature on radial basis function interpolation; see, for instance, Lee et al. (2015) or Barthelmé and Usvich (2021) and the references therein. Flat limits have been recently considered in the context of Gaussian process interpolation by Barthelmé et al. (2022). The conclusion of this body of research is that, under certain assumptions on the kernel and the covariate set geometry, the kernel-dependent term of the conditional mean in (2.6) tends pointwise to (a) a polynomial interpolant if the kernel is infinitely differentiable (Lee et al., 2007, Theorem 3.4) or (b) a polyharmonic spline interpolant if the kernel is finitely differentiable (Song et al., 2012, Theorem 1). Interestingly, if the data are m -constant, we find that it is possible to present a simpler proof that is completely self-contained.

Theorem 2.6 *Suppose that K satisfies Assumption 2.2 and $n \geq 1$. If $Y_m = (c, \dots, c)$ for some $c \in \mathbb{R}$, then*

$$\lim_{\lambda \rightarrow \infty} \mu_\lambda(x) = m(x) + c \quad \text{and} \quad \lim_{\lambda \rightarrow \infty} |P_\lambda(x, y)| = 0$$

for any $x, y \in \mathbb{R}^d$.

Proof See Section 7.5. ■




Figure 3: The conditional mean function for $\lambda \in \{1, 10, 100\}$ when $X = \{-1, 0, 1\} \subset \mathbb{R}$, $Y = (1, 1, 1)$, the prior mean is zero, and K is the Matérn kernel (2.3) with $\sigma = 1$ and $\nu = 3/2$.

Theorem 2.6 is illustrated in Figure 3. Since Φ is assumed continuous in Assumption 2.2, $\mu_\lambda(x)$ and $P_\lambda(x, y)$ are continuous functions of λ for any fixed $x, y \in \mathbb{R}^d$ (which is proved similarly to Lemma 7.1). Therefore Theorem 2.6 justifies writing

$$\mu_{\lambda=\infty}(x) = m(x) + c \quad \text{and} \quad P_{\lambda=\infty}(x, y) = 0. \quad (2.14)$$

2.4 Ill-Posedness

By the classical definition of Hadamard, an inference or estimation problem is *well-posed* if (i) a solution exists, (ii) the solution is unique, and (iii) the solution depends on continuously on the data. If these conditions are not met, the problem is *ill-posed*. In the Bayesian inverse problems literature, where the solution is a posterior measure, the third condition is often strengthened to a requirement that the posterior be locally Lipschitz in the data with respect to the Hellinger distance (Stuart, 2010, Section 4). One of the reasons that make Hellinger distance suitable in our context is that two distributions that are close in Hellinger distance are close also in mean and variance. However, one may also consider other probability metrics (e.g., Latz, 2020). In this section we show that prediction using Gaussian process interpolation is not well-posed in the latter Lipschitz–Hellinger sense if the lengthscale parameter is set using maximum likelihood. Bringing formal notions of well-posedness to bear on Gaussian process interpolation has not, to the best of our knowledge, previously been attempted.

Let Q_1 and Q_2 be two probability distributions on \mathbb{R}^q that are absolutely continuous with respect a reference measure ν on \mathbb{R}^q and let q_1 and q_2 denote their Radon–Nikodym derivatives with respect to ν . The squared Hellinger distance between Q_1 and Q_2 is

$$d_{\text{Hel}}(Q_1, Q_2)^2 = \frac{1}{2} \int_{\mathbb{R}^q} (q_1(x)^{1/2} - q_2(x)^{1/2})^2 d\nu(x). \quad (2.15)$$

The Hellinger distance does not depend on the reference measure ν , which means that for distributions that admit Lebesgue density functions we may set $d\nu(x) = dx$. For univariate Gaussians $Q_1 = N(\mu_1, \Sigma_1)$ and $Q_2 = N(\mu_2, \Sigma_2)$, we have

$$d_{\text{Hel}}(Q_1, Q_2)^2 = 1 - \frac{\sqrt{2}(\Sigma_1 \Sigma_2)^{1/4}}{\sqrt{\Sigma_1 + \Sigma_2}} \exp\left(-\frac{(\mu_1 - \mu_2)^2}{4(\Sigma_1 + \Sigma_2)}\right). \quad (2.16)$$

Let $Q(Y)$ stand for a posterior measure given an observed data vector $Y \in \mathbb{R}^n$. The posterior is said to be well-posed if for every $\varepsilon > 0$ there exists $L > 0$ such that

$$d_{\text{Hel}}(Q(Y), Q(Y')) \leq L \|Y - Y'\| \quad (2.17)$$

for any data vectors $Y, Y' \in \mathbb{R}^n$ for which $\|Y - Y'\| \leq \varepsilon$.

Let us consider the Gaussian process predictive distribution at some unobserved point $x_0 \notin X$ as the posterior and set

$$Q_{\text{GP}}(Y) = N(\mu_{\lambda_{\text{ML}}(Y)}(x_0), P_{\lambda_{\text{ML}}(Y)}(x_0)^2), \quad (2.18)$$

where we use $\lambda_{\text{ML}}(Y)$ to denote that a maximum likelihood estimate depends on the data Y . We may assume that $\lambda_{\text{ML}}(Y)$ (or, if the modified log-likelihood function has multiple global minimum points, the largest of these) is a continuous function of the data, for otherwise predictions would not be continuous in the data, let alone Lipschitz. Let $\varepsilon > 0$ and let $(Y_k)_{k=1}^\infty$ and $(Y'_k)_{k=1}^\infty$ be two data sequences which satisfy $\|Y_k - Y'_k\| \leq \varepsilon$ for every $k \in \mathbb{N}$ and which converge to an m -constant data set:

$$\lim_{k \rightarrow \infty} Y_k - m(X) = \lim_{k \rightarrow \infty} Y'_k - m(X) = (c, \dots, c) \in \mathbb{R}^n$$

for some $c \in \mathbb{R}$. By Theorems 2.3 and 2.6 and the assumed continuity of $\lambda_{\text{ML}}(Y)$ in the data, these sequences can be selected such that

$$\Sigma_k := P_{\lambda_{\text{ML}}(Y_k)}(x_0)^2 = C_1 e^{-k} \quad \text{and} \quad \Sigma'_k := P_{\lambda_{\text{ML}}(Y'_k)}(x_0)^2 = C_2 k^{-1}$$

for some positive constants C_1 and C_2 . Since $e^{-x} \leq 1$ for all $x \geq 0$, we get from (2.16) and (2.18) that

$$\begin{aligned} d_{\text{Hel}}(Q_{\text{GP}}(Y_k), Q_{\text{GP}}(Y'_k))^2 &\geq 1 - \frac{\sqrt{2}(\Sigma_k \Sigma'_k)^{1/4}}{\sqrt{\Sigma_k + \Sigma'_k}} = 1 - \frac{\sqrt{2}(C_1 C_2)^{1/4} k^{-1/4} e^{-k/4}}{\sqrt{C_1 e^{-k} + C_2 k^{-1}}} \\ &\geq 1 - \sqrt{2} C_1^{1/4} C_2^{-1/4} k^{1/4} e^{-k/4}, \end{aligned}$$

where the second term tends to zero as $k \rightarrow \infty$. Therefore

$$d_{\text{Hel}}(Q_{\text{GP}}(Y_k), Q_{\text{GP}}(Y'_k)) \rightarrow 1 \quad \text{as} \quad k \rightarrow \infty$$

even though $\|Y_k - Y'_k\| \rightarrow 0$ as $k \rightarrow \infty$. This shows that the Lipschitz condition (2.17) fails to hold when the data domain is

$$\mathcal{R}^n = \{Y \in \mathbb{R}^n : Y \text{ is not } m\text{-constant}\} \subset \mathbb{R}^n,$$

the set of data sets that are not m -constant. That is, we have shown that the mapping $Q_{\text{GP}}: \mathcal{R}^n \rightarrow \mathcal{P}$ defined in (2.18) is not Lipschitz, where \mathcal{P} is the space of probability distributions on \mathbb{R} equipped with the Hellinger distance.

The above derivation is a consequence of the fact that, from Theorem 2.3 and the interpretation in (2.14),

$$P_{\lambda_{\text{ML}}(Y)}(x_0)^2 = 0$$

if the data Y are m -constant. Then for any data Y' which are not m -constant we compute from (2.16) that

$$d_{\text{Hel}}(Q_{\text{GP}}(Y), Q_{\text{GP}}(Y')) = 1,$$

which means that the predictive distribution is not continuous at any data which are m -constant. Note that this is a purely formal computation because (2.15) and (2.16) are valid only for measures which are absolutely continuous with respect a common reference measure, which is not the case with the degenerate Gaussian predictive distribution Y that arises from m -constant data and the non-degenerate Gaussian Y' . By observing that the argument above uses Theorem 2.3 only to guarantee the existence of data Y for which $\lambda_{\text{ML}}(Y) = \infty$, we may formulate the following generic ill-posedness theorem.

Theorem 2.7 (Ill-posedness) *Suppose that K satisfies Assumption 2.2 and $n \geq 1$. Let $\bar{\lambda}: \mathbb{R}^n \rightarrow [0, \infty]$ be any estimator of λ and define $\mathcal{R}^n = \{Y \in \mathbb{R}^n : \bar{\lambda}(Y) < \infty\} \subset \mathbb{R}^n$. If there are data $Y \in \mathbb{R}^n$ such that $\bar{\lambda}(Y) = \infty$ (i.e., $\mathcal{R}^n \neq \mathbb{R}^n$), then Gaussian process interpolation is ill-posed, in the sense that the predictive distribution mapping $Q_{\text{GP}}: \mathcal{R}^n \rightarrow \mathcal{P}$ defined in (2.18) is not Lipschitz for any $x_0 \notin X$.*

The main message of Theorem 2.7 is that a lengthscale estimator, whatever it might be, must be finite for any data in order for Gaussian process interpolation to be well-posed.

3. What Does Not Help

The Gaussian process model in Theorem 2.3 is fairly simple, having a fixed prior mean function and a single estimated hyperparameter. One might hope that additional modelling choices—or the use of an altogether different parameter estimation method—would yield a well-posed Gaussian process model. In this section we show that this is not to be for several common approaches. Each theorem in this section shows that an estimator of λ is infinite if the data are m -constant, so that Theorem 2.7 consequently establishes that Gaussian process interpolation is ill-posed.

3.1 Cross-Validation

Leave-one-out cross-validation is a popular alternative to maximum likelihood estimation that has been shown to confer robustness when the Gaussian process model is misspecified (Bachoc, 2013). In Gaussian process interpolation the objective function that is typically used is

$$\ell_{\text{CV}}(\theta | Y) = \sum_{k=1}^n \left[\left(\frac{y_k - \mu_{\theta, n, k}(x_k)}{P_{\theta, n, k}(x_k)} \right)^2 + \log[P_{\theta, n, k}(x_k)^2] \right], \quad (3.1)$$

where $\mu_{\theta, n, k} = m + s_{\theta, n, k}$ and $P_{\theta, n, k}$ denote the Gaussian process conditional mean and standard deviation functions in (2.4)–(2.6) based on data at the points $X \setminus \{x_k\}$; see, for example, Section 4.2 in Currin et al. (1988) or Section 5.4.1 in Rasmussen and Williams (2006). Subscripts are again used to make explicit the dependency of these functions on the kernel parameters θ . Any corresponding parameter estimate θ_{CV} satisfies

$$\theta_{\text{CV}} \in \arg \min_{\theta \in \Theta} \ell_{\text{CV}}(\theta | Y).$$

The cross-validation objective function (3.1) is obtained by summing negative predictive log-probabilities of y_k given data at $X \setminus \{x_k\}$ and discarding terms which do not depend on θ . Unfortunately, leave-one-out cross-validation also fails to be well-posed. This may not be surprising given that there is a close connection between maximum likelihood estimation and cross-validation (Fong and Holmes, 2020).

Theorem 3.1 (Cross-validation) *Suppose that K satisfies Assumption 2.2 and $n \geq 2$. If the data Y are m -constant, then*

$$\lim_{\lambda \rightarrow \infty} \ell_{\text{CV}}(\lambda | Y) = -\infty \quad \text{and} \quad \lambda_{\text{CV}} = \infty. \quad (3.2)$$

Proof See Section 7.6. Despite the ostensibly different forms of the objective functions (2.8) and (3.1), the proof is, in consequence of Proposition 7.5, in essence all but identical to the proof of Equation (2.11) in Theorem 2.3. ■

We believe that $\ell_{\text{CV}}(\lambda | Y)$ and λ_{CV} satisfy a version of Equation (2.12) if the data are not m -constant but have been unable to furnish a proof; see Remark 7.8.

Remark 3.2 *A non-probabilistic alternative to (3.1) is to simply minimise the sum of squared leave-one-out errors (e.g., Rippen, 1999):*

$$\ell_{\text{CV}(2)}(\theta | Y) = \sum_{k=1}^n (y_k - \mu_{\theta, n, k}(x_k))^2 \geq 0 \quad \text{and} \quad \theta_{\text{CV}(2)} \in \arg \min_{\theta \in \Theta} \ell_{\text{CV}(2)}(\theta | Y).$$

Consider estimating the lengthscale parameter λ using this procedure. If the data Y are m -constant such that $Y_m = Y - m(X) = (c, \dots, c)$, it follows from Theorem 2.6 that $\mu_{\lambda, n, k}$ tends pointwise to $m + c$ as $\lambda \rightarrow \infty$ if K satisfies Assumption 2.2. Therefore

$$\lim_{\lambda \rightarrow \infty} \ell_{\text{CV}(2)}(\lambda | Y) = \lim_{\lambda \rightarrow \infty} \sum_{k=1}^n (m(x_k) + c - \mu_{\lambda, n, k}(x_k))^2 = 0,$$

from which it follows that $\lambda_{\text{CV}(2)} = \infty$, or at least that $\ell_{\text{CV}(2)}(\lambda | Y)$ has one of its minima at infinity. Therefore also this procedure is ill-posed.

3.2 Unknown Parametric Prior Mean

So far we have considered a setting where the prior mean function m is known and fixed. But in methods such as universal kriging the mean is assumed to be an unknown element of the linear span of a finite number of basis functions, typically polynomials, and its coefficients are estimated from the data. See, for example, O'Hagan (1978) or Chapters 3 and 4 in Santner et al. (2003).

Let the basis functions be $\varphi_1, \dots, \varphi_q$ for $q \leq n$ and define the matrix $V(X) \in \mathbb{R}^{n \times q}$ with elements

$$(V(X))_{i,j} = \varphi_j(x_i).$$

Suppose that the mean function is $m = \sum_{j=1}^q \beta_j \varphi_j$ for unknown coefficients $\beta = (\beta_1, \dots, \beta_q)$ which we wish to estimate using maximum likelihood. The full modified log-likelihood

function for both the kernel parameters θ and the coefficients β is obtained by inserting the parametric prior mean in (2.8):

$$\ell(\theta, \beta | Y) = (Y - V(X)\beta)^\top K_\theta(X, X)^{-1}(Y - V(X)\beta) + \log \det K_\theta(X, X). \quad (3.3)$$

Any maximum likelihood estimates satisfy

$$\{\theta_{\text{ML}}, \beta_{\text{ML}}\} \in \arg \min_{\theta \in \Theta, \beta \in \mathbb{R}^q} \ell(\theta, \beta | Y).$$

The natural generalisation of Definition 2.1 to this setting is that there exist some coefficients for which the data are m -constant. That is, that there exist constants c and $\beta^* = (\beta_1^*, \dots, \beta_q^*)$ (which need not be unique) such that

$$y_i - \sum_{j=1}^q \beta_j^* \varphi_j(x_i) = c \quad \text{for every } i = 1, \dots, n. \quad (3.4)$$

The next theorem shows that the maximum likelihood estimate of the lengthscale parameter λ is badly behaved if the data satisfy the above assumption.

Theorem 3.3 *Suppose that K satisfies Assumption 2.2 and $n \geq 2$. If the data Y satisfy (3.4), then*

$$\lambda_{\text{ML}} = \infty \quad \text{if} \quad \{\lambda_{\text{ML}}, \beta_{\text{ML}}\} \in \arg \min_{\lambda > 0, \beta \in \mathbb{R}^q} \ell(\lambda, \beta | Y).$$

Proof See Section 7.6. The proof is similar to that of Equation (2.11) in Theorem 2.3. ■

Remark 3.4 *If the matrix $V(X)$ has full rank, one can compute that*

$$\beta_{\text{ML}} = [V(X)^\top K_\theta(X, X)^{-1} V(X)]^{-1} V(X)^\top K_\theta(X, X)^{-1} Y$$

for any fixed θ . If $q = n$, the matrix $V(X)$ is square and non-singular so that the above maximum likelihood estimate simplifies to $\beta_{\text{ML}} = V(X)^{-1} Y$. Inserting this to (3.3) eliminates the data-fit term and we are left with $\ell(\theta, \beta_{\text{ML}} | Y) = \log \det K_\theta(X, X)$. The maximum likelihood estimate of θ is therefore obtained by minimising complexity of the model. If $\theta = \lambda$, this naturally leads to $\lambda_{\text{ML}} = \infty$ since this is the only value of the lengthscale parameter for which the covariance matrix becomes singular. The interpretation of this phenomenon is that maximum likelihood estimation picks the simplest possible model if the data are fully explained by the prior mean.

3.3 Simultaneous Estimation of the Scaling Parameter

One typically estimates the lengthscale parameter simultaneously with the *scale* or *magnitude* parameter $\sigma > 0$. Suppose that $\theta = \{\lambda, \sigma\}$ and the covariance kernel is parametrised as $K_\theta(x, y) = \sigma^2 K_\lambda(x, y)$. Then

$$\ell(\lambda, \sigma | Y) = \frac{1}{\sigma^2} Y_m^\top K_\lambda(X, X)^{-1} Y_m + \log \det K_\lambda(X, X) + n \log \sigma^2$$

and maximum likelihood estimates satisfy

$$\{\lambda_{\text{ML}}, \sigma_{\text{ML}}\} \in \arg \min_{\sigma, \lambda > 0} \ell(\lambda, \sigma | Y). \quad (3.5)$$

Unfortunately, the behaviour of the maximum likelihood estimate λ_{ML} is identical to the case in Theorem 2.3 where σ is held fixed. Let $\sigma_{\text{ML}}(\lambda)$ denote the maximum likelihood estimate of σ for a fixed $\lambda > 0$.

Theorem 3.5 (Simultaneous estimation) *Suppose that K satisfies Assumption 2.2 and $n \geq 2$. Consider the maximum likelihood estimates in (3.5). If the data Y are m -constant and $Y_m \neq 0$, then*

$$\lim_{\lambda \rightarrow \infty} \ell(\lambda, \sigma_{\text{ML}}(\lambda) | Y) = -\infty \quad \text{and} \quad \lambda_{\text{ML}} = \infty.$$

Proof See Section 7.6. The proof does not fundamentally differ from that of Theorem 2.3. ■

If it happens that $Y_m = 0$, the modified log-likelihood function is simply

$$\ell(\lambda, \sigma | Y) = \log \det K_\lambda(X, X) + n \log \sigma^2.$$

Then $\ell(\lambda, \sigma | Y) = -\infty$ if and only if $\lambda \rightarrow \infty$ or $\sigma \rightarrow 0$. Either of these cases results zero conditional variance and degenerate predictive distributions.

4. Regularisation

This section discusses two types of regularisation that can be used to ensure the well-posedness of Gaussian process regression or finiteness of a lengthscale estimator. However, these approaches may induce unwanted side effects and are, to some extent, arbitrary.

4.1 Regularisation via Observation Noise

Let $\delta > 0$ be a *regularisation parameter* (alternatively, *smoothing parameter*, *nugget* or *jitter*). Denote the $n \times n$ identity matrix with I_n . The regularised versions of the Gaussian process conditional mean and covariance in (2.4) and (2.5) are

$$\begin{aligned} \mu^\delta(x) &= m(x) + K(x, X)^\top (K(X, X) + \delta^2 I_n)^{-1} Y_m, \\ P^\delta(x, y)^2 &= K(x, y) - K(x, X)^\top (K(X, X) + \delta^2 I_n)^{-1} K(X, y) \end{aligned}$$

and that of the modified log-likelihood function is

$$\ell^\delta(\lambda | Y) = Y_m^\top (K_\lambda(X, X) + \delta^2 I_n)^{-1} Y_m + \log \det(K_\lambda(X, X) + \delta^2 I_n). \quad (4.1)$$

In this section every quantity which is superscripted with δ stands for the corresponding quantity defined in Section 2 but with $K(X, X)$ replaced by $K(X, X) + \delta^2 I_n$. Although the above expressions arise from assuming that the data are corrupted by additive and independent zero-mean Gaussian noise terms with variances δ^2 , regularisation is often used purely out of convenience (that is, even when one does not believe that the data are noisy)




Figure 4: The regularised maximum likelihood estimate $\lambda_{\text{ML}}^\delta$ in (4.2) as a function of the regularisation parameter δ for four different data sets $Y \in \mathbb{R}^3$.

as it improves the condition number of the covariance matrix that needs to be inverted (e.g., Ranjan et al., 2011; Adrianakis and Challenor, 2012).

Because $(K(X, X) + \delta^2 I_n)^{-1} < K(X, X)^{-1}$ in the Loewner ordering of positive-semidefinite matrices, we easily derive that the regularised conditional variance is everywhere positive:

$$P^\delta(x)^2 = P^\delta(x, x)^2 > P(x)^2 \geq 0 \quad \text{for every } x \in \mathbb{R}^d.$$

Consider then the lengthscale parameter λ and the corresponding kernel K_λ and assume for simplicity that $K_\lambda(x, y) \rightarrow 1$ as $\lambda \rightarrow \infty$ for all $x, y \in \Omega$. Then it is straightforward to use the Sherman–Morrison formula to compute that

$$\lim_{\lambda \rightarrow \infty} P_\lambda^\delta(x)^2 = \lim_{\lambda \rightarrow \infty} [K_\lambda(x, y) - K_\lambda(x, X)]^\top (K_\lambda(X, X) + \delta^2 I_n)^{-1} K_\lambda(x, X) = \frac{1}{1 + n\delta^{-2}}$$

and

$$\lim_{\lambda \rightarrow \infty} \mu_\lambda^\delta(x) = m(x) + \frac{\delta^{-2}}{1 + n\delta^{-2}} \sum_{i=1}^n (y_i - m(x_i))$$

for any $x \in \mathbb{R}^d$ and $Y_m \in \mathbb{R}^n$. These estimates and computations establish that under regularisation no finite-dimensional distribution of the conditional Gaussian process can tend to a degenerate Gaussian as λ varies. It follows that regularised Gaussian process interpolation is well-posed in the sense discussed in Section 2.4 as long as the estimator of λ is continuous in the data. However, the price one pays for vanquishing ill-conditioning is that the conditional mean no longer interpolates the data, which may be undesirable in applications where the data are truly noiseless, and that an additional degree of freedom is introduced. Favorable convergence rates of Gaussian process interpolation (asymptotically as $n \rightarrow \infty$) are also lost under regularisation unless one has the regularisation parameter tend to zero with an appropriate rate (Wendland and Rieger, 2005, Section 3).

Even though maximum likelihood estimation of λ cannot cause ill-posedness in the regularised setting (assuming the estimator is continuous in the data), there remains the

interesting question of how $\lambda_{\text{ML}}^\delta$, as computed by minimising the regularised modified log-likelihood function in (4.1), behaves if the data are m -constant. Our attempts at proving any properties of $\lambda_{\text{ML}}^\delta$ have been unsuccessful, and therefore here we limit ourselves to a simple numerical investigation, the results of which are depicted in Figure 4. This figure plots

$$\lambda_{\text{ML}}^\delta = \arg \min_{\lambda > 0} \ell^\delta(\lambda | Y) \quad (4.2)$$

as a function of $\delta > 0$ for four different data vectors $Y \in \mathbb{R}^3$ (we set $m \equiv 0$) when $X = \{1, 1.2, 2.0\} \subset \mathbb{R}$ and K is the Matérn kernel in (2.3) with parameters $\sigma = 1$ and $\nu = 3/2$. Minimisation was performed using grid search. It appears that for each data set the maximum likelihood estimate $\lambda_{\text{ML}}^\delta$ has a singularity at a certain value δ_∞ of δ and that δ_∞ is smaller when the data Y are closer to being constant. This suggests that, for a sufficiently large δ , maximum likelihood estimation always reverts to the simplest possible model, that obtained with $\lambda = \infty$. One way to interpret this observation is that any given data set could have been plausibly generated by a constant function if the data set is assumed to be corrupted by Gaussian noise with sufficiently large variance (i.e., if δ^2 is sufficiently large).

4.2 Regularisation via Lengthscale Hyperprior

Suppose that a prior with a density function p is placed on the parameters θ . Then the posterior for θ has the log-density

$$\log p(\theta | Y) = -\frac{1}{2}\ell(\theta | Y) + \log p(\theta) + \text{constant},$$

where $\ell(\theta | Y)$ is the modified log-likelihood function in (2.8). Any maximiser θ_{MAP} of the posterior density is called a *maximum a posteriori* (MAP) estimate of θ . Equivalently,

$$\theta_{\text{MAP}} \in \arg \min_{\theta \in \Theta} \ell_{\text{MAP}}(\theta | Y) \quad \text{for} \quad \ell_{\text{MAP}}(\theta | Y) = \frac{1}{2}\ell(\theta | Y) - \log p(\theta). \quad (4.3)$$

The following theorem shows that assigning a non-heavy-tailed hyperprior on the lengthscale λ ensures that its MAP estimate is finite regardless of the data.

Theorem 4.1 *Suppose that K satisfies Assumption 2.2, $n \geq 1$, and*

$$\lambda^{(\alpha-d/2)n} p(\lambda) \rightarrow 0 \quad \text{as} \quad \lambda \rightarrow \infty, \quad (4.4)$$

where $\alpha > d/2$ is the constant in Assumption 2.2. If $Y \in \mathbb{R}^n$ are any data, then

$$\lim_{\lambda \rightarrow \infty} \ell_{\text{MAP}}(\lambda | Y) = \infty \quad \text{and} \quad \lambda_{\text{MAP}} < \infty.$$

Proof See Section 7.7. The essence of the proof is that the assumption on tail decay of p ensures that $-\log p(\lambda)$ dominates $\ell_{\text{MAP}}(\lambda | Y)$ as $\lambda \rightarrow \infty$. ■

Because, for fixed Y , the choice of the hyperprior p completely determines the estimator, MAP estimation is rather arbitrary and not well-suited to deployment in general-purpose software.

5. Generalisations and Extensions

This section discusses some generalisations of Theorem 2.3 for (a) general linear data, (b) lengthscale estimation for product kernels, and (c) kernels which do not satisfy the Fourier decay assumption in (2.9).

5.1 Linear Information and General Kernels

In this section we generalise part of Theorem 2.3 by replacing the domain \mathbb{R}^d with an arbitrary vector space Ω , using a more liberal definition of a lengthscale parameter, and considering general linear data, such as derivative evaluations. Though somewhat technical, the assumptions that this generalisation requires can be verified in some settings of practical interest.

Let Ω be a vector space and $K_\theta: \Omega \times \Omega \rightarrow \mathbb{R}$ a positive-definite kernel on Ω for each $\theta \in \Theta$. We use $H(K_\theta, \Omega)$ to denote the RKHS of K_θ on Ω ; see Section 7.1 for details. Let $F(\Omega)$ be a set of real-valued functions defined on Ω which contains $H(K_\theta, \Omega)$ for every $\theta \in \Theta$ as well as all constant functions and the prior mean function m . Let $\mathcal{L} = \{L_1, \dots, L_n\}$ be a collection of $n \geq 2$ non-trivial (i.e., none of them is the zero functional) linear *information functionals* defined on $F(\Omega)$. These functionals are assumed to be linearly independent and bounded on $H(K_\theta, \Omega)$ for every $\theta \in \Theta$, in that there is $C_\theta > 0$ such that $|L_i(f)| \leq C_\theta \|f\|_{H(K_\theta, \Omega)}$ for every $f \in H(K_\theta, \Omega)$. We assume that the data

$$Y_{\mathcal{L}} = (L_1 f, \dots, L_n f) \in \mathbb{R}^n \quad (5.1)$$

consist of applications of the information functionals to an underlying (and unknown) data-generating function $f \in F(\Omega)$. The setting considered earlier in this article is recovered by selecting the point evaluation functionals defined as $L_i f = f(x_i)$. Partial derivative data, defined by $L_i f = \partial_{p_i} f(x_i)$ for some $x_i \in \Omega$ and $p_i \geq 0$, also occurs commonly in Gaussian process applications (e.g., Solak et al., 2002). If necessary to avoid ambiguity, we use superscripts to denote the argument with respect to which an information functional is to be applied.

Let $\mathcal{L}(m) = (L_i m)_{i=1}^n \in \mathbb{R}^n$ and set $Y_{\mathcal{L}, m} = Y_{\mathcal{L}} - \mathcal{L}(m) \in \mathbb{R}^n$. Also set

$$K_\theta(\mathcal{L}, \mathcal{L}) = (L_i^x L_j^y K_\theta(x, y))_{i, j=1}^n \in \mathbb{R}^{n \times n} \quad \text{and} \quad K_\theta(\mathcal{L}, x) = (L_i K_\theta(\cdot, x))_{i=1}^n \in \mathbb{R}^n,$$

which are well-defined by the assumption that the information functionals are bounded on $H(K_\theta, \Omega)$. The vector $K_\theta(x, \mathcal{L})$ is defined analogously to $K_\theta(\mathcal{L}, x)$ but with the information functionals applied to the second argument. The Gaussian process $f_{\text{GP}} \sim \text{GP}(m, K_\theta)$ conditioned on the general linear data $Y_{\mathcal{L}}$ in (5.1) has the mean

$$\mathbb{E}[f_{\text{GP}}(x) | Y_{\mathcal{L}}] = Y_{\mathcal{L}, m}^\top K_\theta(\mathcal{L}, \mathcal{L})^{-1} K_\theta(\mathcal{L}, x),$$

and covariance

$$\text{Cov}[f_{\text{GP}}(x), f_{\text{GP}}(y) | Y_{\mathcal{L}}] = K_\theta(x, y) - K_\theta(x, \mathcal{L})^\top K_\theta(\mathcal{L}, \mathcal{L})^{-1} K_\theta(\mathcal{L}, y).$$

The modified log-likelihood function is

$$\ell(\theta | Y_{\mathcal{L}}) = Y_{\mathcal{L}, m}^\top K_\theta(\mathcal{L}, \mathcal{L})^{-1} Y_{\mathcal{L}, m} + \log \det K_\theta(\mathcal{L}, \mathcal{L}). \quad (5.2)$$

Let $g: (0, \infty) \rightarrow (0, \infty)$ be a continuous function such that (a) $\limsup_{\lambda \rightarrow 0} g(\lambda) < \infty$, (b) $\lim_{\lambda \rightarrow \infty} g(\lambda) = \infty$, and (c) $g(\lambda_1) = 1$ for some $\lambda_1 > 0$. Define

$$K_\lambda(x, y) = K\left(\frac{x}{g(\lambda)}, \frac{y}{g(\lambda)}\right) \quad \text{for any } \lambda > 0 \text{ and } x, y \in \Omega. \quad (5.3)$$

The parameter λ can be interpreted as a generalised version of the lengthscale parameter of a stationary kernel in (1.2). Stationary kernels are recovered from (5.3) by setting $g(\lambda) = \lambda$ and $K(x, y) = \Phi(x - y)$. The exponential kernel $K_\lambda(x, y) = \exp(xy/\lambda)$ is a simple non-stationary example that is occasionally used. We are interested in maximum likelihood estimates $\lambda_{\text{ML}(\mathcal{L})}$ of λ , any of which satisfies

$$\lambda_{\text{ML}(\mathcal{L})} \in \arg \min_{\lambda > 0} \ell(\lambda | Y_{\mathcal{L}}) = \arg \min_{\lambda > 0} \{Y_{\mathcal{L},m}^\top K_\lambda(\mathcal{L}, \mathcal{L})^{-1} Y_{\mathcal{L},m} + \log \det K_\lambda(\mathcal{L}, \mathcal{L})\}.$$

In this setting we say that the data $Y_{\mathcal{L}}$ are m -constant if there is a constant function $c: \Omega \rightarrow \mathbb{R}$ such that

$$Y_{\mathcal{L},m} = (L_1 c, \dots, L_n c). \quad (5.4)$$

To prove a generalisation of Theorem 2.3 we need the following technical assumption.

Assumption 5.1 *Every element of the matrix $K_\lambda(\mathcal{L}, \mathcal{L})$ is a continuous function of $\lambda > 0$ and*

$$(a) \quad \liminf_{\lambda \rightarrow 0} e_{\min}(K_\lambda(\mathcal{L}, \mathcal{L})) > 0 \quad \text{and} \quad (b) \quad \lim_{\lambda \rightarrow \infty} \det K_\lambda(\mathcal{L}, \mathcal{L}) = 0,$$

where $e_{\min}(A)$ denotes the smallest eigenvalue of a matrix A .

Theorem 5.2 *Suppose that Assumption 5.1 holds and let Ω_b be any convex subset of Ω such that (a) $0 \in \Omega_b$ and (b) $Lf_1 = Lf_2$ for every $L \in \mathcal{L}$ and all $f_1, f_2 \in F(\Omega)$ such that $f_1 = f_2$ on Ω_b . If the data $Y_{\mathcal{L}}$ are m -constant and constant functions are contained in $H(K, \Omega_b)$, then*

$$\lim_{\lambda \rightarrow \infty} \ell(\lambda | Y_{\mathcal{L}}) = -\infty \quad \text{and} \quad \lambda_{\text{ML}(\mathcal{L})} = \infty.$$

Proof See Section 7.8. The proof is in essence identical to that of (2.11), but more technical. ■

Assumption 5.1 and other assumptions in Theorem 5.2 may be verified on case-by-case basis. For example, let $\Omega = \mathbb{R}$ and

$$K_\lambda(x, y) = \left(1 + \frac{\sqrt{3}|x-y|}{\lambda}\right) \exp\left(-\frac{\sqrt{3}|x-y|}{\lambda}\right),$$

which is the Matérn kernel in (2.3) with smoothness $\nu = 3/2$. Suppose that $L_1 f = f'(x_1)$ and $L_i f = f(x_i)$ for $i = 2, \dots, n$ and distinct x_2, \dots, x_n . One may easily find a convex $\Omega_b \subset \mathbb{R}$ such that $0, x_1, x_2, \dots, x_n \in \Omega_b$ (e.g., a ball centered at the origin with radius that exceeds the maximal norm of the points). Because this set can be taken to be bounded, it follows

from (2.10) and the results reviewed in Section 7.3 that constant functions are contained in $H(K, \Omega_b)$. To verify Assumption 5.1, observe that the generalised covariance matrix is

$$K_\lambda(\mathcal{L}, \mathcal{L}) = \begin{pmatrix} a_\lambda & b_\lambda^\top \\ b_\lambda & K_\lambda(X', X') \end{pmatrix}, \quad (5.5)$$

where $K_\lambda(X', X') \in \mathbb{R}^{(n-1) \times (n-1)}$ is the regular covariance matrix for the points x_2, \dots, x_n ,

$$a_\lambda = \frac{\partial^2}{\partial x \partial y} K_\lambda(x, y) \Big|_{\substack{x=x_1 \\ y=x_1}} = \frac{3}{\lambda^2},$$

and

$$(b_\lambda)_{i-1} = \frac{\partial}{\partial x} K_\lambda(x, x_i) \Big|_{x=x_1} = -\frac{3}{\lambda^2} (x_1 - x_i) \exp\left(-\frac{\sqrt{3}|x_1 - x_i|}{\lambda}\right)$$

for $i = 2, \dots, n$. The entries of b_λ , as well as a_λ , are continuous in λ on $(0, \infty)$. We compute $\lim_{\lambda \rightarrow 0} a_\lambda = \infty$ and $\lim_{\lambda \rightarrow \infty} a_\lambda = 0$, as well as $\lim_{\lambda \rightarrow 0} (b_\lambda)_{i-1} = 0$ and $\lim_{\lambda \rightarrow \infty} (b_\lambda)_{i-1} = 0$ for each $i = 2, \dots, n$. Therefore $\lim_{\lambda \rightarrow 0} K_\lambda(\mathcal{L}, \mathcal{L}) = \text{diag}(\infty, 1, \dots, 1)$ and $\lim_{\lambda \rightarrow \infty} K_\lambda(\mathcal{L}, \mathcal{L})$ is singular because its first row is zero. Assumption 5.1 thus holds and it follows from Theorem 5.2 that $\lambda_{\text{ML}(\mathcal{L})} = \infty$ if $Y_{\mathcal{L}, m} = (0, c, \dots, c)$ for some $c \in \mathbb{R}$.

5.2 Product Kernels and Multiple Lengthscales

Here we consider a setting where the covariance kernel is a product of stationary kernels equipped with dimensionwise lengthscale parameters. That is, $\theta = \{\lambda_1, \dots, \lambda_d\}$ and the kernel has the product form

$$K_\theta(x, y) = \prod_{i=1}^d K_{i, \lambda_i}(x_i, y_i) = \prod_{i=1}^d \Phi_i\left(\frac{x_i - y_i}{\lambda_i}\right), \quad (5.6)$$

where $x = (x_1, \dots, x_d) \in \mathbb{R}^d$, $y = (y_1, \dots, y_d) \in \mathbb{R}^d$ and $K_i(x_i, y_i) = \Phi_i(x_i - y_i)$ are stationary kernels on \mathbb{R} parametrised by positive lengthscale parameters λ_i . Recall the definition of Matérn kernel from (2.3). The product Matérn kernel

$$K_\theta(x, y) = \sigma^2 \prod_{i=1}^d \left[\frac{2^{1-\nu_i}}{\Gamma(\nu_i)} \left(\frac{\sqrt{2\nu_i} |x_i - y_i|}{\lambda_i} \right)^{\nu_i} \mathcal{K}_{\nu_i} \left(\frac{\sqrt{2\nu_i} |x_i - y_i|}{\lambda_i} \right) \right],$$

where $\nu_i > 0$, is a commonly used product kernel of the form (5.6).

For simplicity, let us consider maximum likelihood estimation of only one of the d lengthscale parameters. For $p \in \{1, \dots, d\}$, we are interested in the behaviour of

$$\lambda_{p, \text{ML}} = \arg \min_{\lambda_p > 0} \ell(\lambda_1, \dots, \lambda_d | Y), \quad (5.7)$$

where $\ell(\lambda_1, \dots, \lambda_d | Y)$ is the modified log-likelihood function in (2.8) for $\theta = \{\lambda_1, \dots, \lambda_d\}$ and the product kernel in (5.6) and λ_i for $i \neq p$ are fixed. For product covariates X of the form

$$X = X_1 \times \dots \times X_d, \quad \text{where } X_i = \{x_{i,1}, \dots, x_{i,n_i}\} \subset \mathbb{R}, \quad (5.8)$$

we say that the associated data Y are m -constant *along dimension p* if

$$y_{i_1, \dots, i_d} - m(x_{1, i_1}, \dots, x_{d, i_d}) \quad \text{does not depend on} \quad i_p = 1, \dots, n_p,$$

where the datum y_{i_1, \dots, i_d} is associated with the covariate $(x_{1, i_1}, \dots, x_{d, i_d}) \in X$. For example, the data

$$Y = (y_{1,1}, y_{1,2}, y_{1,3}, y_{2,1}, y_{2,2}, y_{2,3}) = (0, 1, 2, 0, 1, 2)$$

are constant along dimension $p = 1$ for the product design

$$X = \{x_{1,1}, x_{1,2}\} \times \{x_{2,1}, x_{2,2}, x_{2,3}\} = \{(x_{1,1}, x_{2,1}), (x_{1,1}, x_{2,2}), (x_{1,1}, x_{2,3}), \\ (x_{1,2}, x_{2,1}), (x_{1,2}, x_{2,2}), (x_{1,2}, x_{2,3})\}$$

in \mathbb{R}^2 and the prior mean $m \equiv 0$.

Theorem 5.3 (Estimation of multiple lengthscales) *Consider the product kernel in (5.6) and suppose that the stationary kernels K_1, \dots, K_d on \mathbb{R} satisfy Assumption 2.2. If X has the product form (5.8) with $n_i \geq 1$ for each $i = 1, \dots, d$ and the data Y are m -constant along dimension $p \in \{1, \dots, d\}$, then*

$$\lim_{\lambda_p \rightarrow \infty} \ell(\lambda_1, \dots, \lambda_d | Y) = -\infty \quad \text{and} \quad \lambda_{p, \text{ML}} = \infty,$$

where $\lambda_{p, \text{ML}}$ is the maximum likelihood estimate of the p th lengthscale parameter in (5.7).

Proof See Section 7.9. The product form of the kernel and the covariates allow one to write the full covariance matrix $K_\theta(X, X)$ as a Kronecker product of $K_{i, \lambda}(X_i, X_i)$. One may then utilise the properties of Kronecker products and subsequently follow the proof of Theorem 2.3. ■

Note that the constants C_1, C_2 and α in Assumption 2.2 may differ from one constituent kernel K_i to another. Theorem 5.3 provides some theoretical justification for the use of maximum likelihood estimation of lengthscales as an *automatic relevance determination* method (Rasmussen and Williams, 2006, Section 5.1). When the data are independent of the p th input dimension, the lengthscales for this dimension is set to infinite and the dimension is effectively ignored.

5.3 Infinitely Smooth Stationary Kernels

Commonly used infinitely smooth stationary kernels, such as the Gaussian and the inverse quadratic (or Cauchy) defined by

$$\Phi(z) = \exp(-\|z\|^2) \quad \text{and} \quad \Phi(z) = \frac{1}{1 + \|z\|^2}, \quad (5.9)$$

respectively, do not satisfy Assumption 2.2 because their Fourier transforms decay (at least) exponentially. The exponential decay of their Fourier transforms implies that these kernels are analytic. The purpose of Assumption 2.2 is to guarantee that constant functions are contained in the RKHS of K on a bounded set, a result which in turn can be exploited to

prove that the data-fit term $Y_m^\top K_\lambda(X, X)^{-1} Y_m$ is a bounded function of λ whenever the data are m -constant (see Lemma 7.2). However, it is known that the RKHSs of analytic stationary kernels, such as those in (5.9), do not contain constant functions (Steinwart et al., 2006; Sun and Zhou, 2008; Minh, 2010; Dette and Zhigljavsky, 2021). But this does have to mean that the data-fit term explodes as $\lambda \rightarrow \infty$.

Increasing the lengthscale parameter is equivalent to *coalescence* of the points to the origin. That is, using the kernel K_λ and points X is equivalent to using the kernel K and the scaled points $X_\lambda = \{x_1/\lambda, \dots, x_n/\lambda\}$, each of which tends to the origin as $\lambda \rightarrow \infty$. Suppose for simplicity that $d = 1$. When the points coalesce, one's data effectively comprises the value at the origin of the data-generating function and its successive derivatives up to order $n - 1$. We refer to Section 11 in Schaback (2008) and Chapter 5 in Oettershagen (2017) for more discussion and some results regarding this phenomenon. The computations in Section 2 of Dette and Zhigljavsky (2021) are also relevant. If the data are m -constant such that $Y_m = (c, \dots, c)$ and the kernel $K(x, y) = \Phi(x - y)$ is sufficiently smooth, this reasoning suggests the conjecture that

$$\lim_{\lambda \rightarrow \infty} Y_m^\top K_\lambda(X, X)^{-1} Y_m = \lim_{\lambda \rightarrow \infty} Y_m^\top K(X_\lambda, X_\lambda)^{-1} Y_m = D_0^\top W^{-1} D_0 \quad (5.10)$$

where $D_0 = (c, 0, \dots, 0) \in \mathbb{R}^n$ and the *Wronskian* $W \in \mathbb{R}^{n \times n}$ has the elements

$$(W)_{i+1, j+1} = \frac{\partial^{i+j}}{\partial v^i \partial w^j} K(v, w) \Big|_{\substack{v=0 \\ w=0}} \quad (5.11)$$

for $i, j = 0, \dots, n - 1$. The conjectured limit $D_0^\top W^{-1} D_0$ is the data-fit term in (5.2) for the information functionals defined as $L_i f = f^{(i-1)}(0)$ for $i = 1, \dots, n$ and general m -constant data in (5.4). A proof of (5.10) is the main ingredient in the proof of the following theorem, which partially generalises Theorem 2.3 for infinitely differentiable kernels when $d = 1$.

Theorem 5.4 *Let $d = 1$ and $K(x, y) = \Phi(x - y)$. Suppose that (i) the function $\Phi: \mathbb{R} \rightarrow \mathbb{R}$ is integrable and infinitely differentiable in a neighbourhood of the origin; (ii) $\Phi^{(k)}(0) = 0$ for every odd k ; and (iii) the Fourier transform of Φ is everywhere positive. If $n \geq 1$ and the data Y are m -constant, then*

$$\lim_{\lambda \rightarrow \infty} \ell(\lambda | Y) = -\infty \quad \text{and} \quad \lambda_{\text{ML}} = \infty.$$

Proof See Section 7.10. The proof uses Equation (32) in Barthelmé and Usevich (2021). ■

6. Conclusion and Implications

In this article we have proved that Gaussian process regression with noiseless data and a stationary covariance kernel is ill-posed if the lengthscale parameter of the kernel is estimated using maximum likelihood: When the data differ from the prior mean by a constant mean shift, the maximum likelihood estimate of the lengthscale parameter is infinite (Theorem 2.3) and the conditional Gaussian process is degenerate (Theorem 2.6). As shown in Section 3, these conclusions remain valid under more general parametrisations and also applies to leave-one-out cross-validation.

6.1 Practical Implications

In a way, our results imply a practical simplification. If the data are m -constant (which, when the prior mean m is known and fixed, is trivial to check), there is no need for numerical optimisation of the log-likelihood function as one can use Theorem 2.3 to set $\lambda_{\text{ML}} = \infty$ and Theorem 2.6 to compute the conditional mean and covariance. However, degeneracy of the resulting conditional process implies that there is no predictive uncertainty and the conditional process is therefore useless as a tool for uncertainty quantification. Except for switching to a non-stationary kernel, we do not know of a good approach to fix this, and it may be that some non-stationary kernels are equally problematic and induce similar behaviour when some of their parameters are estimated.

A numerical issue that is encountered when the data are close to being m -constant is that of ill-conditioning of the covariance matrix. As $\lambda \rightarrow \infty$, the condition number of the covariance matrix $K_\lambda(X, X)$ increases with a rate related to the smoothness of the kernel. This means that one cannot compute the modified log-likelihood function for large values of λ . In practice one therefore has to either introduce a regularisation parameter to upper bound the condition number as a function of λ or select a finite maximal lengthscale $\lambda_{\text{max}} > 0$ for which $\ell(\lambda | Y)$ can be reliably computed and find the maximum likelihood estimate in $(0, \lambda_{\text{max}}]$. When the data are m -constant, restricting the feasible set for λ to $(0, \lambda_{\text{max}}]$ is likely to result in $\lambda_{\text{ML}} = \lambda_{\text{max}}$, so that the user effectively selects an arbitrary (though probably fairly large) lengthscale in this case. The dependence of predictions in λ_{max} may or may not be problematic depending on the context.

A practical recommendation borne out by our results is that all Gaussian process implementations which use maximum likelihood but not regularisation should check if the data are m -constant. If the check indicates that the data are m -constant, an implementation should either (a) forgo lengthscale estimation and output a degenerate conditional process or (b) throw an error and inform the user of the problem. If the approach (a) is chosen it should be made clear to the user that the output is degenerate as this may have important ramifications in the applied context. A more general research programme suggests itself:

- (i) To characterise, for each estimator of the kernel parameters θ , the problematic data sets Y which cause Gaussian process regression or interpolation to be ill-posed.
- (ii) To hard-code Gaussian process software to throw an error (or at least a warning) when such data are encountered.

6.2 Theoretical Implications

When analysing the convergence of Gaussian process regression as $n \rightarrow \infty$, it is typically assumed that the covariance kernel is fixed. To the best of our knowledge, in the deterministic interpolation regime only Teckentrup (2020) and Wynne et al. (2021) allow the kernel parameters other than scaling parameter σ from Section 3.3 (which does not affect the conditional mean) to vary. Their results are generic in that no specific parameter estimation method is considered and the parameter estimates are simply assumed to remain within certain sets. In Wynne et al. (2021) only a smoothness parameter, such as the parameter ν of Matérn kernels (2.3), is allowed to vary. Teckentrup (2020, Theorem 3.5) considers kernels which satisfy Assumption 2.2 and proves that the conditional mean in (2.4) tends to the true

data-generating function f such that $Y = f(X)$ if (i) this function has certain smoothness and (ii) there is a compact set which contains the estimate of λ for every n . As we have seen in this article, the second assumption fails if f happens to be a mean shift of the prior mean (i.e., $f = m + c$ for some $c \in \mathbb{R}$) and λ is estimated using maximum likelihood. This demonstrates that unless one imposes an artificial upper bound on the parameter estimates, smoothness assumptions alone are not sufficient for comprehensive convergence analysis of Gaussian process regression.

6.3 On Estimation of Other Parameters

We conclude by pointing out that our ill-posedness results are specific to lengthscales estimation and should not be expected to extend to estimation of other kernel parameters. Two examples serve to illustrate this. First, consider the scale parametrisation $K_\sigma(x, y) = \sigma^2 K(x, y)$ for a scale parameter $\sigma > 0$. From (2.8) it is straightforward to compute that the maximum likelihood estimate of σ is available in closed form:

$$\sigma_{\text{ML}} = \sqrt{\frac{Y_m^\top K(X, X)^{-1} Y_m}{n}}.$$

Here only the data $Y = m(X)$ yield a problematic parameter estimate $\sigma_{\text{ML}} = 0$ that results in degenerate predictive distributions. Consider then estimation of the smoothness parameter $\nu > 0$ of a Matérn kernel K_ν in (2.3). The presence of the coefficient $2^{1-\nu}/\Gamma(\nu)$ ensures that $K_\nu(x, x) = \sigma^2$ for every $\nu > 0$, which has two implications:

- As is well known, $K_\nu(x, y)$ tends to the Gaussian kernel $\sigma^2 \exp(-\|x - y\|^2/(2\lambda^2))$ as $\nu \rightarrow \infty$ for all $x, y \in \mathbb{R}^d$.
- As $\nu \rightarrow 0$, $K_\nu(x, y) \rightarrow \sigma^2$ if $x = y$ and $K_\nu(x, y) \rightarrow 0$ if $x \neq y$. The latter of these claims follows from the facts that $\mathcal{K}_0(z)$, the Bessel function of the second kind of zeroth order, is well-defined if $z \neq 0$ and $\Gamma(\nu) \rightarrow \infty$ as $\nu \rightarrow 0$.

This shows that both potentially problematic limiting cases, $\nu \rightarrow 0$ and $\nu \rightarrow \infty$, yield valid positive-definite kernels. Consequently, degenerate predictive distributions can never arise from estimation of the Matérn smoothness parameter.

7. Proofs

This section contains proofs for the results in Sections 2 to 5.

7.1 Interpolation in Reproducing Kernel Hilbert Spaces

Let Ω be an arbitrary set and $K: \Omega \times \Omega \rightarrow \mathbb{R}$ a positive-definite kernel, which means that

$$\sum_{n=1}^n \sum_{j=1}^n a_i a_j K(x_i, x_j) > 0 \tag{7.1}$$

for any $n \in \mathbb{N}$, any non-zero vector $a = (a_1, \dots, a_n)$, and any distinct points $x_i \in \Omega$. The kernel is positive-semidefinite if the inequality in (7.1) is not required to be strict. Then

K induces a unique *reproducing kernel Hilbert space* (RKHS), $H(K, \Omega)$. This is a Hilbert space consisting of real-valued functions defined on Ω and is equipped with an inner product $\langle \cdot, \cdot \rangle_{H(K, \Omega)}$ and the corresponding norm $\|\cdot\|_{H(K, \Omega)}$. The kernel translate $K(\cdot, x)$ is an element of $H(K, \Omega)$ for every $x \in \Omega$ and the kernel has the *reproducing property*

$$\langle K(\cdot, x), f \rangle_{H(K, \Omega)} = f(x) \quad \text{for all } x \in \Omega \text{ and } f \in H(K, \Omega).$$

It is usually not straightforward to determine whether or not a given function is an element of $H(K, \Omega)$. However, the RKHS of a kernel which satisfies Assumption 2.2 on the rate of decay of its Fourier transform is a Sobolev space; see Section 7.3. For more information on RKHSs we refer the reader to Berlinet and Thomas-Agnan (2004) and Chapters 10 and 16 in Wendland (2005).

We are interested in optimal interpolation in an RKHS. Let $f: \Omega \rightarrow \mathbb{R}$ be any function (i.e., not necessarily an element of the RKHS) that is to be interpolated at a set of distinct points $X = \{x_i\}_{i=1}^n \subset \Omega$. The *kernel interpolant* $s_{f, X}$ is the unique minimum norm interpolant to f at these points:

$$s_{f, X} = \arg \min_{s \in H(K, \Omega)} \{ \|s\|_{H(K, \Omega)} : s(x_i) = f(x_i) \text{ for every } i = 1, \dots, n \}. \quad (7.2)$$

The kernel interpolant has the explicit linear-algebraic form

$$s_{f, X}(x) = K(x, X)^\top K(X, X)^{-1} f(X), \quad (7.3)$$

which equals the conditional mean in (2.4) when $m \equiv 0$. This is the famous equivalence between Gaussian process interpolation and optimal interpolation in an RKHS whose origins can be traced back at least to the work of Kimeldorf and Wahba (1970). From (7.3) it is straightforward to compute that (e.g., Fasshauer, 2011, Section 5.1)

$$\|s_{f, X}\|_{H(K, \Omega)}^2 = f(X)^\top K(X, X)^{-1} f(X), \quad (7.4)$$

which equals the data-fit term in (2.8) for $m \equiv 0$. Note that a particular implication of (7.2) and (7.4) is that $f(X)^\top K(X, X)^{-1} f(X) \leq \|f\|_{H(K, \Omega)}^2$ if $f \in H(K, \Omega)$. How these properties of $s_{f, X}$ follow is explained in more detail in the proof of Proposition 7.10 concerning interpolation based on general linear data. For the conditional variance we use the notation

$$P_X(x)^2 = K(x, x) - K(x, X)^\top K(X, X)^{-1} K(X, x), \quad (7.5)$$

which makes the dependency on the points X explicit. Now, for every $x \in \Omega$ it holds that (e.g., Wendland, 2005, Theorem 11.4)

$$|f(x) - s_{f, X}(x)| \leq \|f\|_{H(K, \Omega)} P_X(x) \quad (7.6)$$

if $f \in H(K, \Omega)$, so that the conditional standard deviation controls the interpolation error.

7.2 On Notation

The proofs require notation that is more expressive than what we have used elsewhere in this article. Therefore the conditional variance in (2.5) equals the conditional variance (7.5)

whose dependency on the covariates has been made explicit. Similarly, the function s in (2.6) equals the kernel interpolant $s_{f,X}$ in (7.3) for any function f such that $f(X) = Y_m$. It is often necessary or useful to indicate that various quantities depend on the kernel parameters (either θ or λ). We use subscripts for this purpose. Subscripts are also used as shorthands for point sets formed by removing some elements of $X = \{x_i\}_{i=1}^n$ in the following way: $X_k = \{x_i\}_{i=1}^k$ and $X_{n,k} = X \setminus \{x_k\}$. Analogous notation is used for the conditional standard deviation and mean and the kernel interpolant constructed at these point sets, so that

$$P_{\theta,k} = P_{\theta,X_k} = P_{\theta,\{x_1,\dots,x_k\}} \quad \text{and} \quad P_{\theta,n,k} = P_{\theta,X_{n,k}} = P_{\theta,X \setminus \{x_k\}}$$

and

$$s_{\theta,k} = s_{\theta,f,X_k} = s_{\theta,f,\{x_1,\dots,x_k\}} \quad \text{and} \quad s_{\theta,n,k} = s_{\theta,f,X_{n,k}} = s_{\theta,f,X \setminus \{x_k\}}$$

for $k \leq n$ and f such that $f(X) = Y_m$. These notational conventions are reintroduced preceding their use in the proofs.

7.3 Sobolev Spaces

For $\alpha \geq 0$, the Sobolev space $W_2^\alpha(\mathbb{R}^d)$ consists of square-integrable functions $f: \mathbb{R}^d \rightarrow \mathbb{R}$ such that

$$\|f\|_{W_2^\alpha(\mathbb{R}^d)}^2 = \int_{\mathbb{R}^d} (1 + \|\xi\|^2)^\alpha |\widehat{f}(\xi)|^2 d\xi < \infty. \quad (7.7)$$

On $\Omega \subset \mathbb{R}^d$, the Sobolev space $W_2^\alpha(\Omega)$ consists of those $f: \Omega \rightarrow \mathbb{R}$ which admit an extension $f_e \in W_2^\alpha(\mathbb{R}^d)$ such that $f_e|_\Omega = f$. The norm of $W_2^\alpha(\Omega)$ is

$$\|f\|_{W_2^\alpha(\Omega)} = \inf_{f_e \in W_2^\alpha(\mathbb{R}^d)} \{ \|f_e\|_{W_2^\alpha(\mathbb{R}^d)} : f_e|_\Omega = f \}. \quad (7.8)$$

If $\alpha \in \mathbb{N}_0$, $W_2^\alpha(\mathbb{R}^d)$ consists of functions whose weak derivatives up to order α exist and are square-integrable. It is a standard result (e.g., Wendland, 2005, Corollary 10.13) that for a kernel K which satisfies Assumption 2.2 the RKHS $H(K, \mathbb{R}^d)$ is norm-equivalent to $W_2^\alpha(\mathbb{R}^d)$. This is to say that $H(K, \mathbb{R}^d)$ and $W_2^\alpha(\mathbb{R}^d)$ are equal as sets and that there are positive constants C_1 and C_2 such that

$$C_1 \|f\|_{W_2^\alpha(\mathbb{R}^d)} \leq \|f\|_{H(K, \mathbb{R}^d)} \leq C_2 \|f\|_{W_2^\alpha(\mathbb{R}^d)}$$

for every $f \in H(K, \mathbb{R}^d)$. An analogous result carries over to $H(K, \Omega)$, which is related to $H(K, \mathbb{R}^d)$ in the same way as $W_2^\alpha(\Omega)$ is to $W_2^\alpha(\mathbb{R}^d)$ via (7.8). We use the following two facts in the proof of Theorem 2.3:

- Let B be any open ball centered at the origin. Then constant functions are contained in $W_2^\alpha(\Omega)$ for any $\alpha \geq 0$ because one can construct a bump function which is constant in B and whose Fourier transform decays with a super-algebraic rate.
- The function Φ , which defines K in Assumption 2.2, is Hölder continuous with the exponent $\beta(\alpha) = \min\{1, \alpha - d/2\} > 0$ on any sufficiently regular domain Ω (e.g., an open ball). That is, there is a positive constant C such that

$$|\Phi(0) - \Phi(x)| \leq C \|x\|^{\beta(\alpha)} \quad (7.9)$$

for any $x \in \Omega$. This assertion is a consequence of the classical inclusion relation between Sobolev and Hölder spaces (e.g., Triebel, 1978, Remark 2 on p. 206). That Φ is an element of $W_2^\alpha(\mathbb{R}^d)$ is easy to verify using (2.9) and (7.7).

7.4 Proof of Theorem 2.3

We split the proof of Theorem 2.3, which is repeated below, in two.

Theorem 2.3 (Maximum likelihood estimation) *Suppose that the kernel K satisfies Assumption 2.2 and $n \geq 2$. If the data Y are m -constant, then*

$$\lim_{\lambda \rightarrow \infty} \ell(\lambda | Y) = -\infty \quad \text{and} \quad \lambda_{\text{ML}} = \infty. \quad (2.11)$$

If the data Y are not m -constant, then

$$\lim_{\lambda \rightarrow \infty} \ell(\lambda | Y) = \infty \quad \text{and} \quad \lambda_{\text{ML}} < \infty. \quad (2.12)$$

Equation (2.11) is a rather straightforward consequence of (7.4) when one interprets λ as a scaling of the covariate set instead of a kernel parameter, while proving Equation (2.12) requires some more work, including upper and lower bounds for the conditional variance. Most of the proof of Equation (2.11) is contained in the following lemmas, which will be used again in Section 7.6.

Lemma 7.1 (Continuity of the data-fit and model complexity) *Suppose that Φ is continuous and $n \geq 1$. Then the functions*

$$f_{\text{df}}(\lambda) = Y_m^\top K_\lambda(X, X)^{-1} Y_m \quad \text{and} \quad f_{\text{mc}}(\lambda) = \log \det K_\lambda(X, X)$$

are well-defined and continuous on $(0, \infty)$.

Proof Define $X_\lambda = \{x_i/\lambda\}_{i=1}^n$ and observe that $K_\lambda(X, X) = K(X_\lambda, X_\lambda)$. Since K is a positive-definite kernel and the covariates X are distinct, this shows that $K_\lambda(X, X)$ is positive-definite and hence non-singular for every $\lambda > 0$. In particular, $\det K_\lambda(X, X) > 0$ for every $\lambda > 0$ by positive-definiteness. Therefore the functions f_{df} and f_{mc} are well-defined. Because $K(x, y) = \Phi(x - y)$ is continuous, each element of $K_\lambda(X, X)$ is a continuous function of λ . From the definition of the determinant it immediately follows that f_{mc} is continuous. The continuity of f_{df} is then a consequence of, for example, Cramer's rule and the positivity of $\det K_\lambda(X, X)$. ■

Lemma 7.2 (Boundedness of the data-fit term) *Suppose that K satisfies Assumption 2.2 and $n \geq 1$. If the data Y are m -constant, then there is a constant $a > 0$ such that*

$$\sup_{\lambda > 0} Y_m^\top K_\lambda(X, X)^{-1} Y_m \leq a.$$

Proof Because the data Y are m -constant, we can write $Y_m = Y - m(X) = f(X)$ for some constant function f . Define $X_\lambda = \{x_i/\lambda\}_{i=1}^n$ and observe that, since $f(X_\lambda)$ does not depend on λ ,

$$Y_m^\top K_\lambda(X, X)^{-1} Y_m = f(X_\lambda)^\top K(X_\lambda, X_\lambda)^{-1} f(X_\lambda).$$

The RKHS $H(K, B)$ contains constant functions if B is any open ball centered at the origin by the results in Section 7.3. Because X contains a finite number of points, we can trivially select B such that $X \subset B$. Then the set X_λ is also contained in B whenever $\lambda \geq 1$. Therefore $f \in H(K, B)$ and it follows from (7.2) and (7.4) that

$$f(X_\lambda)^\top K(X_\lambda, X_\lambda)^{-1} f(X_\lambda) \leq \|f\|_{H(K, B)}^2 \quad (7.10)$$

if $\lambda \geq 1$. Because $K(x, y) = \Phi(x - y)$ for Φ which is continuous and integrable on \mathbb{R}^d ,

$$K\left(\frac{x_i}{\lambda}, \frac{x_i}{\lambda}\right) = \Phi(0) > 0 \quad \text{and} \quad \lim_{\lambda \rightarrow 0} K\left(\frac{x_i}{\lambda}, \frac{x_j}{\lambda}\right) = \lim_{\lambda \rightarrow 0} \Phi\left(\frac{x_i - x_j}{\lambda}\right) = 0$$

for all $i \neq j$. That is, $K_\lambda(X, X) = K(X_\lambda, X_\lambda)$ tends to a non-zero diagonal matrix as $\lambda \rightarrow 0$. Thus by Lemma 7.1, which guarantees the λ -continuity of the data-fit term,

$$\lim_{\lambda \rightarrow 0} Y_m^\top K_\lambda(X, X)^{-1} Y_m = \Phi(0)^{-1} \|Y_m\|^2 < \infty. \quad (7.11)$$

From (7.10) and (7.11) and Lemma 7.1 we conclude that there is a constant $a > 0$ such that

$$Y_m^\top K_\lambda(X, X)^{-1} Y_m = f(X_\lambda)^\top K(X_\lambda, X_\lambda)^{-1} f(X_\lambda) \leq a$$

for every $\lambda > 0$. ■

Proof of Equation (2.11) By Lemma 7.2, there is $a > 0$ such that

$$0 \leq Y_m^\top K_\lambda(X, X)^{-1} Y_m \leq a \quad (7.12)$$

for all $\lambda > 0$. Stationarity and continuity of Φ imply that $K_\lambda(X, X)$ converges to the identity matrix times $\Phi(0)$ as $\lambda \rightarrow 0$ and to the singular matrix of $\Phi(0)$'s as $\lambda \rightarrow \infty$. Thus it follows from Lemma 7.1 that

$$\log \det K_\lambda(X, X) \rightarrow -\infty \quad (7.13)$$

if and only if $\lambda \rightarrow \infty$. By combining (7.12) and (7.13) we conclude that

$$\ell(\lambda | Y) = Y_m^\top K_\lambda(X, X)^{-1} Y_m + \log \det K_\lambda(X, X) \rightarrow -\infty$$

if and only if $\lambda \rightarrow \infty$. Therefore $\lambda_{\text{ML}} = \arg \min_{\lambda > 0} \ell(\lambda | Y) = \infty$. ■

Three auxiliary results are needed to prove Equation (2.12). The first of these—or its variants—is well known in scattered data approximation literature (e.g., Schaback, 1995). The version that we need here is contained in the proof of Theorem 4.4 in Karvonen et al. (2020).

Proposition 7.3 *Suppose that K satisfies Assumption 2.2 and let $X = \{x_i\}_{i=1}^n$ be any set of distinct points in \mathbb{R}^d . Define $\|x - X\| = \min_{i=1, \dots, n} \|x - x_i\|$. Then there is a positive constant C , which does not depend on x or X , such that*

$$P_X(x)^2 \geq C \|x - X\|^{2\alpha-d}$$

for any $x \in \mathbb{R}^d$ for which $\|x - X\| \leq 1$, where $\alpha > d/2$ is the constant in Assumption 2.2.

The second auxiliary result, which is standard and essentially Exercise 8.66 in Iske (2018), gives a rough upper bound on the conditional variance under Assumption 2.2.

Proposition 7.4 *Suppose that K satisfies Assumption 2.2 and let $X = \{x_i\}_{i=1}^n$ be any set of distinct points in \mathbb{R}^d . Let $\|x - X\| = \min_{i=1, \dots, n} \|x - x_i\|$ and $\beta(\alpha) = \min\{1, \alpha - d/2\} > 0$, where $\alpha > d/2$ is the constant in Assumption 2.2. Then there is a positive constant C , which does not depend on x or X , such that*

$$P_X(x)^2 \leq C \|x - X\|^{\beta(\alpha)}$$

for any $x \in \mathbb{R}^d$ for which $\|x - X\|$ is sufficiently small.

Proof Let $x^* \in X$ be such that $\|x - x^*\| = \|x - X\|$. Because the standard deviation is a non-decreasing function in that $P_X(x) \leq P_{X'}(x)$ for any x if $X' \subset X$ (e.g., Wendland, 2005, Theorem 16.11), we have

$$P_X(x) \leq P_{\{x^*\}}(x). \tag{7.14}$$

Using the stationarity assumption and (7.5) we write

$$\begin{aligned} P_{\{x^*\}}(x)^2 &= K(x, x) - \frac{K(x, x^*)^2}{K(x^*, x^*)} = \Phi(0) - \frac{\Phi(x - x^*)^2}{\Phi(0)} \\ &= \frac{1}{\Phi(0)} (\Phi(0) + \Phi(x - x^*)) (\Phi(0) - \Phi(x - x^*)). \end{aligned}$$

Applying the Hölder condition (7.9) to $|\Phi(0) - \Phi(x - x^*)|$ and using $\Phi(0) \geq \Phi(x - x^*)$, which follows from the positive-definiteness of $K(x, y) = \Phi(x - y)$, yields the estimate

$$P_{\{x^*\}}(x)^2 \leq C \|x - x^*\|^{\beta(\alpha)} = C \|x - X\|^{\beta(\alpha)}$$

for a certain positive constant C which depends only on Φ . Using this bound in (7.14) concludes the proof. ■

Our third auxiliary result is an expression for the modified log-likelihood function. Although this expression has appeared in the literature (e.g., Xu and Stein, 2017, Section 4.2.2), we have not encountered its proof and therefore provide one based entirely on linear algebra, block matrix inversion, and determinantal identities. The expressions for the individual terms of the modified log-likelihood function are relatively well known. For the data-fit term, see Schaback and Werner (2006, Theorem 6) or Müller (2008, Bemerkung 3.1.4). We also point the reader to Section 3 in Scheuerer (2011). The expression for the model complexity term appear in literature on determinantal point processes (e.g., Bardenet and Hardy, 2020,

Section 2.4). For the purposes of this proposition and the proof of Equation (2.12) we use the notation

$$P_{\theta,k}(x) = P_{\theta,\{x_1,\dots,x_k\}}(x) \quad \text{and} \quad s_{\theta,k}(x) = s_{\theta,f,\{x_1,\dots,x_k\}}(x)$$

for $k \leq n$.

Proposition 7.5 *If the points $X = \{x_i\}_{i=1}^n$ are distinct, then*

$$Y_m^\top K_\theta(X, X)^{-1} Y_m = \sum_{k=0}^{n-1} \left(\frac{y_{k+1} - m(x_{k+1}) - s_{\theta,k}(x_{k+1})}{P_{\theta,k}(x_{k+1})} \right)^2 \quad (7.15)$$

and

$$\det K_\theta(X, X) = \prod_{k=0}^{n-1} P_{\theta,k}(x_{k+1})^2, \quad (7.16)$$

where $P_{\theta,0}(x_1)^2 = K_\theta(x_1, x_1)$ and $s_{\theta,0}(x_1) = 0$.

Proof Denote $y_{m,k} = y_k - m(x_k)$ and $Y_{m,k} = (y_{m,1}, \dots, y_{m,k}) \in \mathbb{R}^k$ so that

$$Y_m^\top K_\theta(X, X)^{-1} Y_m = \begin{pmatrix} y_{m,n} \\ Y_{m,n-1} \end{pmatrix}^\top \begin{pmatrix} K_\theta(x_n, x_n) & K_\theta(x_n, X_{n-1})^\top \\ K_\theta(x_n, X_{n-1}) & K_\theta(X_{n-1}, X_{n-1}) \end{pmatrix}^{-1} \begin{pmatrix} y_{m,n} \\ Y_{m,n-1} \end{pmatrix}.$$

The block matrix inversion formula, a few lines of straightforward algebra, and (2.4) and (2.5) then yield

$$Y_m^\top K_\theta(X, X)^{-1} Y_m = \left(\frac{y_{m,n} - s_{\theta,n-1}(x_n)}{P_{\theta,n-1}(x_n)} \right)^2 + Y_{m,n-1}^\top K_\theta(X_{n-1}, X_{n-1})^{-1} Y_{m,n-1},$$

iteration of which yields the form (7.15) for the data-fit term. The block determinant identity and the expression (2.5) for the conditional variance yield

$$\begin{aligned} \det K_\theta(X, X) &= \begin{vmatrix} K_\theta(x_n, x_n) & K_\theta(x_n, X_{n-1})^\top \\ K_\theta(x_n, X_{n-1}) & K_\theta(X_{n-1}, X_{n-1}) \end{vmatrix} \\ &= [K_\theta(x_n, x_n) - K_\theta(x_n, X_{n-1})^\top K_\theta(X_{n-1}, X_{n-1})^{-1} K_\theta(x_n, X_{n-1})] \\ &\quad \times \det K_\theta(X_{n-1}, X_{n-1}) \\ &= P_{\theta,n-1}(x_n)^2 \det K_\theta(X_{n-1}, X_{n-1}), \end{aligned}$$

iteration of which produces (7.16). ■

Lemma 7.6 *Suppose that K satisfies Assumption 2.2 and let $X = \{x_i\}_{i=1}^n$ be any set of distinct points in \mathbb{R}^d . Then there is a constant C , which does not depend on λ , such that*

$$\log \det K_\lambda(X, X) \geq -(2\alpha - d)n \log \lambda + C$$

for every sufficiently large $\lambda > 0$, where $\alpha > d/2$ is the constant in Assumption 2.2.

Proof Let $X_k = \{x_1, \dots, x_k\}$ and $X_{\lambda,k} = \lambda^{-1}X_k = \{\lambda^{-1}x_i\}_{i=1}^k$. Because

$$P_{\lambda,k}(x_{k+1}) = P_{X_{\lambda,k}}(\lambda^{-1}x_{k+1})$$

and

$$\begin{aligned} \|\lambda^{-1}x_{k+1} - X_{\lambda,k}\| &= \min_{i=1,\dots,k} \|\lambda^{-1}(x_{k+1} - x_i)\| = \lambda^{-1} \min_{i=1,\dots,k} \|x_{k+1} - x_i\| \\ &= \lambda^{-1} \|x_{k+1} - X_k\|, \end{aligned}$$

it follows from Proposition 7.3, when applied to the points $X_{\lambda,k}$, that, for a positive constant C which does not depend on k or λ ,

$$\begin{aligned} \log[P_{\lambda,k}(x_{k+1})^2] &= \log[P_{X_{\lambda,k}}(\lambda^{-1}x_{k+1})^2] \\ &\geq -(2\alpha - d) \log \lambda + (2\alpha - d) \log \|x_{k+1} - X_k\| + \log C \end{aligned}$$

when λ is large enough that $\|\lambda^{-1}x_{k+1} - X_{\lambda,k}\| \leq 1$. Equation (7.16) implies that

$$\log \det K_\lambda(X, X) = \sum_{k=0}^{n-1} \log[P_{\lambda,k}(x_{k+1})^2],$$

which yields the claim. ■

Proof of Equation (2.12) Under the assumption that the data are not m -constant we can freely order $\{x_1, \dots, x_n\}$ such that $y_1 - m(x_1) \neq y_2 - m(x_2)$. Then for the second term of the data-fit in (7.15) we have

$$\frac{y_2 - m(x_2) - s_{\lambda,1}(x_2)}{P_{\lambda,1}(x_2)} = \frac{y_2 - m(x_2) - K_\lambda(x_2, x_1)K_\lambda(x_1, x_1)^{-1}(y_1 - m(x_1))}{P_{\lambda,1}(x_2)}. \quad (7.17)$$

By the stationarity assumption, the numerator on the right-hand side of (7.17) is

$$y_2 - m(x_2) - \Phi\left(\frac{x_2 - x_1}{\lambda}\right)\Phi(0)^{-1}(y_1 - m(x_1)).$$

As $\lambda \rightarrow \infty$, the numerator therefore tends to $y_2 - m(x_2) - (y_1 - m(x_1)) \neq 0$. Consequently, there is a positive constant C_1 independent of λ such that

$$\begin{aligned} \ell(\lambda | Y) &= Y_m^\top K_\lambda(X, X)^{-1}Y_m + \log \det K_\lambda(X, X) \\ &= \sum_{k=0}^{n-1} \left(\frac{y_{k+1} - m(x_{k+1}) - s_{\lambda,k}(x_{k+1})}{P_{\lambda,k}(x_{k+1})} \right)^2 + \log \det K_\lambda(X, X) \\ &\geq \frac{C_1}{P_{\lambda,1}(x_2)^2} + \log \det K_\lambda(X, X) \end{aligned}$$

for all sufficiently large λ , where we have discarded all other terms of the data-fit in (7.15) except the $k = 1$ term. Proposition 7.4 yields the upper bound

$$P_{\lambda,1}(x_2)^2 \leq C_2 \lambda^{-\beta(\alpha)} \quad (7.18)$$

for $\beta(\alpha) > 0$, a certain positive constant C_2 independent of λ , and all sufficiently large λ . This bound and Lemma 7.6 then give

$$\ell(\lambda | Y) \geq \frac{C_1}{P_{\lambda,1}(x_2)^2} + \log \det K_\lambda(X, X) \geq C_1 C_2^{-1} \lambda^{\beta(\alpha)} - (2\alpha - d)n \log \lambda + C_3$$

when λ is sufficiently large and where none of the constants depend on λ . Therefore

$$\lim_{\lambda \rightarrow \infty} \ell(\lambda | Y) \geq \lim_{\lambda \rightarrow \infty} (C_1 C_2^{-1} \lambda^{\beta(\alpha)} - (2\alpha - d)n \log \lambda) = \infty$$

since $\beta(\alpha) > 0$. This concludes the proof. \blacksquare

We point out that the lower bound on $\log \det K_\lambda(X, X)$ in Lemma 7.6 is of independent interest. See Theorems 4.5 and 6.3 in Barthelmé and Usevich (2021) for other results on the behaviour of $\det K_\lambda(X, X)$ as $\lambda \rightarrow \infty$.

7.5 Proof of Theorem 2.6

Theorem 2.6 *Suppose that K satisfies Assumption 2.2 and $n \geq 1$. If $Y_m = (c, \dots, c)$ for some $c \in \mathbb{R}$, then*

$$\lim_{\lambda \rightarrow \infty} \mu_\lambda(x) = m(x) + c \quad \text{and} \quad \lim_{\lambda \rightarrow \infty} |P_\lambda(x, y)| = 0$$

for any $x, y \in \mathbb{R}^d$.

Proof Because the data are m -constant, there is a constant function $f \equiv c$ for some $c \in \mathbb{R}$ such that $Y_m = Y - m(X) = f(X)$. Recall from Section 7.3 that Assumption 2.2 implies that f is an element of $H(K, B)$ for an open ball B that can be selected such that $X \subset B$. Let $X_\lambda = \{x_i/\lambda\}_{i=1}^n$ and write

$$\mu_\lambda(x) = m(x) + s_\lambda(x) = m(x) + s_{f, X_\lambda}(\lambda^{-1}x).$$

The RKHS error estimate (7.6) yields, for any $\lambda \geq 1$,

$$|m(x) + c - \mu_\lambda(x)| = |f(\lambda^{-1}x) - s_{f, X_\lambda}(\lambda^{-1}x)| \leq \|f\|_{H(K, B)} P_{X_\lambda}(\lambda^{-1}x). \quad (7.19)$$

As in the proof of Equation (2.12), Proposition 7.4 yields the bound

$$P_\lambda(x)^2 = P_{X_\lambda}(\lambda^{-1}x)^2 \leq C \lambda^{-\beta(\alpha)} \quad (7.20)$$

for $\beta = \min\{1, \alpha - d/2\} > 0$ and a positive constant C which does not depend on λ . The claim $\lim_{\lambda \rightarrow \infty} \mu_\lambda(x) = m(x) + c$ is proved by inserting this bound in (7.19). The claim for the conditional covariance follows from (7.20) and the Cauchy–Schwarz covariance inequality $P_\lambda(x, y)^2 \leq P_\lambda(x)P_\lambda(y)$. \blacksquare

Remark 7.7 Note in the above proof it is sufficient that $P_\lambda(x)$ tends to zero as $\lambda \rightarrow \infty$. Unlike in the proof of Equation (2.12), the rate with which this convergence occurs is of no importance. From the proof of Proposition 7.4 it is easy to see that a sufficient condition for $\lim_{\lambda \rightarrow \infty} P_\lambda(x) = 0$ is that $K(x, y) = \Phi(x - y)$ for a continuous $\Phi: \mathbb{R}^d \rightarrow \mathbb{R}$. Therefore Theorem 2.6 holds if K is a continuous stationary kernel such that constant functions are contained in $H(K, B)$ for some open ball B centered at the origin. This generalisation can be extended also to Equation (2.11) and Theorem 3.1. Further generalisations are discussed in Section 5.

7.6 Proofs for Section 3

Recall the notational conventions reviewed in Section 7.2. In addition, denote

$$X_{\lambda, n, k} = \lambda^{-1} X_{n, k} = \{\lambda^{-1} x_i\}_{i \neq k}.$$

Theorem 3.1 (Cross-validation) Suppose that K satisfies Assumption 2.2 and $n \geq 2$. If the data Y are m -constant, then

$$\lim_{\lambda \rightarrow \infty} \ell_{\text{CV}}(\lambda | Y) = -\infty \quad \text{and} \quad \lambda_{\text{CV}} = \infty. \quad (3.2)$$

Proof Since the data are m -constant, there is a constant function f such that $Y_m = Y - m(X) = f(X)$. Recall from Section 7.3 that Assumption 2.2 implies that f is an element of $H(K, B)$ for an open ball B that can be selected such that $X \subset B$. Observe that $\mu_{\lambda, n, k}(x) = m(x) + s_{\lambda, n, k}(x)$. Proposition 7.4 and the RKHS error estimate (7.6) yield, for any $\lambda \geq 1$ and a constant C_1 which does not depend on λ ,

$$\begin{aligned} \ell_{\text{CV}}(\lambda | Y) &= \sum_{k=1}^n \left[\left(\frac{y_k - \mu_{\lambda, n, k}(x_k)}{P_{\lambda, n, k}(x_k)} \right)^2 + \log[P_{\lambda, n, k}(x_k)^2] \right] \\ &= \sum_{k=1}^n \left[\left(\frac{f(\lambda^{-1} x_k) - s_{f, X_{\lambda, n, k}}(\lambda^{-1} x_k)}{P_{X_{\lambda, n, k}}(\lambda^{-1} x_k)} \right)^2 + \log[P_{X_{\lambda, n, k}}(\lambda^{-1} x_k)^2] \right] \\ &\leq n \|f\|_{H(K, B)}^2 - \beta(\alpha) n \log \lambda + C_1, \end{aligned}$$

which establishes that

$$\lim_{\lambda \rightarrow \infty} \ell_{\text{CV}}(\lambda | Y) = -\infty \quad (7.21)$$

since $\beta(\alpha) = \min\{1, \alpha - d/2\} > 0$. Because $K_\lambda(x, y) \rightarrow 0$ as $\lambda \rightarrow 0$ if $x \neq y$, it is easy to compute from (2.4) and (2.5) that

$$\lim_{\lambda \rightarrow 0} s_{\lambda, n, k}(x_k) = 0 \quad \text{and} \quad \lim_{\lambda \rightarrow 0} P_{\lambda, n, k}(x_k) = \sqrt{\Phi(0)} > 0$$

for every k . Therefore

$$\lim_{\lambda \rightarrow 0} \ell_{\text{CV}}(\lambda | Y) = cn\Phi(0)^{-1} + n \log \Phi(0) > -\infty,$$

so that it follows from Equation (7.21) and the fact that $\ell_{\text{CV}}(\lambda | Y)$ is continuous that $\lambda_{\text{CV}} = \arg \min_{\lambda > 0} \ell_{\text{CV}}(\lambda | Y) = \infty$. \blacksquare

Remark 7.8 *Proving that $\lim_{\lambda \rightarrow \infty} \ell_{\text{CV}}(\lambda | Y) = \infty$ and $\lambda_{\text{CV}} < \infty$ if the data are not m -constant is more challenging than the same for the maximum likelihood estimate. This is because in the proof of Equation (2.12) one could exploit a recursive form of the modified log-likelihood function that follows from Proposition 7.5 and obtain a closed-form expression for the term in which the conditional mean and variance are based on one data point. If $n > 2$, it is more difficult to analyse any of the analogous terms in $\ell_{\text{CV}}(\lambda | Y)$ because each of these terms is built out of conditional moments based on $n - 1$ points.*

Theorem 3.3 *Suppose that K satisfies Assumption 2.2 and $n \geq 2$. If the data Y satisfy (3.4), then*

$$\lambda_{\text{ML}} = \infty \quad \text{if} \quad \{\lambda_{\text{ML}}, \beta_{\text{ML}}\} \in \arg \min_{\lambda > 0, \beta \in \mathbb{R}^q} \ell(\lambda, \beta | Y).$$

Proof The modified log-likelihood function is

$$\ell(\lambda, \beta | Y) = (Y - V(X)\beta)^\top K_\lambda(X, X)^{-1} (Y - V(X)\beta) + \log \det K_\lambda(X, X).$$

By Theorem 2.3 and identification of $Y - V(X)\beta$ as the vector Y_m , $\ell(\lambda, \beta | Y)$ tends to negative infinity for a fixed $\beta \in \mathbb{R}^q$ if and only if $Y - V(X)\beta = (c, \dots, c) \in \mathbb{R}^n$ for some $c \in \mathbb{R}$ and $\lambda \rightarrow \infty$. Furthermore, because the data-fit term is non-negative and model complexity does not depend on β , $\ell(\lambda, \beta | Y)$ cannot tend to negative infinity if λ is kept fixed. We conclude that $\lambda_{\text{ML}} = \infty$ because by the assumption in (3.4) there is $\beta^* \in \mathbb{R}^q$ for which $Y - V(X)\beta^* = (c, \dots, c)$. \blacksquare

Theorem 3.5 (Simultaneous estimation) *Suppose that K satisfies Assumption 2.2 and $n \geq 2$. Consider the maximum likelihood estimates in (3.5). If the data Y are m -constant and $Y_m \neq 0$, then*

$$\lim_{\lambda \rightarrow \infty} \ell(\lambda, \sigma_{\text{ML}}(\lambda) | Y) = -\infty \quad \text{and} \quad \lambda_{\text{ML}} = \infty.$$

Proof Recall that σ_{ML} and λ_{ML} are any minimisers of

$$\ell(\lambda, \sigma | Y) = \frac{1}{\sigma^2} Y_m^\top K_\lambda(X, X)^{-1} Y_m + \log \det K_\lambda(X, X) + n \log \sigma^2. \quad (7.22)$$

For a fixed λ it is straightforward to compute that the maximum likelihood estimate of σ is

$$\sigma_{\text{ML}}(\lambda) = \sqrt{\frac{Y_m^\top K_\lambda(X, X)^{-1} Y_m}{n}}.$$

Plugging this estimate in (7.22) yields that λ_{ML} is any minimiser of

$$\ell(\lambda, \sigma_{\text{ML}}(\lambda) | Y) = n + \log \det K_\lambda(X, X) + n \log(Y_m^\top K_\lambda(X, X)^{-1} Y_m) - n \log n.$$

The term $n \log(Y_m^\top K_\lambda(X, X)^{-1} Y_m)$ is bounded as a function of λ by Lemma 7.2 and because we have assumed that $Y_m \neq 0$. A straightforward adaptation of the proof of (2.11) then establishes the claim. \blacksquare

7.7 Proof of Theorem 4.1

Theorem 4.1 *Suppose that K satisfies Assumption 2.2, $n \geq 1$, and*

$$\lambda^{(\alpha-d/2)n} p(\lambda) \rightarrow 0 \quad \text{as} \quad \lambda \rightarrow \infty, \quad (4.4)$$

where $\alpha > d/2$ is the constant in Assumption 2.2. If $Y \in \mathbb{R}^n$ are any data, then

$$\lim_{\lambda \rightarrow \infty} \ell_{\text{MAP}}(\lambda | Y) = \infty \quad \text{and} \quad \lambda_{\text{MAP}} < \infty.$$

Proof Recall that the MAP estimator is any minimiser of

$$\ell_{\text{MAP}}(\lambda | Y) = \frac{1}{2} \ell(\lambda | Y) - \log p(\lambda),$$

where $\ell(\lambda | Y)$ is the modified log-likelihood function in (2.8). If the data are not m -constant, it follows from Theorem 2.3 that $\ell(\lambda | Y) \rightarrow \infty$ as $\lambda \rightarrow \infty$. Since p is a probability density function, it holds that $-\log p(\lambda) \rightarrow \infty$ as $\lambda \rightarrow \infty$. Consequently,

$$\lim_{\lambda \rightarrow \infty} \ell_{\text{MAP}}(\lambda | Y) = \infty \quad \text{and} \quad \lambda_{\text{MAP}} < \infty \quad (7.23)$$

if the data are not m -constant. Suppose then that the data are m -constant. The non-negativity of the data-fit term and Lemma 7.6 yield

$$\ell_{\text{MAP}}(\lambda | Y) \geq -(\alpha - d/2)n \log \lambda + C - \log p(\lambda) = -\log[\lambda^{(\alpha-d/2)n} p(\lambda)] + C$$

for a constant C that does not depend on λ . That $-\log[\lambda^{(\alpha-d/2)n} p(\lambda)] \rightarrow \infty$ as $\lambda \rightarrow \infty$ is true by assumption (4.4). Therefore (7.23) holds also when the data are m -constant. \blacksquare

7.8 Proof of Theorem 5.2

We begin with a technical lemma which is used to establish some properties of the information functionals.

Lemma 7.9 *Let Ω be a vector space and K and K_γ two positive-semidefinite kernels on Ω such that*

$$K_\gamma(x, y) = K(\gamma x, \gamma y) \quad \text{for some} \quad \gamma > 0 \quad \text{and all} \quad x, y \in \Omega.$$

Let f_γ be the function $x \mapsto f(\gamma x)$ for any $f: \Omega \rightarrow \mathbb{R}$ and, given a linear functional I , define I_γ via $I_\gamma f = I f_\gamma$. Then it holds that

1. if $f \in H(K, \Omega)$, then $f_\gamma \in H(K_\gamma, \Omega)$;
2. if I is a bounded linear functional on $H(K_\gamma, \Omega)$, then I_γ is bounded on $H(K, \Omega)$;
3. if linear functionals I_1, \dots, I_n are linearly independent on $H(K_\gamma, \Omega)$, then $I_{1,\gamma}, \dots, I_{n,\gamma}$ are linearly independent on $H(K, \Omega)$.

Proof By a classical result (e.g., Paulsen and Raghupathi, 2016, Theorem 3.11) a function f is an element of $H(K, \Omega)$ if and only if there is $c \geq 0$ such that $c^2 K(x, y) - f(x)f(y)$ defines a positive-semidefinite kernel. If $f \in H(K, \Omega)$, then $\|f\|_{H(K, \Omega)}$ equals the smallest c for which this kernel is positive-semidefinite. Let $f \in H(K, \Omega)$. Therefore $c^2 K(x, y) - f(x)f(y)$ defines a positive-semidefinite kernel for some $c \geq 0$ and

$$c^2 K(\gamma x, \gamma y) - f(\gamma x)f(\gamma y) = c^2 K_\gamma(x, y) - f_\gamma(x)f_\gamma(y) \quad (7.24)$$

also defines a positive-semidefinite kernel. Consequently, $f_\gamma \in \mathcal{H}(K_\gamma, \Omega)$. Furthermore, because $c = \|f\|_{H(K, \Omega)}$ is the smallest c for which the kernels in (7.24) are positive-semidefinite, we conclude that $\|f_\gamma\|_{H(K_\gamma, \Omega)} = \|f\|_{H(K, \Omega)}$. Because I is bounded on $H(K_\gamma, \Omega)$ and f_γ is an element of $H(K_\gamma, \Omega)$, there is a constant $C > 0$, which does not depend on f , such that

$$|I_\gamma f| = |I f_\gamma| \leq C \|f_\gamma\|_{H(K_\gamma, \Omega)} = C \|f\|_{H(K, \Omega)}.$$

This concludes the proof of the first two claims. The third claim also follows because we have proved that $f_\gamma \in H(K_\gamma, \Omega)$ if $f \in H(K, \Omega)$ and it holds that

$$a_1 I_{1, \gamma}(f) + \cdots + a_n I_{n, \gamma}(f) = a_1 I_1(f_\gamma) + \cdots + a_n I_n(f_\gamma)$$

for any $a_1, \dots, a_n \in \mathbb{R}$. ■

The following proposition is a generalisation of the minimum-norm property for interpolation based on point evaluations that was reviewed in Section 7.1. We have not been able to locate a convenient reference and therefore provide a proof.

Proposition 7.10 *Let K be a positive-definite kernel on a set Ω . If $f \in H(K, \Omega)$, then*

$$\mathcal{I}(f)^\top K(\mathcal{I}, \mathcal{I})^{-1} \mathcal{I}(f) \leq \|f\|_{H(K, \Omega)}^2 \quad (7.25)$$

for every $n \geq 1$ and every collection $\mathcal{I} = \{I_1, \dots, I_n\}$ of n non-trivial linear functionals which are linearly independent and bounded on $H(K, \Omega)$.

Proof Let $s_{\mathcal{I}} f$ be the minimum-norm interpolant to f in $H(K, \Omega)$, which is to say that

$$s_{\mathcal{I}} f = \arg \min_{s \in H(K, \Omega)} \{ \|s\|_{H(K, \Omega)} : I s = I f \text{ for all } I \in \mathcal{I} \}.$$

By, for example, Theorem 16.1 in Wendland (2005) or Section 3.3 in Oettershagen (2017) this interpolant is unique and has the explicit form

$$(s_{\mathcal{I}} f)(x) = I(f)^\top K(\mathcal{I}, \mathcal{I})^{-1} K(\mathcal{I}, x) = \sum_{i=1}^n I_i^y(K(x, y))(K(\mathcal{I}, \mathcal{I})^{-1} \mathcal{I}(f))_i.$$

Recall that we use superscripts to indicate the argument with respect to which a functional is to be applied. Because $\langle I_j^y K(\cdot, y), I_i^x K(x, \cdot) \rangle_{H(K, \Omega)} = I_i^x I_j^y K(x, y)$, from the reproducing

property we compute that

$$\begin{aligned}
 \|s_{\mathcal{I}}f\|_{H(K,\Omega)}^2 &= \langle s_{\mathcal{I}}f, s_{\mathcal{I}}f \rangle_{H(K,\Omega)} \\
 &= \sum_{i,j=1}^m \langle I_j^y K(\cdot, y), I_i^x K(x, \cdot) \rangle_{H(K,\Omega)} (K(\mathcal{I}, \mathcal{I})^{-1} \mathcal{I}(f))_i (K(\mathcal{I}, \mathcal{I})^{-1} \mathcal{I}(f))_j \\
 &= \sum_{i,j=1}^n I_i^x I_j^y K(x, y) (K(\mathcal{I}, \mathcal{I})^{-1} \mathcal{I}(f))_i (K(\mathcal{I}, \mathcal{I})^{-1} \mathcal{I}(f))_j \\
 &= \mathcal{I}(f)^\top K(\mathcal{I}, \mathcal{I})^{-1} K(\mathcal{I}, \mathcal{I}) K(\mathcal{I}, \mathcal{I})^{-1} \mathcal{I}(f) \\
 &= \mathcal{I}(f)^\top K(\mathcal{I}, \mathcal{I})^{-1} \mathcal{I}(f).
 \end{aligned}$$

If $f \in H(K, \Omega)$, the condition $Is = If$ for all $I \in \mathcal{I}$ is trivially satisfied by $s = f$ and therefore $\|s_{\mathcal{I}}f\|_{H(K,\Omega)}^2 = \mathcal{I}(f)^\top K(\mathcal{I}, \mathcal{I})^{-1} \mathcal{I}(f) \leq \|f\|_{H(K,\Omega)}^2$. \blacksquare

These preliminaries suffice to prove Theorem 5.2.

Theorem 5.2 *Suppose that Assumption 5.1 holds and let Ω_b be any convex subset of Ω such that (a) $0 \in \Omega_b$ and (b) $Lf_1 = Lf_2$ for every $L \in \mathcal{L}$ and all $f_1, f_2 \in F(\Omega)$ such that $f_1 = f_2$ on Ω_b . If the data $Y_{\mathcal{L}}$ are m -constant and constant functions are contained in $H(K, \Omega_b)$, then*

$$\lim_{\lambda \rightarrow \infty} \ell(\lambda | Y_{\mathcal{L}}) = -\infty \quad \text{and} \quad \lambda_{\text{ML}(\mathcal{L})} = \infty.$$

Proof Let $L_{i,\lambda}$ be linear functionals that are defined as $L_{i,\lambda}f = L_i f(\cdot/g(\lambda))$ for $f \in H(K, \Omega)$. These functionals are well-defined because $f(\cdot/g(\lambda)) \in H(K_{1/\lambda}, \Omega) \subset F(\Omega)$ by Lemma 7.9. Moreover, by Lemma 7.9 and the assumption that there exists $\lambda_1 > 0$ such that $g(\lambda_1) = 1$, the linear functionals $\mathcal{L}_\lambda = \{L_{1,\lambda}, \dots, L_{n,\lambda}\}$ are linearly independent and bounded on $H(K, \Omega) = H(K_{\lambda_1}, \Omega)$ for every $\lambda > 0$. Because g is continuous and $g(\lambda) \rightarrow \infty$ as $\lambda \rightarrow \infty$, there is $\lambda_0 > 0$ such that $g(\lambda) \geq 1$ for all $\lambda \geq \lambda_0$. The assumptions on Ω_b imply that $\Omega_b \subset \Omega_{b,\lambda} = \{g(\lambda)x : x \in \Omega_b\}$ if $\lambda \geq \lambda_0$. Let $f_1, f_2 \in H(K, \Omega)$ be such that $f_1 = f_2$ on Ω_b . Then $f_1(\cdot/g(\lambda)) = f_2(\cdot/g(\lambda))$ on $\Omega_b \subset \Omega_{b,\lambda}$ and thus

$$L_{i,\lambda}f_1 = L_i \left[f_1 \left(\frac{\cdot}{g(\lambda)} \right) \right] = L_i \left[f_2 \left(\frac{\cdot}{g(\lambda)} \right) \right] = L_{i,\lambda}f_2$$

by the assumption that $L_i f_1 = L_i f_2$ for every $i \in \{1, \dots, n\}$ and all $f_1, f_2 \in F(\Omega)$ such that $f_1 = f_2$ on Ω_b . Therefore the functionals \mathcal{L}_λ admit well-defined, bounded, and linearly independent restrictions on $H(K, \Omega_b)$ for every $\lambda \geq \lambda_0$. Let $f \equiv c$ for $c \in \mathbb{R}$ be the constant function such that $Y_{\mathcal{L},m} = \mathcal{L}(f)$. Then

$$Y_{\mathcal{L},m}^\top K_\lambda(\mathcal{L}, \mathcal{L})^{-1} Y_{\mathcal{L},m} = \mathcal{L}(f)^\top K_\lambda(\mathcal{L}, \mathcal{L})^{-1} \mathcal{L}(f) = \mathcal{L}_\lambda(f)^\top K(\mathcal{L}_\lambda, \mathcal{L}_\lambda)^{-1} \mathcal{L}_\lambda(f)$$

since $f(\cdot/g(\lambda)) = f$ for constant functions. The assumption that $f \in H(K, \Omega_b)$ and Proposition 7.10 then imply that

$$\sup_{\lambda \geq \lambda_0} Y_{\mathcal{L},m}^\top K_\lambda(\mathcal{L}, \mathcal{L})^{-1} Y_{\mathcal{L},m} \leq \|f\|_{H(K,\Omega_b)}^2. \quad (7.26)$$

By Assumption 5.1, the data-fit term $Y_{\mathcal{L},m}^\top K_\lambda(\mathcal{L}, \mathcal{L})^{-1} Y_{\mathcal{L},m}$ is continuous in λ (which can be proved similarly to Lemma 7.1) and satisfies

$$\limsup_{\lambda \rightarrow 0} Y_{\mathcal{L},m}^\top K_\lambda(\mathcal{L}, \mathcal{L})^{-1} Y_{\mathcal{L},m} \leq \frac{1}{\rho} \|Y_{\mathcal{L},m}\|^2 < \infty,$$

where $\rho = \liminf_{\lambda \rightarrow 0} e_{\min}(K_\lambda(\mathcal{L}, \mathcal{L})) > 0$. From this and (7.26) it follows that there is $a \geq 0$ such that

$$0 \leq Y_{\mathcal{L},m}^\top K_\lambda(\mathcal{L}, \mathcal{L})^{-1} Y_{\mathcal{L},m} \leq a \quad (7.27)$$

for all $\lambda > 0$, which is a generalisation of Lemma 7.2. Finally, Assumption 5.1 implies that $\det K_\lambda(\mathcal{L}, \mathcal{L}) \rightarrow 0$ if and only if $\lambda \rightarrow \infty$ since $\det K_\lambda(\mathcal{L}, \mathcal{L})$ is a continuous function of λ and positive for every $\lambda > 0$ by positive-definiteness of $K_\lambda(\mathcal{L}, \mathcal{L}) = K(\mathcal{L}_\lambda, \mathcal{L}_\lambda)$. We thus conclude from (7.27) that

$$\ell(\lambda | Y_{\mathcal{L}}) = Y_{\mathcal{L},m}^\top K_\lambda(\mathcal{L}, \mathcal{L})^{-1} Y_{\mathcal{L},m} + \log \det K_\lambda(\mathcal{L}, \mathcal{L})$$

is a sum two λ -continuous terms, the first of which is bounded while the second tends to negative infinity if and only if $\lambda \rightarrow \infty$. Therefore $\lambda_{\text{ML}(\mathcal{L})} = \arg \min_{\lambda > 0} \ell(\lambda | Y_{\mathcal{L}}) = \infty$. ■

7.9 Proof of Theorem 5.3

Theorem 5.3 (Estimation of multiple lengthscales) *Consider the product kernel in (5.6) and suppose that the stationary kernels K_1, \dots, K_d on \mathbb{R} satisfy Assumption 2.2. If X has the product form (5.8) with $n_i \geq 1$ for each $i = 1, \dots, d$ and the data Y are m -constant along dimension $p \in \{1, \dots, d\}$, then*

$$\lim_{\lambda_p \rightarrow \infty} \ell(\lambda_1, \dots, \lambda_d | Y) = -\infty \quad \text{and} \quad \lambda_{p,\text{ML}} = \infty,$$

where $\lambda_{p,\text{ML}}$ is the maximum likelihood estimate of the p th lengthscale parameter in (5.7).

Proof Because both the kernel K_θ and the covariates X have product forms, we may write the covariance matrix as

$$K_\theta(X, X) = K_{1,\lambda_1}(X_1, X_1) \otimes \dots \otimes K_{d,\lambda_d}(X_d, X_d),$$

where \otimes denotes the Kronecker product. We may assume that $p = 1$ without loss of generality. This means that we may write

$$Y_m = Y - m(X) = 1_{n_1} \otimes Y'_m,$$

where $1_{n_1} = (1, \dots, 1) \in \mathbb{R}^{n_1}$ and $Y'_m \in \mathbb{R}^{n'}$ for $n' = n_2 \times \dots \times n_d$ is a certain vector. Let $A = K_{2,\lambda_2}(X_2, X_2) \otimes \dots \otimes K_{d,\lambda_d}(X_d, X_d) \in \mathbb{R}^{n' \times n'}$. The properties of the Kronecker product yield

$$Y_m^\top K_\theta(X, X)^{-1} Y_m = [1_{n_1}^\top K_{1,\lambda_1}(X_1, X_1)^{-1} 1_{n_1}] \times [(Y'_m)^\top A^{-1} Y'_m]$$

and

$$\det K_\theta(X, X) = [\det K_{1,\lambda_1}(X_1, X_1)]^{n'} \times [\det A]^{n_1}.$$

Because neither Y'_m nor A depends on λ_1 , we may proceed as in the proof of Theorem 2.3 to show that $\mathbf{1}_{n_1}^\top K_{1,\lambda_1}(X_1, X_1)^{-1} \mathbf{1}_{n_1}$ is a bounded function of λ_1 while $\log \det K_{1,\lambda_1}(X_1, X_1)$ tends to negative infinity as $\lambda_1 \rightarrow \infty$. \blacksquare

7.10 Proof of Theorem 5.4

Let $d = 1$ and $K(x, y) = \Phi(x - y)$ for an integrable function $\Phi: \mathbb{R} \rightarrow \mathbb{R}$ that is infinitely differentiable in a neighbourhood of the origin. Define the diagonal matrix

$$\Delta_\lambda = \text{diag}(1, \lambda^{-1}, \dots, \lambda^{n-1})$$

and the *Vandermonde matrix*

$$V(X) = \begin{pmatrix} 1 & x_1 & \cdots & x_1^{n-1} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & x_n & \cdots & x_n^{n-1} \end{pmatrix}.$$

It is a standard result in polynomial interpolation that the Vandermonde matrix is non-singular when x_i are distinct. Recall that W is the Wronskian defined in (5.11). The Wronskian is non-singular if the Fourier transform of Φ is positive (Lee et al., 2007, Lemma 3.3). To see this, note first that

$$(-1)^j (-i)^{i+j} \frac{\partial^{i+j}}{\partial v^i \partial w^j} K(v, w) \Big|_{\substack{v=0 \\ w=0}} = (-i)^{i+j} \Phi^{(i+j)}(0) = \frac{1}{2\pi} \int_{\mathbb{R}} \xi^{i+j} \widehat{\Phi}(\xi) \, d\xi$$

and observe then that

$$\sum_{i=0}^{n-1} \sum_{j=0}^{n-1} a_i a_j \Phi^{(i+j)}(0) = \frac{1}{2\pi} \int_{\mathbb{R}} \widehat{\Phi}(\xi) \left(\sum_{i=1}^n a_i \xi^i \right)^2 \, d\xi$$

is positive for any $n \in \mathbb{N}$ and any non-zero $a = (a_0, \dots, a_{n-1}) \in \mathbb{R}^n$. It follows from these equations that the Wronskian is non-singular because it can be written as a product of a positive-definite matrix with elements $\Phi^{(i+j)}(0)$ and two non-singular diagonal matrices.

Theorem 5.4 *Let $d = 1$ and $K(x, y) = \Phi(x - y)$. Suppose that (i) the function $\Phi: \mathbb{R} \rightarrow \mathbb{R}$ is integrable and infinitely differentiable in a neighbourhood of the origin; (ii) $\Phi^{(k)}(0) = 0$ for every odd k ; and (iii) the Fourier transform of Φ is everywhere positive. If $n \geq 1$ and the data Y are m -constant, then*

$$\lim_{\lambda \rightarrow \infty} \ell(\lambda | Y) = -\infty \quad \text{and} \quad \lambda_{\text{ML}} = \infty.$$

Proof We proceed as in the proof of Equation (2.11) in Section 7.4 and conclude that to prove the claim it is sufficient to show that the limit $\lim_{\lambda \rightarrow \infty} Y_m^\top K_\lambda(X, X)^{-1} Y_m$ exists and

is finite. To prove that this is so we use Equation (32) in Barthelmé and Usevich (2021), which states that

$$K_\lambda(X, X) = V(X)\Delta_\lambda W \Delta_\lambda V(X)^\top + \lambda^{-n}[V(X)\Delta_\lambda W_{1,\lambda} + W_{2,\lambda}\Delta_\lambda V(X)^\top] + \lambda^{-2n}W_{3,\lambda},$$

where the matrices $W_{1,\lambda}$, $W_{2,\lambda}$, and $W_{3,\lambda}$ are bounded as $\lambda \rightarrow \infty$. We may now use this equation to write

$$Y_m^\top K_\lambda(X, X)^{-1} Y_m = Y_m^\top V(X)^{-\top} \Delta_\lambda^{-1} (W + A_\lambda)^{-1} \Delta_\lambda^{-1} V(X)^{-1} Y_m, \quad (7.28)$$

where

$$A_\lambda = \lambda^{-n}[W_{1,\lambda}V(X)^{-\top} \Delta_\lambda^{-1} + \Delta_\lambda^{-1}V(X)^{-1}W_{2,\lambda}] + \lambda^{-2n}\Delta_\lambda^{-1}V(X)^{-1}W_{3,\lambda}V(X)^{-\top} \Delta_\lambda^{-1}.$$

Since $\Delta_\lambda^{-1} = \mathcal{O}(\lambda^{n-1})$, we conclude that $A_\lambda = \mathcal{O}(\lambda^{-1})$. Furthermore, from the assumption that $Y_m = (c, \dots, c)$ for some $c \in \mathbb{R}$ it follows that $V(X)^{-1}Y_m = (c, 0, \dots, 0) \in \mathbb{R}^n$. Consequently, $\Delta_\lambda^{-1}V(X)^{-1}Y_m = (c, 0, \dots, 0) = D_0$. Therefore (7.28) yields

$$Y_m^\top K_\lambda(X, X)^{-1} Y_m = D_0^\top (W + A_\lambda)^{-1} D_0 \rightarrow D_0^\top W^{-1} D_0 < \infty \quad \text{as } \lambda \rightarrow \infty,$$

where the Wronskian W is, as noted earlier, non-singular. ■

Acknowledgements

TK was supported by the Academy of Finland postdoctoral researcher grant #338567 “Scalable, adaptive and reliable probabilistic integration”. CJO was supported by the Alan Turing Institute, United Kingdom, and the Engineering & Physical Sciences Research Council grant EP/WO19590/1. We are grateful to Philipp Hennig, Motonobu Kanagawa, and François-Xavier Briol for helpful general comments and to Simon Barthelmé and Konstantin Usevich for pointing out how to prove Theorem 5.4. Comments by an anonymous reviewer served as inspiration for Sections 5.2 and 6.3.

References

- I. Adrianakis and P. G. Challenor. The effect of the nugget on Gaussian process emulators for computer models. *Computational Statistics and Data Analysis*, 56(12):4215–4228, 2012.
- E. Anderes. On the consistent separation of scale and variance for Gaussian random fields. *The Annals of Statistics*, 38(2):870–893, 2010.
- F. Bachoc. Cross validation and maximum likelihood estimations of hyper-parameters of Gaussian processes with model misspecification. *Computational Statistics & Data Analysis*, 66:55–69, 2013.
- F. Bachoc, A. Lagnoux, and T. M. N. Nguyen. Cross-validation estimation of covariance parameters under fixed-domain asymptotics. *Journal of Multivariate Analysis*, 160:42–67, 2017.

- R. Bardenet and A. Hardy. Monte Carlo with determinantal point processes. *The Annals of Applied Probability*, 30(1):368–417, 2020.
- S. Barthelmé and K. Usevich. Spectral properties of kernel matrices in the flat limit. *SIAM Journal on Matrix Analysis and Applications*, 42(1):17–57, 2021.
- S. Barthelmé, P.-O. Amblard, Nicolas Tremblay, and Konstantin Usevich. Gaussian process regression in the flat limit. *arXiv:2201.01074v2*, 2022.
- A. Berlinet and C. Thomas-Agnan. *Reproducing Kernel Hilbert Spaces in Probability and Statistics*. Springer, 2004.
- M. Bevilacqua, T. Faouzi, R. Furrer, and E. Porcu. Estimation and prediction using generalized Wendland covariance functions under fixed domain asymptotics. *The Annals of Statistics*, 47(2):828–856, 2019.
- L. Brevault, M. Balesdent, and A. Hebbal. Overview of Gaussian process based multi-fidelity techniques with variable relationship between fidelities, application to aerospace systems. *Aerospace Science and Technology*, 107:106339, 2020.
- F.-X. Briol, C. J. Oates, M. Girolami, M. A. Osborne, and D. Sejdinovic. Probabilistic integration: A role in statistical computation? *Statistical Science*, 34(1):1–22, 2019.
- R. Cavoretto. Adaptive radial basis function partition of unity interpolation: A bivariate algorithm for unstructured data. *Journal of Scientific Computing*, 87(41), 2021.
- Y. Chen, H. Owhadi, and A. M. Stuart. Consistency of empirical Bayes and kernel flow for hierarchical parameter estimation. *Mathematics of Computation*, 90:2527–2578, 2021.
- J. Cockayne, C. J. Oates, T. J. Sullivan, and M. Girolami. Bayesian probabilistic numerical methods. *SIAM Review*, 61(4):756–789, 2019.
- J. P. Cunningham, K. V. Shenoy, and M. Sahani. Fast Gaussian process methods for point process intensity estimation. In *Proceedings of the 25th International Conference on Machine learning*, pages 192–199, 2008.
- C. Currin, T. Mitchell, M. Morris, and D. Ylvisaker. A Bayesian approach to the design and analysis of computer experiments. ORNL-6498, Oak Ridge National Laboratory, 1988.
- H. Dette and A. Zhigljavsky. Reproducing kernel Hilbert spaces, polynomials, and the classical moment problem. *SIAM/ASA Journal on Uncertainty Quantification*, 9(4):1589–1614, 2021.
- P. Diaconis. Bayesian numerical analysis. In *Statistical decision theory and related topics IV*, volume 1, pages 163–175. Springer-Verlag New York, 1988.
- D. Dong. Mine gas emission prediction based on Gaussian process model. *Procedia Engineering*, 45:334–338, 2012.
- J. Du, H. Zhang, and V. S. Mandrekar. Fixed-domain asymptotic properties of tapered maximum likelihood estimators. *The Annals of Statistics*, 37(6A):3330–3361, 2009.

- G. E. Fasshauer. Positive definite kernels: past, present and future. *Dolomite Research Notes on Approximation*, 4:21–63, 2011.
- G. E. Fasshauer and M. McCourt. *Kernel-based Approximation Methods Using MATLAB*. Number 19 in Interdisciplinary Mathematical Sciences. World Scientific Publishing, 2015.
- E. Fong and C. C. Holmes. On the marginal likelihood and cross-validation. *Biometrika*, 107(2):489–496, 2020.
- S. Geisser and W. F. Eddy. A predictive approach to model selection. *Journal of the American Statistical Association*, 74(365):153–160, 1979.
- GPY. GPY: A Gaussian process framework in Python. <http://github.com/SheffieldML/GPY>, since 2012.
- M. Gu, X. Wang, and J. O. Berger. Robust Gaussian stochastic process emulation. *The Annals of Statistics*, 46(6A):3038–3066, 2018.
- I. Henderson. Sobolev regularity of Gaussian random fields. *arXiv:2209.02703v1*, 2022.
- P. Hennig, M. A. Osborne, and H. P. Kersting. *Probabilistic Numerics: Computation as Machine Learning*. Cambridge University Press, 2022.
- A. Iske. *Approximation Theory and Algorithms for Data Analysis*. Number 68 in Texts in Applied Mathematics. Springer, 2018.
- M. Kanagawa, P. Hennig, D. Sejdinovic, and B. K. Sriperumbudur. Gaussian processes and kernel methods: A review on connections and equivalences. *arXiv:1807.02582v1*, 2018.
- T. Karvonen. Asymptotic bounds for smoothness parameter estimates in Gaussian process interpolation. *arXiv:2203.05400v3*, 2023.
- T. Karvonen, F. Tronarp, and S. Särkkä. Asymptotics of maximum likelihood parameter estimates for Gaussian processes: The Ornstein–Uhlenbeck prior. In *29th IEEE International Workshop on Machine Learning for Signal Processing*, 2019.
- 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. *SIAM/ASA Journal on Uncertainty Quantification*, 8(3):926–958, 2020.
- C. G. Kaufman and B. A. Shaby. The role of the range parameter for estimation and prediction in geostatistics. *Biometrika*, 100(2):473–484, 2013.
- M. C. Kennedy and A. O’Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(3):425–464, 2002.
- G. S. Kimeldorf and G. Wahba. A correspondence between Bayesian estimation on stochastic processes and smoothing by splines. *The Annals of Mathematical Statistics*, 41(2):495–502, 1970.

- 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(3–4):771–813, 2016.
- K. Kowalska and L. Peel. Maritime anomaly detection using Gaussian process active learning. In *Proceedings of the 15th International Conference on Information Fusion*, pages 1164–1171, 2012.
- J. Latz. On the well-posedness of Bayesian inverse problems. *SIAM/ASA Journal on Uncertainty Quantification*, 8(1):451–482, 2020.
- Y. J. Lee, G. J. Yoon, and J. Yoon. Convergence of increasingly flat radial basis interpolants to polynomial interpolants. *SIAM Journal on Mathematical Analysis*, 39(2):537–553, 2007.
- Y. J. Lee, C. A. Micchelli, and J. Yoon. A study on multivariate interpolation by increasingly flat kernel functions. *Journal of Mathematical Analysis and Applications*, 427(1):74–87, 2015.
- W.-L. Loh. Fixed-domain asymptotics for a subclass of Matérn-type Gaussian random fields. *The Annals of Statistics*, 33(5):2344–2394, 2005.
- D. J. C. MacKay. Bayesian interpolation. *Neural Computation*, 4(3):415–447, 1992.
- K. V. Mardia and R. J. Marshall. Maximum likelihood estimation of models for residual covariance in spatial regression. *Biometrika*, 71(1):135–146, 1984.
- A. G. de G. Matthews, M. van der Wilk, T. Nickson, K. Fujii, A. Boukouvalas, P. León-Villagrà, Z. Ghahramani, and J. Hensman. GPflow: A Gaussian process library using TensorFlow. *Journal of Machine Learning Research*, 18(40):1–6, apr 2017. URL <http://jmlr.org/papers/v18/16-537.html>.
- H. Q. Minh. Some properties of Gaussian reproducing kernel Hilbert spaces and their implications for function approximation and learning theory. *Constructive Approximation*, 32(2):307–338, 2010.
- S. Müller. *Komplexität und Stabilität von kernbasierten Rekonstruktionsmethoden*. PhD thesis, University of Göttingen, 2008.
- J. Oettershagen. *Construction of Optimal Cubature Algorithms with Applications to Econometrics and Uncertainty Quantification*. PhD thesis, Faculty of Mathematics and Natural Sciences, University of Bonn, 2017.
- A. O’Hagan. Curve fitting and optimal design for prediction. *Journal of the Royal Statistical Society. Series B (Methodological)*, 40(1):1–42, 1978.
- V. I. Paulsen and M. Raghupathi. *An Introduction to the Theory of Reproducing Kernel Hilbert Spaces*. Number 152 in Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2016.

- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- P. Ranjan, R. Haynes, and R. Karsten. A computationally stable approach to Gaussian process interpolation of deterministic computer simulation data. *Technometrics*, 53(4): 366–378, 2011.
- C. E. Rasmussen and H. Nickisch. Gaussian processes for machine learning (GPML) toolbox. *Journal of Machine Learning Research*, 11:3011–3015, 2010.
- C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. Adaptive Computation and Machine Learning. 2006.
- L. E. Revell, A. Stenke, F. Tummon, A. Feinberg, E. Rozanov, T. PEter, N. L. Abraham, H. Akiyoshi, A. T. Archibald, N. Butchart, M. Deushi, P. Jöckel, D. Kinnison, M. Michou, O. Morgenstern, F. M. O’Connor, L. D. Oman, G. Pitari, D. A. Plummer, R. Schofield, K. Stone, S. Tilmes, D. Vioni, Y. Yamashita, and G. Zeng. Tropospheric ozone in CCMI models and Gaussian process emulation to understand biases in the SOCOLv3 chemistry–climate model. *Atmospheric Chemistry and Physics*, 18(21):16155–16172, 2018.
- S. Rippa. An algorithm for selecting a good value for the parameter c in radial basis function interpolation. *Advances in Computational Mathematics*, 11(2):193–210, 1999.
- O. Roustant, D. Ginsbourger, and Y. Deville. DiceKriging, DiceOptim: Two R packages for the analysis of computer experiments by kriging-based metamodeling and optimization. *Journal of Statistical Software*, 51(1), 2012.
- J. Sacks, Mitchell T. J. Welch, W. J., and H. P. Wynn. Design and analysis of computer experiments. *Statistical Science*, 4(4):409–435, 1989.
- T. J. Santner, B. J. Williams, and W. I. Notz. *The Design and Analysis of Computer Experiments*. Springer Series in Statistics. Springer, 2003.
- R. Schaback. Error estimates and condition numbers for radial basis function interpolation. *Advances in Computational Mathematics*, 3(3):251–264, 1995.
- R. Schaback. Limit problems for interpolation by analytical radial basis functions. *Journal of Computational and Applied Mathematics*, 212(2):127–149, 2008.
- R. Schaback and J. Werner. Linearly constrained reconstruction of functions by kernels with applications to machine learning. *Advances in Computational Mathematics*, 25:237, 2006.
- M. Scheuerer. An alternative procedure for selecting a good value for the parameter c in RBF-interpolation. *Advances in Computational Mathematics*, 34(1):105–126, 2011.
- M. Scheuerer, R. Schaback, and M. Schlather. Interpolation of spatial data – A stochastic or a deterministic problem? *European Journal of Applied Mathematics*, 24(4):601–629, 2013.

- J. Snoek, H. Larochelle, and R. P. Adams. Practical Bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems*, volume 25, pages 2951–2959, 2012.
- E. Solak, R. Murray-Smith, W. E. Leithead, D. J. Leith, and C.E. Rasmussen. Derivative observations in Gaussian process models of dynamic systems. In *Advances in Neural Information Processing Systems*, volume 15, pages 1057–1064, 2002.
- G. Song, J. Riddle, G. E. Fasshauer, and F. J. Hickernell. Multivariate interpolation with increasingly flat radial basis functions of finite smoothness. *Advances in Computational Mathematics*, 36(3):485–501, 2012.
- M. L. Stein. *Interpolation of Spatial Data: Some Theory for Kriging*. Springer Series in Statistics. Springer, 1999.
- I. Steinwart. Convergence types and rates in generic Karhunen-Loève expansions with applications to sample path properties. *Potential Analysis*, 51:361–395, 2019.
- I. Steinwart, D. Hush, and C. Scovel. An explicit description of the reproducing kernel Hilbert spaces of Gaussian RBF kernels. *IEEE Transactions on Information Theory*, 52(10):4635–4643, 2006.
- A. M. Stuart. Inverse problems: a Bayesian perspective. *Acta Numerica*, 19:451–559, 2010.
- Hong-Wei Sun and Ding-Xuan Zhou. Reproducing kernel Hilbert spaces associated with analytic translation-invariant Mercer kernels. *Journal of Fourier Analysis and Applications*, 14(1):89–101, 2008.
- B. Szabó, A. W. van der Vaart, and J. H. van Zanten. Frequentist coverage of adaptive nonparametric Bayesian credible sets. *The Annals of Statistics*, 43(4):1391–1428, 2015.
- A. L. Teckentrup. Convergence of Gaussian process regression with estimated hyperparameters and applications in Bayesian inverse problems. *SIAM/ASA Journal on Uncertainty Quantification*, 8(4):1310–1337, 2020.
- H. Triebel. *Interpolation Theory, Function Spaces, Differential Operators*. Number 18 in North-Holland Mathematical Library. North-Holland Publishing Company, 1978.
- A. van der Vaart and H. van Zanten. Information rates of nonparametric Gaussian process methods. *Journal of Machine Learning Research*, 12(6):2095–2119, 2011.
- W. Wang. On the inference of applying Gaussian process modeling to a deterministic function. *Electronic Journal of Statistics*, 15(2):5014–5066, 2021.
- H. Wendland. *Scattered Data Approximation*. Number 17 in Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press, 2005.
- H. Wendland and C. Rieger. Approximate interpolation with applications to selecting smoothing parameters. *Numerische Mathematik*, 101(4):729–748, 2005.

- A. G. Wilson, Z. Hu, R. Salakhutdinov, and E. P. Xing. Deep kernel learning. In *Artificial Intelligence and Statistics*, pages 370–378. PMLR, 2016.
- G. Wynne, F.-X Briol, and M. Girolami. Convergence guarantees for Gaussian process means with misspecified likelihoods and smoothness. *Journal of Machine Learning Research*, 22(123):1–40, 2021.
- W. Xu and M. L. Stein. Maximum likelihood estimation for a smooth Gaussian random field model. *SIAM/ASA Journal on Uncertainty Quantification*, 5(1):138–175, 2017.
- Z. Ying. Asymptotic properties of a maximum likelihood estimator with data from a Gaussian process. *Journal of Multivariate Analysis*, 36(2):280–296, 1991.
- H. Zhang. Inconsistent estimation and asymptotically equal interpolations in model-based geostatistics. *Journal of the American Statistical Association*, 99(465):250–261, 2004.