



HAL
open science

Robust prediction interval estimation for Gaussian processes by cross-validation method

Naoufal Acharki, Antoine Bertoncello, Josselin Garnier

► **To cite this version:**

Naoufal Acharki, Antoine Bertoncello, Josselin Garnier. Robust prediction interval estimation for Gaussian processes by cross-validation method. *Computational Statistics and Data Analysis*, 2023, 178, pp.107597. 10.1016/j.csda.2022.107597 . hal-03897311

HAL Id: hal-03897311

<https://inria.hal.science/hal-03897311v1>

Submitted on 17 Sep 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License



Robust prediction interval estimation for Gaussian processes by cross-validation method



Naoufal Acharki ^{a,b,*}, Antoine Bertonecello ^b, Josselin Garnier ^a

^a Centre de Mathématiques Appliquées, Ecole Polytechnique, 91128 Palaiseau, France

^b TotalEnergies OneTech, 91120 Palaiseau, France

ARTICLE INFO

Article history:

Received 10 June 2021

Received in revised form 12 August 2022

Accepted 14 August 2022

Available online 23 August 2022

Keywords:

Cross-validation

Coverage probability

Gaussian processes

Prediction intervals

ABSTRACT

Probabilistic regression models typically use the Maximum Likelihood Estimation or Cross-Validation to fit parameters. These methods can give an advantage to the solutions that fit observations on average, but they do not pay attention to the coverage and the width of Prediction Intervals. A robust two-step approach is used to address the problem of adjusting and calibrating Prediction Intervals for Gaussian Processes Regression. First, the covariance hyperparameters are determined by a standard Cross-Validation or Maximum Likelihood Estimation method. A Leave-One-Out Coverage Probability is introduced as a metric to adjust the covariance hyperparameters and assess the optimal type II Coverage Probability to a nominal level. Then a relaxation method is applied to choose the hyperparameters that minimize the Wasserstein distance between the Gaussian distribution with the initial hyperparameters (obtained by Cross-Validation or Maximum Likelihood Estimation) and the proposed Gaussian distribution with the hyperparameters that achieve the desired Coverage Probability. The method gives Prediction Intervals with appropriate coverage probabilities and small widths.

© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Many approaches of supervised learning focus on point prediction by producing a single value for a new point and do not provide information about how far those predictions may be from true response values. This may be inadmissible, especially for systems that require risk management. Indeed, an interval is crucial and offers valuable information that helps for better management than just predicting a single value.

The Prediction Intervals are well-known tools to provide more information by quantifying and representing the level of uncertainty associated with predictions. One existing and popular approach for prediction models without predictive distribution (e.g. Random Forest or Gradient Boosting models) is the bootstrap, starting from the Traditional bootstrap (Efron and Tibshirani, 1994; Heskes, 1997) to Improved Bootstrap (Li et al., 2018). It is considered as one of the most used methods (Efron and Tibshirani, 1994) for estimating empirical variances and for constructing Predictions Intervals, it is claimed to achieve good performance under some asymptotic framework.

A set of empirical methods have been proposed for these models to build Prediction Intervals like the Infinitesimal Jackknife (Wager et al., 2014), Jackknife-after-Bootstrap methods (Efron, 1992), Quantile Random Forest (Meinshausen, 2006) Out-Of-Bag intervals (Zhang et al., 2020) and Conformal prediction (Lei et al., 2018; Romano et al., 2019). In the Deep Learn-

* Corresponding author at: TotalEnergies OneTech, 91120 Palaiseau, France.

E-mail address: naoufal.acharki@polytechnique.edu (N. Acharki).

ing field, many recent methods have been also developed to quantify the uncertainty in Neural networks: The Delta method (Hwang and Ding, 1997), Mean-Variance Estimation (Nix and Weigend, 1994), the Bayesian approach (MacKay, 1992; Gal and Ghahramani, 2016), Lower Upper Bound Estimation (Khosravi et al., 2011) and High-Quality Driven ensemble approach (Pearce et al., 2018). Most methods estimate the Coverage Probability (CP) (Landon and Singpurwalla, 2008) and the mean Prediction Interval width (MPIW) (Khosravi et al., 2010) by using the combinational Coverage Width-based Criterion (CWC) as a metric to identify model's parameters or define a loss function with a Lagrangian controlling the importance of the width and coverage. Pang et al. (2018) suggest the Receiver Operating Characteristic curve of Prediction Interval (ROC-PI), a graphic indicator that serves as a trade-off between the intervals width and CP for identifying the best parameters.

Unlike Ensemble methods or Neural Networks, there exist several prediction models with a probabilistic framework like the Gaussian Processes (GP) model (Rasmussen and Williams, 2005) which are able to compute an efficient predictor with associated uncertainty. These models are more suitable for uncertainty quantification. They provide a predictive distribution with both points prediction and interval estimation and do not require any empirical approach such as the bootstrap. In most cases, the predictive distributions of GP models are obtained either with a *plug-in* method that takes the Maximum Likelihood Estimator (MLE) (Mardia and Marshall, 1984; Stein, 1999) of the model's hyperparameters or by using a Full-Bayesian approach that takes into account the posterior distribution of the hyperparameters and propagates it into the predictive distribution. However, both methods suffer from some limitations. Indeed, the MLE approach works well only when the model is well-specified and may fail in case of model misspecification (Bachoc, 2013b). At the same time, the Full-Bayesian approach is very complex to implement, typically with a Markov chain Monte Carlo (MCMC) algorithm and is sensitive to the choice of the *prior* distribution of the hyperparameters (Filippone et al., 2013; Muré, 2018). On the other hand, the calibration of Prediction Intervals is little studied in the literature. Lawless and Fredette (2005) proposed a frequentist approach to predictive distribution to build and calibrate the Prediction Intervals. However, to the best of our knowledge, this approach has not yet been extended to models with a predictive distribution and we do not have any guarantees that it can work in the case of a misspecified model. Furthermore, improving the modelling of the covariance function seems to be efficient in overcoming the issue of a misspecified model. However, it may lead to complex covariance models and, consequently, severe difficulties in estimating the covariance function's hyperparameters, especially in high dimensions. Moreover, sometimes, it is challenging to find proper modelling without further knowledge of the system and the sources of uncertainty. In this work, we propose a method based on Cross-Validation (CV) on the GP model to address the problem of model misspecification and calibrate Prediction Intervals by adjusting the upper and lower bounds to satisfy the desired level of CP. The method gives Prediction Intervals with appropriate coverage probabilities and small widths.

The paper is organized as follows. Section 2 formulates the problem of Prediction Intervals estimation. Section 3 introduces the Gaussian Process regression model and its training methods. In Section 4, we present a method for estimating robust Prediction Intervals supported by theoretical results. We show in Section 5 the application of this method to academic examples and to an industrial example. Finally, we present our conclusions in Section 6.

2. Problem formulation

We consider n observations of an empirical model or computer code f . Each observation of the output corresponds to a d -dimensional input vector $\mathbf{x} = (x_1, \dots, x_d)^\top \in \mathcal{D} \subseteq \mathbb{R}^d$. The n points corresponding to the model/code runs are called an experimental design $\mathbf{X} = (\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)})$ where $\mathbf{x}^{(i)} = (x_1^{(i)}, \dots, x_d^{(i)})^\top \in \mathcal{D}$. The outputs are denoted by $\mathbf{y} = (y^{(1)}, \dots, y^{(n)}) \in \mathbb{R}^n$ with $y^{(i)} = f(\mathbf{x}^{(i)})$. We seek to estimate the unobserved function $\mathbf{x} \in \mathcal{D} \mapsto f(\mathbf{x})$ from the data \mathbf{y} and make accurate predictions with the associated uncertainty.

Formally, let assume that f is a realization of random process Y and let $Y(\mathbf{x})$ be the value of model output at a point $\mathbf{x} \in \mathcal{D}$, let $\alpha \in [0, 1]$ describe the nominal level of confidence. We wish to estimate the interval $\mathcal{PI}_{1-\alpha}$ with respect to the type II CP (the conditional Coverage Probability given the training set) such that the probability

$$\mathbb{P}(Y(\mathbf{x}) \in \mathcal{PI}_{1-\alpha}(\mathbf{x}) \mid \mathbf{X}, \mathbf{y}) \tag{1}$$

is as close as possible to $1 - \alpha$. In most cases, $\mathcal{PI}_{1-\alpha}$ is a two-sided interval delimited by two bounds at $\mathbf{x} \in \mathcal{D}$

$$\mathcal{PI}_{1-\alpha}(\mathbf{x}) := [y_{\alpha/2}(\mathbf{x}), y_{1-\alpha/2}(\mathbf{x})], \tag{2}$$

where $y_{\alpha/2}(\mathbf{x}) = \tilde{y}(\mathbf{x}) + z_{\alpha/2} \times \tilde{\sigma}(\mathbf{x})$ is the lower bound, $y_{1-\alpha/2}(\mathbf{x}) = \tilde{y}(\mathbf{x}) + z_{1-\alpha/2} \times \tilde{\sigma}(\mathbf{x})$ is the upper bound, $z_{\alpha/2}$ (resp. $z_{1-\alpha/2}$) is the $\alpha/2$ (resp. $1 - \alpha/2$) quantile of the normalized predictive distribution (e.g. t -distribution for regression prediction), $\tilde{y}(\mathbf{x}) = \mathbb{E}(Y(\mathbf{x}) \mid \mathbf{X}, \mathbf{y})$ and $\tilde{\sigma}^2(\mathbf{x}) = \text{Var}(Y(\mathbf{x}) \mid \mathbf{X}, \mathbf{y})$ are the predictive mean and variance.

In the framework of kriging, the prior distribution of the process Y is Gaussian characterized by a mean and covariance. The Cumulative Distribution Function (CDF) of the predictive variable $Y(\mathbf{x})$ given \mathbf{X} and \mathbf{y} is well-defined and continuous with the Gaussian distribution. The quantile function is defined then as the inverse of the CDF and the quantiles $z_{\alpha/2}$ and $z_{1-\alpha/2}$ are fully characterized. Thus, estimating the interval $\mathcal{PI}_{1-\alpha}$ in equation (2) is equivalent to estimate the predictive mean $\tilde{y}(\mathbf{x})$ and variance $\tilde{\sigma}^2(\mathbf{x})$.

Therefore, the objective is to build a surrogate model to estimate correctly the upper and lower bounds of Prediction Intervals $\mathcal{PI}_{1-\alpha}$. This goes through the CV method with respect to the CP. In the following sections, $\|\cdot\|$ refers to the

Euclidean norm $\|\cdot\|_2$ if applied to a vector and to the Frobenius norm, defined by $\|\mathbf{M}\|_F = (\text{Tr}(\mathbf{M}\mathbf{M}^\top))^{1/2}$, if applied to a matrix.

3. Modeling with Gaussian processes

We use the GP model to learn the unobserved function f . It is a Bayesian non-parametric regression (see Tipping (2004) for Bayesian inference) which employs GP *prior* over the regression functions. It will be converted into a *posterior* over functions once some data has been observed. In the kriging framework (Rasmussen and Williams, 2005; Stein, 1999), Y is assumed *a priori* to be a GP with mean $\mu(\mathbf{x})$ and covariance function $\mathbf{k}(\mathbf{x}, \mathbf{x}') + \sigma_\epsilon^2 \mathbf{1}\{\mathbf{x} = \mathbf{x}'\}$ for all $\mathbf{x}, \mathbf{x}' \in \mathcal{D}$. $\sigma_\epsilon^2 \geq 0$ is the variance of measurement error, also called the nugget effect.

3.1. The mean and covariance functions

The assumption made on the existing knowledge of the model Y and the mean function μ defines three sub-cases of kriging

- The Simple Kriging: μ is assumed to be known, usually null $\mu = 0$.
- The Ordinary Kriging: μ is assumed to be constant but unknown.
- The Universal Kriging: μ is assumed to be of the form $\sum_{j=1}^p \beta_j f_{j-1}(\mathbf{x})$, where f_j are predefined (e.g. polynomial functions $f_0(\mathbf{x}) = 1, f_j(\mathbf{x}) = x_j, j = 1, \dots, p - 1$) and unknown scalar coefficients β_j .

The covariance function \mathbf{k} is a map that is symmetric positive semi-definite, usually stationary $\mathbf{k}(\mathbf{x}, \mathbf{x}') = \mathbf{r}(\mathbf{x} - \mathbf{x}')$. The most commonly used kernel in \mathbb{R} is the Matérn kernel class given by

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

for $x, y \in \mathbb{R}$. Here $\sigma^2 > 0$ is the amplitude, $\theta > 0$ is the length-scale, Γ is the complete Gamma function and K_ν is the modified Bessel function of the second kind. (σ^2, θ) are called hyperparameters. Some particular cases of Matérn kernel are when $\nu = \frac{1}{2}$ (Exponential), $\nu = \frac{3}{2}$ (Matérn 3/2), $\nu = \frac{5}{2}$ (Matérn 5/2) and $\nu \rightarrow \infty$ (Gaussian or Squared-Exponential).

The choice of kernels is important in the kriging scheme and requires prior knowledge of the smoothness of the function f . For example, the choice of the Gaussian kernel assumes that the function is very smooth of class \mathcal{C}^∞ (infinitely differentiable) which is often too strict as a condition. A common alternative is the functions Matérn 5/2 or Matérn 3/2 kernel

It is possible to build high-dimensional covariance models in \mathbb{R}^d based on classical kernels in \mathbb{R} . In particular, the Matérn anisotropic geometric model (radial model), which we consider in the following in this paper, defined as

$$\mathbf{k}_{\sigma^2, \theta}^{\text{radial}}(\mathbf{x}, \mathbf{x}') = \mathbf{r}_{\sigma^2, \theta}^\nu \left(\sqrt{\sum_{j=1}^d \frac{|x_j - x'_j|^2}{\theta_j^2}} \right), \quad (4)$$

where \mathbf{r} is a Matérn kernel \mathbb{R} as defined in (3) and $\theta = (\theta_1, \dots, \theta_d)$ the length-scale vector. The described method can be applied to other forms of covariance models like the tensorized product model with d -dimensional kernels (as the product of kernels is also a kernel) or the Power-Exponential model. In the following sections, instead of writing $\mathbf{k}_{\sigma^2, \theta}$, we denote simply \mathbf{k} when there is no possible confusion.

3.2. Gaussian process regression model

The *prior* distribution of Y on the learning experimental design \mathbf{X} is multivariate Gaussian

$$\mathbf{y} \mid \boldsymbol{\beta}, \sigma^2, \boldsymbol{\theta}, \sigma_\epsilon^2 \sim \mathcal{N}(\mathbf{F}\boldsymbol{\beta}, \mathbf{K}), \quad (5)$$

where

- $\mathbf{F} = (F_{ij}) \in \mathbb{R}^{n \times p}$ is the regression matrix such that $F_{ij} = f_j(\mathbf{x}^{(i)})$.
- $\boldsymbol{\beta} = \{\beta_1, \dots, \beta_p\} \in \mathbb{R}^p$ are the regression coefficients.
- $\mathbf{K} = (\mathbf{k}(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}))_{1 \leq i, j \leq n} + \sigma_\epsilon^2 \mathbf{I}_n \in \mathbb{R}^{n \times n}$ is the covariance matrix of the learning design \mathbf{X} .

Hypothesis \mathcal{H}_1 : In the case of ordinary or universal kriging, we assume that $n \geq p$, \mathbf{F} is a full rank matrix, and $\mathbf{e} \in \text{Im } \mathbf{F}$ where $\mathbf{e} = (1, \dots, 1)^\top$.

Remark 1. In Ordinary Kriging, the hypothesis \mathcal{H}_1 is always satisfied. In the Universal Kriging, the hypothesis $\mathbf{e} \in \text{Im } \mathbf{F}$ is satisfied as soon as the constant function $f_0(\mathbf{x}) = C$ is included in the chosen family of functions f_i .

3.3. Prediction

The Gaussian conditioning theorem is useful to deduce the *posterior* distribution. By considering a new point \mathbf{x}_{new} , it can be shown that the predictive distribution of $Y(\mathbf{x}_{\text{new}})$ conditioned on the learning sample \mathbf{X}, \mathbf{y} is also Gaussian

$$Y(\mathbf{x}_{\text{new}}) \mid \mathbf{X}, \mathbf{y}, \sigma^2, \boldsymbol{\theta}, \sigma_\epsilon^2 \sim \mathcal{N}\left(\tilde{y}(\mathbf{x}_{\text{new}}), \tilde{\sigma}^2(\mathbf{x}_{\text{new}})\right), \tag{6}$$

where, in the case of Ordinary or Universal Kriging and by denoting $f_{\text{trend}}(\mathbf{x}_{\text{new}}) = (f_j(\mathbf{x}_{\text{new}}))_{j=0}^{p-1}$, $y(\mathbf{x}_{\text{new}})$ and $\tilde{\sigma}^2(\mathbf{x}_{\text{new}})$ are given by the Best Linear Unbiased Predictor (BLUP),

$$\tilde{y}_{\sigma^2, \boldsymbol{\theta}, \sigma_\epsilon^2}(\mathbf{x}_{\text{new}}) = f_{\text{trend}}(\mathbf{x}_{\text{new}})^\top \hat{\boldsymbol{\beta}} + \mathbf{k}(\mathbf{x}_{\text{new}}, \mathbf{X})^\top \mathbf{K}^{-1}(\mathbf{y} - \mathbf{F}\hat{\boldsymbol{\beta}}), \tag{7}$$

$$\begin{aligned} \tilde{\sigma}_{\sigma^2, \boldsymbol{\theta}, \sigma_\epsilon^2}^2(\mathbf{x}_{\text{new}}) &= \mathbf{k}(\mathbf{x}_{\text{new}}, \mathbf{x}_{\text{new}}) + \sigma_\epsilon^2 - \mathbf{k}(\mathbf{x}_{\text{new}}, \mathbf{X})^\top \mathbf{K}^{-1} \mathbf{k}(\mathbf{x}_{\text{new}}, \mathbf{X}) + (f_{\text{trend}}(\mathbf{x}_{\text{new}}) - \\ &\quad \mathbf{F}\mathbf{K}^{-1} \mathbf{k}(\mathbf{x}_{\text{new}}, \mathbf{X}))^\top (\mathbf{F}^\top \mathbf{K}^{-1} \mathbf{F})^{-1} (f_{\text{trend}}(\mathbf{x}_{\text{new}}) - \mathbf{F}\mathbf{K}^{-1} \mathbf{k}(\mathbf{x}_{\text{new}}, \mathbf{X})), \end{aligned} \tag{8}$$

and,

$$\hat{\boldsymbol{\beta}} = (\mathbf{F}^\top \mathbf{K}^{-1} \mathbf{F})^{-1} \mathbf{F}^\top \mathbf{K}^{-1} \mathbf{y}. \tag{9}$$

We refer to Santner et al. (2003) in Chapter 4 for a detailed proof of Equations (6)–(9). In particular, we note that the additional term of the predictive variance in (8) is due to the propagation of the non-informative improper form of prior distribution on the estimation of $\boldsymbol{\beta}$.

In the following and when there is no possible confusion, $\tilde{y}_{\sigma^2, \boldsymbol{\theta}, \sigma_\epsilon^2}$ (resp. $\tilde{\sigma}_{\sigma^2, \boldsymbol{\theta}, \sigma_\epsilon^2}^2$) will be also denoted by \tilde{y} (resp. $\tilde{\sigma}^2$) without specifying its dependence on hyperparameters or the nugget effect.

The most outstanding advantage of GP model compared to other models relies on the previous equations (7) and (8). The predictive distribution can be used in a wide variety of applications such as uncertainty quantification, sensitivity analysis (Oakley and O’Hagan, 2004) and quantile evaluation. Other possible considerations and extensions of GP modelling are described in Currin et al. (1991) and Rasmussen and Williams (2005).

Given a GP regression model and a point $\mathbf{x}_{\text{new}} \in \mathcal{D}$, the *posterior* distribution of prediction in (6) can be standardized into

$$\tilde{Z}(\mathbf{x}_{\text{new}}) = \frac{Y(\mathbf{x}_{\text{new}}) - \tilde{y}(\mathbf{x}_{\text{new}})}{\tilde{\sigma}(\mathbf{x}_{\text{new}})} \mid \mathbf{X}, \mathbf{y}, \sigma^2, \boldsymbol{\theta}, \sigma_\epsilon^2 \sim \mathcal{N}(0, 1). \tag{10}$$

By considering the standardized variable $\tilde{Z}(\mathbf{x}_{\text{new}})$, the α -quantiles z_α are those of the standard normal distribution: $q_{1-\alpha/2} = \Phi^{-1}(1 - \alpha/2)$ and $q_{\alpha/2} = \Phi^{-1}(\alpha/2) = -q_{1-\alpha/2}$ where Φ is the CDF of the standard normal distribution, such that the Prediction Intervals $\mathcal{PI}_{1-\alpha}$ in (2) can be written as

$$\mathcal{PI}_{1-\alpha}(\mathbf{x}_{\text{new}}) = [\tilde{y}(\mathbf{x}_{\text{new}}) - q_{1-\alpha/2} \times \tilde{\sigma}(\mathbf{x}_{\text{new}}); \tilde{y}(\mathbf{x}_{\text{new}}) + q_{1-\alpha/2} \times \tilde{\sigma}(\mathbf{x}_{\text{new}})], \tag{11}$$

which gives a natural definition for $y_{\alpha/2}$ and $y_{1-\alpha/2}$

$$y_{\alpha/2}(\mathbf{x}) = \tilde{y}(\mathbf{x}) - q_{1-\alpha/2} \times \tilde{\sigma}(\mathbf{x}); \quad y_{1-\alpha/2}(\mathbf{x}) = \tilde{y}(\mathbf{x}) + q_{1-\alpha/2} \times \tilde{\sigma}(\mathbf{x}). \tag{12}$$

3.4. Training model with maximum likelihood method

Constructing a GP model and computing the kriging mean and variance as shown in (7) and (8) implies estimating the nugget effect σ_ϵ^2 and the covariance parameters $(\sigma^2, \boldsymbol{\theta})$. Here, we assume that σ_ϵ^2 is known or has been estimated by the method proposed in looss and Marrel (2017) for instance.

The Maximum Likelihood Estimator (MLE) $\hat{\sigma}_{ML}^2$ and $\hat{\boldsymbol{\theta}}_{ML}$ of σ^2 and $\boldsymbol{\theta}$ is given by (Santner et al., 2003)

$$(\hat{\sigma}_{ML}^2, \hat{\boldsymbol{\theta}}_{ML}) \in \operatorname{argmin}_{\sigma^2, \boldsymbol{\theta}} \mathbf{y}^\top \left(\mathbf{K}^{-1} - \mathbf{K}^{-1} \mathbf{F} (\mathbf{F}^\top \mathbf{K}^{-1} \mathbf{F})^{-1} \mathbf{F}^\top \mathbf{K}^{-1} \right) \mathbf{y} + \log(\det \mathbf{K}). \tag{13}$$

The MLE method is optimal when the covariance function is well-specified (Bachoc, 2013a) (i.e. when the data \mathbf{y} comes from a function f that is a realization of a GP with covariance function that belongs to the family of covariance functions in section 3.1).

However, there is no guarantee that the MLE method would perform optimally as this method is poorly robust with respect to model misspecifications. Besides, training and assessing the quality of a predictor should not be done on the same data (Hastie et al. (2009) in chapter 7). In particular, the MLE method does not show how well the model will do when it is asked to make new predictions for data it has not already seen. The CV method represents an alternative to estimate the covariance hyperparameters $(\sigma^2, \boldsymbol{\theta})$ for prediction purposes (Zhang and Wang, 2010; Bachoc, 2013a) and has the advantage of being more efficient and robust when the covariance function is misspecified (Bachoc, 2013a).

3.5. Training model with cross-validation method for point-wise prediction

We consider the same learning set of n observations $\mathbf{D}_{\text{learn}} = (\mathbf{X}, \mathbf{y}) = \{(\mathbf{x}^{(i)}, y^{(i)}), i \in \{1, \dots, n\}\}$ and we assume that the value of $\sigma_\epsilon^2 \in [0, +\infty)$ is known. The Leave-One-Out method (i.e. n -Cross-Validation) consists in predicting $y^{(i)}$ by building a GP model, denoted \mathcal{GP}_{-i} and trained on $\mathbf{D}_{-i} = \{(\mathbf{x}^{(j)}, y^{(j)})\}_{j \in \{1, \dots, n\} \setminus \{i\}}$. The obtained prediction mean and variance are functions of parameters $(\sigma^2, \boldsymbol{\theta})$ as shown in (7) and (8) and are used to assess the predictive capability of the global GP model.

In the case of the Leave-One-Out method, the Mean Squared prediction Error (MSE) is used to assess the quality of the point-wise prediction (see Wallach and Goffinet (1989) for more details about this metric) of the GP model, it can be expressed as

$$\mathcal{LOO}_{MSE}(\sigma^2, \boldsymbol{\theta}) := \frac{1}{n} \sum_{i=1}^n (y^{(i)} - \tilde{y}_i)^2, \tag{14}$$

where \tilde{y}_i and $\tilde{\sigma}_i^2$ are the Leave-One-Out predictive mean and variance of $f(\mathbf{x}^{(i)})$ by a GP model trained on \mathbf{D}_{-i} with the hyperparameters $(\sigma^2, \boldsymbol{\theta})$.

Hypothesis \mathcal{H}_2 : Let $(\mathbf{e}_i)_{i=1}^n$ be the canonical basis of \mathbb{R}^n . We assume that $\mathbf{e}_i \notin \text{Im } \mathbf{F}$ for all $i \in \{1, \dots, n\}$.

Let $\bar{\mathbf{K}}$ be the matrix defined by

$$\bar{\mathbf{K}} = \mathbf{K}^{-1} - \mathbf{K}^{-1} \mathbf{F} (\mathbf{F}^\top \mathbf{K}^{-1} \mathbf{F})^{-1} \mathbf{F}^\top \mathbf{K}^{-1}. \tag{15}$$

For all $i \in \{1, \dots, n\}$, we have $(\bar{\mathbf{K}})_{i,i} > 0$ by Lemma 3 (see Appendix A), and, in the case of Ordinary or Universal Kriging, the Virtual Cross-Validation formulas (Dubrule, 1983) of the predictive mean \tilde{y}_i and variance $\tilde{\sigma}_i^2$ are given by

$$y^{(i)} - \tilde{y}_i = \frac{(\bar{\mathbf{K}}\mathbf{y})_i}{(\bar{\mathbf{K}})_{i,i}}, \tag{16}$$

and

$$\tilde{\sigma}_i^2 = \frac{1}{(\bar{\mathbf{K}})_{i,i}}. \tag{17}$$

With the presence of the nugget effect, the GP regressor does not interpolate the training data \mathbf{y} but approximates them as best as possible. The Leave-One-Out method looks for the best approximation by minimizing the \mathcal{LOO}_{MSE} criterion. The criterion (14) can be written with explicit quadratic forms in \mathbf{y} and therefore

$$(\hat{\sigma}_{MSE}^2, \hat{\boldsymbol{\theta}}_{MSE}) \in \text{argmin}_{\sigma^2, \boldsymbol{\theta}} \mathbf{y}^\top \bar{\mathbf{K}} \text{Diag}(\bar{\mathbf{K}})^{-2} \bar{\mathbf{K}} \mathbf{y}. \tag{18}$$

Note that in the absence of the nugget effect $\sigma_\epsilon^2 = 0$, $\bar{\mathbf{K}}$ is of the form $\sigma^{-2} \bar{\mathbf{R}}_\theta$ where $\bar{\mathbf{R}}_\theta$ does not depend on σ^2 . The predictive variance $\hat{\sigma}_{MSE}^2$ can then be computed by the following explicit quadratic form (Bachoc, 2013a)

$$\hat{\sigma}_{MSE}^2 = \frac{1}{n} \mathbf{y}^\top \bar{\mathbf{R}}_{\hat{\boldsymbol{\theta}}_{MSE}} \text{Diag}(\bar{\mathbf{R}}_{\hat{\boldsymbol{\theta}}_{MSE}})^{-1} \bar{\mathbf{R}}_{\hat{\boldsymbol{\theta}}_{MSE}} \mathbf{y}, \tag{19}$$

and the optimal length-scale vector $\hat{\boldsymbol{\theta}}_{MSE}$ is obtained by solving

$$\hat{\boldsymbol{\theta}}_{MSE} \in \text{argmin}_{\boldsymbol{\theta}} \mathbf{y}^\top \bar{\mathbf{R}}_\theta \text{Diag}(\bar{\mathbf{R}}_\theta)^{-2} \bar{\mathbf{R}}_\theta \mathbf{y}. \tag{20}$$

3.6. Full-Bayesian approach

In this subsection, we consider the full-Bayesian treatment of GP models (Williams and Barber, 1998). Indeed, the full-Bayesian approach integrates the uncertainty about the unknown hyperparameters and assumes a *prior* on the hyperparameters $(\sigma^2, \boldsymbol{\theta}) \sim p(\sigma^2, \boldsymbol{\theta})$. Consequently, the probability density function (pdf) of the *posterior* predictive distribution of $Y(\mathbf{x}_{\text{new}})$ at a new point \mathbf{x}_{new} can be expressed as an integral over the hyperparameters (we omit the conditioning over inputs \mathbf{X} and \mathbf{x}_{new}):

$$p(y_{\text{new}} | \mathbf{y}) = \iint p(y_{\text{new}} | \mathbf{y}, \sigma^2, \boldsymbol{\theta}) p(\sigma^2, \boldsymbol{\theta} | \mathbf{y}) d\sigma^2 d\boldsymbol{\theta}, \tag{21}$$

where $p(y_{\text{new}} | \sigma^2, \boldsymbol{\theta})$ is the pdf of $Y(\mathbf{x}_{\text{new}})$ given \mathbf{y}, σ^2 and $\boldsymbol{\theta}$, and $p(\sigma^2, \boldsymbol{\theta} | \mathbf{y}) \propto p(\mathbf{y} | \sigma^2, \boldsymbol{\theta}) p(\sigma^2, \boldsymbol{\theta})$ is the hyperparameters' posterior distribution.

The implementation of the full-Bayesian approach requires the evaluation of the previous integral and the posterior $p(\sigma^2, \theta | \mathbf{y})$, which cannot be computed directly. It is common to use Markov chain Monte Carlo (MCMC) methods for sampling and inference from the posterior distribution of the hyperparameters to overcome this issue, using, in particular, the Metropolis-Hastings (MH) algorithm (Robert and Casella, 2004) or Hamiltonian Monte Carlo (HMC) (Neal, 1993, 1996).

Therefore, the predictive distribution is obtained by Monte Carlo

$$p(y_{\text{new}} | \mathbf{y}) \simeq \frac{1}{N} \sum_{i=1}^N p(y_{\text{new}} | \mathbf{y}, \sigma_i^2, \theta_i), \tag{22}$$

where N denotes the MCMC sample size and (σ_i^2, θ_i) is the i -th sample drawn from the posterior distribution $p(\sigma^2, \theta | \mathbf{y})$.

Finally, one can draw a sample $(Y_i(\mathbf{x}_{\text{new}}))_{i=1}^N$ of $Y(\mathbf{x}_{\text{new}})$ following the posterior distribution $p(y_{\text{new}} | \sigma_i^2, \theta_i)$ as in (6) for each $i = 1, \dots, N$ and build the Prediction Intervals $\mathcal{PI}_{1-\alpha}$ by taking the empirical quantiles of order $\alpha/2$ and $1 - \alpha/2$ of the sample $(Y_i(\mathbf{x}_{\text{new}}))_{i=1}^N$.

Note that the *plug-in* approaches (e.g. the MLE method in 3.4) consider (21) and replace $p(\sigma^2, \theta | \mathbf{y})$ by a Dirac distribution centered on a value such as $(\hat{\sigma}_{ML}^2, \hat{\theta}_{ML})$ that maximizes the likelihood function.

4. Prediction intervals estimation for Gaussian processes

Using the Cross-Validation method, the MSE hyperparameters $(\hat{\sigma}_{MSE}^2, \hat{\theta}_{MSE})$ are obtained from a point-wise prediction metric and do not focus on Prediction Intervals neither on quantifying the uncertainty of the model. For these purposes, using the CP is more appropriate.

The *Coverage Probability* (CP) is defined as the probability that the Prediction Interval procedure will produce an interval that captures what it is intended to capture (Hong et al., 2009). In the Leave-One-Out framework, we keep the notations of \tilde{y}_i and $\tilde{\sigma}_i^2$: the predictive mean and variance on $\mathbf{x}^{(i)} \in \mathbf{X}$ using the learning set $\mathbf{D}_{-i} = \{(\mathbf{x}^{(j)}, y^{(j)})\}_{j \in \{1, \dots, n\} \setminus \{i\}}$. We define then the Leave-One-Out CP $\tilde{\mathbb{P}}_{1-\alpha}$ as the percentage of observed values \mathbf{y} belonging to Prediction Intervals $\mathcal{PI}_{1-\alpha}$ of \tilde{y}_i for all $i \in \{1, \dots, n\}$

$$\begin{aligned} \tilde{\mathbb{P}}_{1-\alpha} &= \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{y^{(i)} \in \mathcal{PI}_{1-\alpha}(\mathbf{x}^{(i)})\}, \\ &= \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{\tilde{y}_i + q_{\alpha/2} \times \tilde{\sigma}_i < y^{(i)} \leq \tilde{y}_i + q_{1-\alpha/2} \times \tilde{\sigma}_i\}, \end{aligned} \tag{23}$$

where q_a is the a -quantile of the standard normal distribution and $\mathbf{1}\{A\}$ is the indicator function of A . We introduce the Heaviside step function h

$$h(x) = \mathbf{1}\{x \geq 0\} = \begin{cases} 1 & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases} \tag{24}$$

The Leave-One-Out CP $\tilde{\mathbb{P}}_{1-\alpha}$ can be written as

$$\tilde{\mathbb{P}}_{1-\alpha} = \frac{1}{n} \sum_{i=1}^n h\left(q_{1-\alpha/2} - \frac{y^{(i)} - \tilde{y}_i}{\tilde{\sigma}_i}\right) - \frac{1}{n} \sum_{i=1}^n h\left(q_{\alpha/2} - \frac{y^{(i)} - \tilde{y}_i}{\tilde{\sigma}_i}\right). \tag{25}$$

When the model is well-specified, the coverage of the Prediction Intervals $\mathcal{PI}_{1-\alpha}$ is optimal as the predictive distribution is fully characterized by the Gaussian distribution (see section 3.1), each term of the right-hand side of (25) is an unbiased estimator of the probability

$$\mathbb{P}\left(\frac{Y(\mathbf{x}^{(i)}) - \tilde{y}_i}{\tilde{\sigma}_i} \leq q_{1-\alpha/2} \mid \mathbf{D}_{-i}\right) = 1 - \alpha/2, \tag{26}$$

and

$$\mathbb{P}\left(\frac{Y(\mathbf{x}^{(i)}) - \tilde{y}_i}{\tilde{\sigma}_i} \leq q_{\alpha/2} \mid \mathbf{D}_{-i}\right) = \alpha/2. \tag{27}$$

Conversely, if the model is misspecified, each predictive quantile, needs to be quantified properly with respect to the normal distribution quantile so that the CP as described in section 2 achieves the desired level.

Let $a \in (0, 1/2) \cup (1/2, 1)$ describe a nominal level of quantile. We define the *quasi-Gaussian* proportion ψ_a as a map from $[0, +\infty) \times (0, +\infty)^d$ to $[0, 1]$

$$\psi_a(\sigma^2, \boldsymbol{\theta}) = \frac{1}{n} \sum_{i=1}^n h\left(q_a - \frac{y^{(i)} - \tilde{y}_i}{\tilde{\sigma}_i}\right), \tag{28}$$

where \tilde{y}_i and $\tilde{\sigma}_i$ are the predictive mean and variance at $\mathbf{x}^{(i)}$ using the learning set \mathbf{D}_{-i} and the hyperparameters $(\sigma^2, \boldsymbol{\theta})$. Given the Virtual Cross-Validation formulas (Dubrule, 1983), ψ_a can be written in terms of the covariance matrix $\bar{\mathbf{K}}$

$$\psi_a(\sigma^2, \boldsymbol{\theta}) = \frac{1}{n} \sum_{i=1}^n h\left(q_a - \frac{(\bar{\mathbf{K}}\mathbf{y})_i}{\sqrt{(\bar{\mathbf{K}})_{i,i}}}\right). \tag{29}$$

The *quasi-Gaussian* proportion ψ_a describes how close the a -quantile q_a of the standardized predictive distribution is to the level a (ideally, it should correspond to a). Therefore, the objective is to fit the hyperparameters $(\sigma^2, \boldsymbol{\theta})$ according to the *quasi-Gaussian* proportions and to find two pairs $(\bar{\sigma}^2, \bar{\boldsymbol{\theta}})$ and $(\underline{\sigma}^2, \underline{\boldsymbol{\theta}})$ such that $\psi_{1-\alpha/2}(\bar{\sigma}^2, \bar{\boldsymbol{\theta}}) = 1 - \alpha/2$ and $\psi_{\alpha/2}(\underline{\sigma}^2, \underline{\boldsymbol{\theta}}) = \alpha/2$. This allows us to get the optimal Leave-One-Out CP by respecting the nominal confidence level $(1 - \alpha)$, that is $\mathbb{P}_{1-\alpha} = 1 - \alpha$.

4.1. Presence of nugget effect

In this section, we assume $\sigma_\epsilon^2 > 0$. The *quasi-Gaussian* proportion ψ_a is, however, piecewise constant and can take values only in the finite set $\{k/n, k \in \{0, \dots, n\}\}$. We first need to modify the problem $\psi_a(\sigma^2, \boldsymbol{\theta}) = a$. Let $\delta > 0$, we define the continuous functions h_δ^- and h_δ^+

$$h_\delta^+(x) = \begin{cases} 1 & \text{if } x > \delta \\ x/\delta & \text{if } 0 < x \leq \delta \\ 0 & \text{otherwise} \end{cases} \tag{30}$$

$$h_\delta^-(x) = \begin{cases} 1 & \text{if } x \geq 0 \\ 1 + x/\delta & \text{if } -\delta \leq x < 0 \\ 0 & \text{otherwise} \end{cases}$$

If $a > 1/2$ we define

$$\psi_a^{(\delta)}(\sigma^2, \boldsymbol{\theta}) = \frac{1}{n} \sum_{i=1}^n h_\delta^+\left(q_a - \frac{(\bar{\mathbf{K}}\mathbf{y})_i}{\sqrt{(\bar{\mathbf{K}})_{i,i}}}\right). \tag{31}$$

If $a < 1/2$ we define

$$\psi_a^{(\delta)}(\sigma^2, \boldsymbol{\theta}) = \frac{1}{n} \sum_{i=1}^n h_\delta^-\left(q_a - \frac{(\bar{\mathbf{K}}\mathbf{y})_i}{\sqrt{(\bar{\mathbf{K}})_{i,i}}}\right). \tag{32}$$

Let $\delta > 0$ be a small enough so that $\delta < q_a$ if $a > 1/2$ (respectively, $\delta < q_{1-a}$ if $a < 1/2$) in such a way that $h_\delta^+(q_a) = 1$ (respectively, $h_\delta^-(q_a) = 0$). We consider the problem

$$\psi_a^{(\delta)}(\sigma^2, \boldsymbol{\theta}) = a, \tag{33}$$

and we denote by $\mathcal{A}_{a,\delta}$ the solution set of the problem (33)

$$\mathcal{A}_{a,\delta} := \left\{ (\sigma^2, \boldsymbol{\theta}) \in [0, +\infty) \times (0, +\infty)^d, \psi_a^{(\delta)}(\sigma^2, \boldsymbol{\theta}) = a \right\}. \tag{34}$$

Hypothesis \mathcal{H}_3 : Let $k_\epsilon = \text{Card}\{i \in \{1, \dots, n\}, \frac{(\boldsymbol{\Pi}\mathbf{y})_i}{\sqrt{(\boldsymbol{\Pi})_{ii}}} \leq \sigma_\epsilon q_a\}$ where $\boldsymbol{\Pi}$ is the orthogonal projection matrix on $(\text{Im } \mathbf{F})^\perp$ such that $\boldsymbol{\Pi} = \mathbf{I}_n - \mathbf{F}(\mathbf{F}^\top \mathbf{F})^{-1} \mathbf{F}^\top$. We assume that $k_\epsilon < na$ if $a > 1/2$ and $k_\epsilon > na$ if $a < 1/2$.

Remark 2. The hypothesis \mathcal{H}_3 is typically satisfied in Ordinary and Universal Kriging. Indeed, $\boldsymbol{\Pi}$ is the projection on the space $(\text{Im } \mathbf{F})^\perp$ and is expected to remove the trend of the model. It is reasonable to think that $(\boldsymbol{\Pi}\mathbf{y})$ is centered and that

$$\text{Card}\{i \in \{1, \dots, n\}, (\boldsymbol{\Pi}\mathbf{y})_i \leq 0\} \approx \frac{n}{2}. \tag{35}$$

If σ_ϵ^2 is smaller than σ^2 , then we should also have

$$\text{Card}\{i \in \{1, \dots, n\}, \frac{(\mathbf{\Pi} \mathbf{y})_i}{\sqrt{(\mathbf{\Pi})_{ii}}} \leq \sigma_\epsilon q_a\} \approx \frac{n}{2}, \tag{36}$$

so that the hypothesis \mathcal{H}_3 should be fulfilled.

Proposition 1. *Let us assume the hypotheses $\mathcal{H}_1, \mathcal{H}_2$ and \mathcal{H}_3 , then $\mathcal{A}_{a,\delta}$ is non-empty.*

Proof. In Appendix A. \square

The challenge now is to identify and choose wisely the optimal solutions $(\sigma_{\text{opt}}^2, \boldsymbol{\theta}_{\text{opt}}) \in \mathcal{A}_{a,\delta}$. Some authors (Khosravi et al., 2010) suggest the mean Prediction Intervals width (MPIW) of Prediction Intervals $\mathcal{PI}_{1-\alpha}$ as an additional constraint to reduce the set of solutions. However, this constraint may not work when dealing with quantile estimation because the lower bound of the corresponding interval may be infinite.

Instead, we will compare these solutions with MLE's solution $(\hat{\sigma}_{ML}^2, \hat{\boldsymbol{\theta}}_{ML})$ (subsection 3.4) or MSE-CV solution $(\hat{\sigma}_{MSE}^2, \hat{\boldsymbol{\theta}}_{MSE})$ (subsection 3.5) and we will take the closest pair $(\sigma_{\text{opt}}^2, \boldsymbol{\theta}_{\text{opt}})$ by using an appropriate notion of similarity between multivariate Gaussian distributions. Ideally, we aim to solve the following problem

$$\text{argmin}_{(\sigma^2, \boldsymbol{\theta}) \in \mathcal{A}_{a,\delta}} d^2 \left((\sigma^2, \boldsymbol{\theta}), (\sigma_0^2, \boldsymbol{\theta}_0) \right), \tag{37}$$

where d is a continuous similarity measure of hyperparameters $(\sigma^2, \boldsymbol{\theta})$ operating on the mean vector \mathbf{m} and covariance matrix \mathbf{K} , and $(\sigma_0^2, \boldsymbol{\theta}_0) = (\hat{\sigma}_{ML}^2, \hat{\boldsymbol{\theta}}_{ML})$ or $(\hat{\sigma}_{MSE}^2, \hat{\boldsymbol{\theta}}_{MSE})$ as described in (13) or (18).

The resolution of the problem (37) may be too costly and heavy to solve when the dimension is high, say $d \geq 10$. An alternative is to apply *the relaxation* method where we redefine this optimization problem of $\boldsymbol{\theta}$ from $(0, +\infty)^d$ to $(0, +\infty)$ by shifting the length-scale vector $\boldsymbol{\theta}_0$ by a parameter $\lambda \in (0, +\infty)$.

Let $\boldsymbol{\theta}_0$ denote a solution of the problems (13) or (18) and for $\lambda \in (0, +\infty)$, let $H_\delta(\lambda)$ denote the subset

$$H_\delta(\lambda) = \{\sigma^2 \in [0, +\infty), \psi_a^{(\delta)}(\sigma^2, \lambda \boldsymbol{\theta}_0) = a\}. \tag{38}$$

Hypothesis \mathcal{H}_4 : The set-valued mapping (the so-called correspondence function) $H_\delta : (0, +\infty) \rightarrow \mathcal{P}((0, +\infty))$, where $\mathcal{P}(S)$ denotes the power set of a set S , is lower semi-continuous, that is, for all $\lambda \in (0, +\infty)$, for each open set \mathcal{U} with $H_\delta(\lambda) \cap \mathcal{U} \neq \emptyset$, there exists a neighborhood $\mathcal{O}(\lambda)$ such that if $\lambda^* \in \mathcal{O}(\lambda)$ then $H_\delta(\lambda^*) \cap \mathcal{U} \neq \emptyset$.

In the kriging framework, σ^2 should be as small as possible to reduce the uncertainty of the model, a natural choice of σ_{opt}^2 is

$$\forall \lambda \in (0, +\infty) : \sigma_{\text{opt}}^2(\lambda) := \min\{\sigma^2 \in [0, +\infty), \psi_a^{(\delta)}(\sigma^2, \lambda \boldsymbol{\theta}_0) = a\}. \tag{39}$$

Proposition 2. *The function $\lambda \mapsto \sigma_{\text{opt}}^2(\lambda)$ is well-defined under hypotheses \mathcal{H}_1 to \mathcal{H}_3 , and continuous on $(0, +\infty)$ under the additional hypothesis \mathcal{H}_4 .*

Proof. In Appendix A. \square

Concerning the choice of d , one known similarity measure between probability distributions is the Wasserstein distance, widely used in optimal transportation problems (see Chapter 6 of Villani (2009) for more details). In case of two Gaussian random distributions $\mathcal{N}(\mathbf{m}_1, \mathbf{K}_1)$ and $\mathcal{N}(\mathbf{m}_2, \mathbf{K}_2)$, the second Wasserstein distance is equal to

$$W_2^2(\mathcal{N}(\mathbf{m}_1, \mathbf{K}_1), \mathcal{N}(\mathbf{m}_2, \mathbf{K}_2)) = \|\mathbf{m}_1 - \mathbf{m}_2\|^2 + \text{Tr} \left(\mathbf{K}_1 + \mathbf{K}_2 - 2\sqrt{\mathbf{K}_1^{1/2} \mathbf{K}_2 \mathbf{K}_1^{1/2}} \right), \tag{40}$$

where, in our setting, $\mathbf{m}_1 = \mathbf{F}\hat{\boldsymbol{\beta}}_1 = (\mathbf{F}^\top \mathbf{K}_1^{-1} \mathbf{F})^{-1} \mathbf{F}^\top \mathbf{K}_1^{-1} \mathbf{y}$ and $\mathbf{m}_2 = \mathbf{F}\hat{\boldsymbol{\beta}}_2 = (\mathbf{F}^\top \mathbf{K}_2^{-1} \mathbf{F})^{-1} \mathbf{F}^\top \mathbf{K}_2^{-1} \mathbf{y}$.

Therefore, each pair $(\sigma^2, \boldsymbol{\theta})$ is associated to a Gaussian distribution $\mathcal{N}(\mathbf{m}, \mathbf{K})$ and we define the similarity measure d as

$$d^2 \left((\sigma^2, \boldsymbol{\theta}), (\sigma_0^2, \boldsymbol{\theta}_0) \right) = W_2^2(\mathcal{N}(\mathbf{m}, \mathbf{K}), \mathcal{N}(\mathbf{m}_0, \mathbf{K}_0)). \tag{41}$$

The choice of the second Wasserstein distance d^2 and σ_{opt}^2 makes the Prediction Intervals $\mathcal{PI}_{1-\alpha}$ shorter without the need for an additional metric like the MPIW and without modifying the distribution of the obtained model significantly. We will see in Section 5.2 that, empirically, the barycenters of Prediction Intervals are not far from the predictive means obtained by MLE or MSE-CV methods.

The relaxed optimization problem in (33) for the quantile estimation is given by the problem \mathcal{P}_λ

$$\mathcal{P}_\lambda : \text{argmin}_{\lambda \in (0, +\infty)} \mathcal{L}(\lambda) := d^2 \left((\sigma_{\text{opt}}^2(\lambda), \lambda \boldsymbol{\theta}_0), (\sigma_0^2, \boldsymbol{\theta}_0) \right). \tag{42}$$




Fig. 1. Illustration of the relaxation effect on the ECDF of the Leave-One-Out standardized predictive distribution on the quantile of level $a = 90\%$; The relaxed standardized predictive distribution coincides with the standard normal distribution on point $(q_a, a) = (1.28, 0.90)$ instead of $(\psi_a, a) = (0.78, 0.90)$. Green: standard normal distribution - MLE standardized predictive distribution when the model is well-specified; Orange: MLE standardized predictive distribution when the model is misspecified; Blue: standardized predictive distribution after relaxing model's hyperparameters. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

Proposition 3. Under hypotheses \mathcal{H}_1 to \mathcal{H}_4 , the function $\mathcal{L} : (0, +\infty) \rightarrow \mathbb{R}^+$ is continuous and coercive on $(0, +\infty)$. The problem \mathcal{P}_λ admits at least one global minimizer λ^* in $(0, +\infty)$.

Proof. See Appendix A. \square

Remark 3. The coercivity of the function \mathcal{L} is guaranteed by the hypotheses \mathcal{H}_1 to \mathcal{H}_3 (see Appendix A). The function \mathcal{L} is also upper semi-continuous (Zhao, 1997). The hypothesis \mathcal{H}_4 insures that \mathcal{L} is continuous and that there exists a global minimizer. This hypothesis is not easy to check. If it does not hold or if it cannot be checked, then it is possible to solve the problem (42) on a regular grid by a grid search method.

Let $\hat{\beta}_{\text{opt}}$ denote the corresponding regression parameter

$$\hat{\beta}_{\text{opt}} = \left(\mathbf{F}^\top \mathbf{K}_{\sigma_{\text{opt}}^2(\lambda^*), \lambda^* \theta_0}^{-1} \mathbf{F} \right)^{-1} \mathbf{F}^\top \mathbf{K}_{\sigma_{\text{opt}}^2(\lambda^*), \lambda^* \theta_0}^{-1} \mathbf{y}. \tag{43}$$

The purpose of this resolution is to create a GP model with hyperparameters $(\hat{\beta}_{\text{opt}}, \sigma_{\text{opt}}^2(\lambda^*), \lambda^* \theta_0)$ able to predict the quantile \tilde{y}_a such that a proportion a of true values are below \tilde{y}_a with respect to the constraint of quasi-Gaussian proportion ψ_a (see Fig. 1). Finally, the Prediction Intervals $\mathcal{PI}_{1-\alpha}$ will be obtained using two GP models built with the same method, one for the upper quantile $1 - \alpha/2$ with optimal relaxation parameter $\bar{\lambda}^*$ and the other for the lower quantile of $\alpha/2$ with parameter $\underline{\lambda}^*$. The CP of $\mathcal{PI}_{1-\alpha}$ is optimal and insured by respecting the coverage of each quantile as shown in (25). In the following, we call this method Robust Prediction Intervals Estimation (RPIE).

Remark 4. It is clear that the GP hyperparameters selected by the RPIE method depends on the level a . Given the continuity properties of the different steps of the RPIE method, one may expect, however, that the hyperparameters selected for a specific level a should also give good CP locally for other levels of coverage a' close to a . Nevertheless, this local property is certainly not global. This sensitivity can also be related to the known observation that guarantees for conditional coverage are more challenging to obtain than for marginal coverage (Barber et al., 2020).

4.2. Absence of nugget effect

When the nugget effect is null $\sigma_\epsilon^2 = 0$, the set of solutions $\mathcal{A}_{a,\delta}$ is still non-empty because one can show that, for θ in the neighborhood of $\mathbf{0} \in \mathbb{R}^d$, the problem $\psi_a^{(\delta)}(\sigma^2, \theta) = a$ has a solution $\sigma^2 \in (0, +\infty)$ (see Appendix B). In particular, the correspondence function H_δ is non-empty valued for $\lambda > 0$ small enough and it may be empty-valued for some large $\lambda \in (0, +\infty)$. We may think, however, that H_δ is non-empty valued and that $\sigma_{\text{opt}}^2(\lambda)$ exists for λ close to one. Indeed, assume

for a while that the model is well-specified, that is, there exist hyperparameters $(\beta_*, \sigma_*^2, \theta_*)$ such that \mathbf{y} corresponds to a realization of a random vector $\mathbf{Y} \sim \mathcal{N}(\mathbf{F}\beta_*, \sigma_*^2 \mathbf{R}_{\theta_*})$. The existence of $H_\delta(\lambda)$ and $\sigma_{\text{opt}}^2(\lambda)$ depend on the condition $k_\lambda \leq na$, where k_λ is the integer defined by

$$k_\lambda := \text{Card} \{i \in \{1, \dots, n\}, (\bar{\mathbf{R}}_\lambda \theta_0 \mathbf{y})_i \leq 0\}. \tag{44}$$

Since $\bar{\mathbf{R}}_{\theta_*} \mathbf{Y}$ is centered, we can anticipate that

$$\text{Card} \{i \in \{1, \dots, n\}, (\bar{\mathbf{R}}_{\theta_*} \mathbf{y})_i \leq 0\} \approx \frac{n}{2}. \tag{45}$$

Hence, the condition $n/2 < k_\lambda \leq na$ should be satisfied in a neighborhood of $\lambda = 1$ since θ_0 should be close to θ_* . Finally, even though the function \mathcal{L} is not defined on $(0, +\infty)$, we can solve (42) by a grid search method on its domain.

5. Numerical results

5.1. Test cases with analytical functions

In this section, we give three numerical examples to illustrate Prediction Intervals estimation by the RPIE method. We show that for the *Wing-Weight* function, the model is well-specified as the CP is optimal for different levels, hence, no robust calibration of Prediction Intervals is required. However, for Zhou (1998) and Morokoff and Caflisch (1995) functions where the model is misspecified and for a given confidence level α , we apply the RPIE method as described in section 4 to estimate both upper and lower bounds of Predictions Intervals. The following metrics: the Leave-One-Out CP $\tilde{\mathbb{P}}_{1-\alpha}$ defined in (25), the Coverage Probability (CP), the mean (MPIW) and standard-deviation (SdPIW) of the Prediction Interval width, and the accuracy Q^2 (Kleijnen and Sargent, 2000) are used to assess and compare GP models built by MLE or MSE-CV methods, full Bayesian approach or the RPIE method. They can be used either for point-wise prediction comparison (Q^2 will be given in some cases for information, it does not represent the main metric of this section):

$$Q^2 = 1 - \frac{\sum_{i=1}^{n_{\text{test}}} (y_{\text{test}}^{(i)} - \tilde{y}_{i,\text{test}})^2}{\sum_{i=1}^{n_{\text{test}}} (y_{\text{test}}^{(i)} - \bar{y})^2}, \tag{46}$$

or for quantifying the goodness of Prediction Intervals:

$$\tilde{\mathbb{P}}_{1-\alpha} = \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{y^{(i)} \in \mathcal{PI}_{1-\alpha}(\mathbf{x}^{(i)})\}, \tag{47}$$

$$\text{CP}_{1-\alpha} = \frac{1}{n_{\text{test}}} \sum_{i=1}^{n_{\text{test}}} \mathbf{1}\{y_{\text{test}}^{(i)} \in \mathcal{PI}_{1-\alpha}(\mathbf{x}_{\text{test}}^{(i)})\}, \tag{48}$$

$$\text{MPIW}_{1-\alpha} = \frac{1}{n_{\text{test}}} \sum_{i=1}^{n_{\text{test}}} |\mathcal{PI}_{1-\alpha}(\mathbf{x}_{\text{test}}^{(i)})|, \tag{49}$$

and,

$$\text{SdPIW}_{1-\alpha} = \sqrt{\frac{1}{n_{\text{test}}} \sum_{i=1}^{n_{\text{test}}} [|\mathcal{PI}_{1-\alpha}(\mathbf{x}_{\text{test}}^{(i)})| - \text{MPIW}_{1-\alpha}]^2}, \tag{50}$$

where $\mathbf{y}_{\text{test}} = (y_{\text{test}}^{(1)}, \dots, y_{\text{test}}^{(n_{\text{test}})})$ is the vector to predict at $(\mathbf{x}_{\text{test}}^{(1)}, \dots, \mathbf{x}_{\text{test}}^{(n_{\text{test}})})$, $\mathcal{PI}_{1-\alpha}$ is the $(1 - \alpha) \times 100\%$ confidence Prediction Interval delimited by the quantiles $q_{1-\alpha/2}$ and $q_{\alpha/2}$, and $|\mathcal{PI}_{1-\alpha}|$ is the length of the interval.

Note that the $\text{CP}_{1-\alpha}$ may be different from the Leave-One-Out CP $\tilde{\mathbb{P}}_{1-\alpha}$, this case can happen when the distributions of the training and testing sets are different. However, a Leave-One-Out CP $\tilde{\mathbb{P}}_{1-\alpha}$ close to $1 - \alpha$ insures that, if the assumption of *i.i.d.* distributions is respected i.e. $p_{\text{train}}(\mathbf{X}, \mathbf{y}) = p_{\text{test}}(\mathbf{X}, \mathbf{y})$, $\text{CP}_{1-\alpha}$ should be also close to $1 - \alpha$.

This subsection provides results obtained on $d = 10$ -dimensional GP with constant mean function (Ordinary Kriging). The value of δ is fixed at $\delta = 10^{-2}$. We implement our methods using the package *kergp* (Roustant et al., 2020) on R. For the computational time, we use an Intel(R) Core(TM) i5-9400H CPU @ 2.50 GHz with a RAM of 32 Go.

Table 1
The input variables x_j and their domain ranges $[a_j; b_j]$.

Component	Domain	Component	Domain
x_1	[150; 200]	x_6	[0.5; 1]
x_2	[220; 300]	x_7	[0.08; 0.18]
x_3	[6, 10]	x_8	[2.5; 6]
x_4	[-10; 10]	x_9	[1700, 2500]
x_5	[16; 45]	x_{10}	[0.025; 0.08]

Table 2
Performances of methods (MLE, MSE-CV and Full-Bayesian) for Wing Weight function.

	Before RPIE		After RPIE		Full-Bayesian
	MLE	MSE-CV	MLE	MSE-CV	
Q^2	0.563	0.764	n.c	n.c	0.562
$\tilde{P}_{99\%}$	99.1	99.8	98.9	98.9	99.1
CP _{99%}	98.7	100	98.7	98.0	98.7
$\tilde{P}_{95\%}$	94.0	98.9	94.9	94.9	94.2
CP _{95%}	95.3	99.3	96.7	96.0	95.3
$\tilde{P}_{90\%}$	90.1	96.9	90.0	90.0	90.9
CP _{90%}	91.3	96.0	89.3	90.0	91.3
Ct	2 min 12 s	32 min 42 s	6 min*	37 min*	4 h 39 min 27 s

Q^2 : Accuracy; $\tilde{P}_{1-\alpha}$: The Leave-One-Out CP in % on the training set; CP_{1- α} : The CP in % on the testing set and Ct: computational time.

* The approximated cumulative computational time after running the RPIE method for all levels.

Example 1: well-specified model - the wing weight function

The Wing Weight function is a model in dimension $d = 10$ proposed by Forrester et al. (2008) that estimates the weight of a light aircraft wing. For an input vector $\mathbf{x} \in \mathbb{R}^{10}$, the response y is:

$$f(\mathbf{x}) = 0.036x_1^{0.758}x_2^{0.0035} \left(\frac{x_3}{\cos^2(x_4)} \right)^{0.6} x_5^{0.006}x_6^{0.04} \left(\frac{100 x_7}{\cos(x_4)} \right)^{-0.3} (x_8x_9)^{0.49} + x_1x_{10}. \tag{51}$$

The components x_i are assumed to vary over the ranges given in Table 1 (see Forrester et al. (2008) and Moon (2010) for details).

We create an experimental design \mathbf{X} of $n = 600$ observations and $d = 10$ variables where observations $\mathbf{x}^{(i)} = (x_1^{(i)}, \dots, x_d^{(i)})$ are sampled i.i.d. with uniform distribution over $\otimes_{j=1}^d [a_j, b_j]$. We generate the response $\mathbf{y} = (y^{(1)}, \dots, y^{(n)})$ such that $y^{(i)} = f(\mathbf{x}^{(i)}) + \epsilon^{(i)}$ with f defined in (51) and $\epsilon^{(i)}$ are sampled i.i.d. with the distribution $\mathcal{N}(0, \sigma_\epsilon^2 = 25)$. Here the nugget effect is estimated with the methodology described in Iooss and Marrel (2017) and the covariance kernel is the Matérn 3/2.

The GP model is trained on 75% of the data (25% of data is left for testing). The diagnostics of the model are presented in Table 2 with the metrics described above. The accuracy Q^2 is moderate for MLE and Full-Bayesian methods. The MSE-CV does much better, an expected result since the MSE-CV method is more adapted for point-wise prediction criterion. However, the Leave-One-Out CP $\tilde{P}_{1-\alpha}$ for two different levels $\alpha = 5\%, 10\%$ is far from the required level, which means that they were poorly estimated with point-wise prediction criterion. In addition, Table 2 shows in particular that the model is well-specified for Matérn 3/2 correlation kernel with the MLE method since the CPs are optimal and close to the required level. This claim is empirical and can be verified either by comparing the standardized predictive distribution with the standard normal distribution as in Fig. 1 or using Shapiro and Wilk (1965) normality test (in this example, p -value = 0.203). The Full-Bayesian approach also does well in estimating Prediction Intervals in the case of a well-specified model. Indeed, the hyperparameters' posterior distribution $p(\sigma^2, \theta | \mathbf{y})$ is concentrated around the MLE estimator, so the plug-in MLE approach and the Full-Bayesian approach give similar predictive distributions and Prediction Intervals. However, its computational time is extremely long compared to other methods (e.g. 100 times longer than the MLE method). Concerning the RPIE method, one can notice that it provides the optimal coverage at each required level, either on training or testing sets. However, we do not see significant interest in applying it here (except for the MSE-CV solution).

Example 1 is a case of well-specified model in which the CPs obtained by the MLE method satisfy the nominal value and the RPIE method does not bring a significant additional value (at least for the MLE solution).

Table 3
Performances of methods before and after RPIE for Morokoff & Caflisch function; here $1 - \alpha = 90\%$.

	Before RPIE		After RPIE		Full-Bayesian
	MLE	MSE-CV	MLE	MSE-CV	
Q^2	0.892	0.895	n.c	n.c	0.891
$\tilde{\mathbb{P}}_{1-\alpha}$	93.6	98.3	90.0	90.0	93.8
$CP_{1-\alpha}$	94.0	98.0	92.6	87.3	93.3
$MPIW_{1-\alpha}$	$1.68 \cdot 10^{-1}$	$1.81 \cdot 10^{-1}$	$5.51 \cdot 10^{-2}$	$5.78 \cdot 10^{-2}$	$1.66 \cdot 10^{-1}$
$SdPIW_{1-\alpha}$	$9.61 \cdot 10^{-3}$	$4.16 \cdot 10^{-2}$	$1.29 \cdot 10^{-2}$	$1.41 \cdot 10^{-2}$	$9.27 \cdot 10^{-3}$
Ct	1 min 16 s	24 min 18 s	3 min 55 s	27 min 43 s	4 h 43 min 38 s

Q^2 : Accuracy; $\tilde{\mathbb{P}}_{1-\alpha}$: The Leave-One-Out CP in % on the training set; $CP_{1-\alpha}$: CP in % on the testing set; MPIW: Mean of Prediction Interval widths; SdPIW: standard deviation of Prediction Interval widths and Ct: computational time.

Example 2: misspecified model with noise - Morokoff & Caflisch function -

We consider the Morokoff and Caflisch (1995) function defined on $[0, 1]^d$ by

$$f(\mathbf{x}) = \frac{1}{2} \left(1 + \frac{1}{d}\right)^d \prod_{i=1}^d (x_i)^{1/d}. \tag{52}$$

In Example 2, we consider an experimental design \mathbf{X} of $n = 600$ observations and $d = 10$ correlated inputs. Each observation has the form $\mathbf{x}^{(i)} = (\Phi(z_1^{(i)}), \dots, \Phi(z_d^{(i)})) \in \mathbb{R}^d$, Φ is the CDF of the standard normal distribution, $z^{(i)}$ are sampled from the multivariate distribution $\mathcal{N}(\mathbf{0}, \mathbf{C})$ and $\mathbf{C} \in \mathbb{R}^{d \times d}$ is the following covariance matrix:

$$\mathbf{C} = \begin{bmatrix} 1 & 0.90 & 0 & 0 & 0 & 0.50 & -0.30 & 0 & 0 & 0 \\ 0.90 & 1 & 0 & 0 & 0 & 0 & 0 & 0.10 & 0 & 0 \\ 0 & 0 & 1 & 0 & -0.30 & 0.10 & 0.40 & 0 & 0.05 & 0 \\ 0 & 0 & 0 & 1 & 0.40 & 0 & 0 & -0.35 & 0 & 0 \\ 0 & 0 & -0.30 & 0.40 & 1 & 0 & 0 & 0 & 0.10 & 0 \\ 0.05 & 0 & 0.10 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ -0.30 & 0 & 0.40 & 0 & 0 & 0 & 1 & 0 & 0 & -0.30 \\ 0 & 0.1 & 0 & -0.35 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0.05 & 0 & 0.10 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -0.3 & 0 & 0 & 1 \end{bmatrix}.$$

The response vector \mathbf{y} is generated as $y^{(i)} = f(\mathbf{x}^{(i)}) + \epsilon^{(i)}$ with f the Morokoff & Caflisch function defined in (52) and $\epsilon^{(i)}$ are sampled i.i.d. with the distribution $\mathcal{N}(0, \sigma_\epsilon^2 = 10^{-4})$. We consider the Matérn anisotropic geometric correlation model with smoothness $5/2$ as covariance model and we study the Prediction Interval’s problem with a nugget effect estimated with the methodology looss and Marrel (2017).

The model is not well-specified as Example 1 and the Shapiro and Wilk (1965) test gives p -value = $1.253 \cdot 10^{-7}$. Table 3 summarizes the results of MLE and MSE-CV estimations before and after applying the RPIE, compared with the Full-Bayesian approach. The accuracy Q^2 of both models is satisfactory and is slightly improved when using the MSE-CV method. However, before applying the RPIE, the Prediction Intervals are overestimated for both models. The CP does not correspond to the required level of 90%, and the MSE-CV model performs even worse. We note that the Full-Bayesian approach does not improve the quality of estimated Prediction Intervals for the same reason as explained before: the hyperparameters’ posterior distribution $p(\sigma^2, \boldsymbol{\theta} \mid \mathbf{y})$ is concentrated around the MLE estimator and the performances of both approaches are similar. We will see that this claim is also valid in Example 3.

We now address the problem of Prediction Intervals Estimation for each solution of MLE $\hat{\boldsymbol{\theta}}_{ML}$ and MSE-CV $\hat{\boldsymbol{\theta}}_{MSE}$. We consider the upper and lower bounds $1 - \alpha/2 = 95\%$ and $\alpha/2 = 5\%$ and we apply the RPIE method as described in section 4. The optimal values $\bar{\lambda}^*$ and $\underline{\lambda}^*$ obtained from the resolution of the problem (42) are used to build two GP models to estimate each bound. Fig. 2 shows the variation of the function \mathcal{L} for Morokoff & Caflisch example while solving the problem (42) on the upper bound $1 - \alpha/2 = 95\%$, it illustrates the statement of Proposition 3: \mathcal{L} is continuous and coercive on $(0, +\infty)$ and reaches a global minimum.

We consider now the Prediction Intervals built according to the RPIE method. In Table 3, one observes that these Prediction Intervals are three times shorter than those of MLE, MSE-CV models or Full-Bayesian approaches and have appropriate variances (e.g. more heterogeneous than MLE or Full-Bayesian method’s Prediction Intervals). The coverage rate of $1 - \alpha = 90\%$ on the training set is achieved, which is the main objective of the RPIE method, and the CP on the testing set is very close to this level. Concerning the computational time, it appears that applying the RPIE method to MLE or MSE-CV solutions counts for a short computational time (only a few minutes to run in this example). The Full-Bayesian approach is still computationally heavy, as already discussed in the previous example and section 1.




Fig. 2. The variation of the relaxed Wasserstein distance \mathcal{L} for Morokoff & Caflisch function; $a = 1 - \alpha/2 = 95\%$.

Table 4
Performances of methods for Zhou function (53) in the first setting ($\sigma_\epsilon^2 = 0$); here $1 - \alpha = 90\%$.

	Before RPIE		After RPIE		Full-Bayesian
	MLE	MSE-CV	MLE	MSE-CV	
Q^2	0.947	0.947	n.c	n.c	0.948
$\tilde{\mathbb{P}}_{1-\alpha}$	92.0	42.1	90.0	90.0	92.0
$CP_{1-\alpha}$	92.7	45.3	90.0	88.0	92.9
$MPIW_{1-\alpha}$	$4.60 \cdot 10^{-1}$	$1.46 \cdot 10^{-1}$	$4.35 \cdot 10^{-1}$	$4.32 \cdot 10^{-1}$	$4.59 \cdot 10^{-1}$
$SdPIW_{1-\alpha}$	$1.06 \cdot 10^{-1}$	$3.48 \cdot 10^{-2}$	$1.00 \cdot 10^{-1}$	$1.00 \cdot 10^{-1}$	$1.08 \cdot 10^{-1}$
Ct	10 s	31 min 2 s	2 min 31 s	33 min 32 s	4 h 56 min 15 s

Q^2 : Accuracy; $\tilde{\mathbb{P}}_{1-\alpha}$: The Leave-One-Out CP in % on the training set; $CP_{1-\alpha}$: The CP in % on the testing set; MPIW: Mean of Prediction Interval widths; SdPIW: standard deviation of Prediction Interval widths and Ct: computational time.

Example 2 is a case of misspecified model with noise in which the CP obtained by MLE, MSE-CV and Full-Bayesian methods are not good. The RPIE method fulfills its purpose: its reduces Prediction Intervals width and improves the robustness of Prediction Intervals in such a way that they achieve the optimal coverage rate.

Example 3: misspecified model without noise - Zhou function -

The Zhou (1998) function, considered initially for the numerical integration of spiky functions, is defined on $[0, 1]^d$ by

$$f(\mathbf{x}) = \frac{10^d}{2} \left[\phi\left(10\left(\mathbf{x} - \frac{1}{3}\right)\right) + \phi\left(10\left(\mathbf{x} - \frac{2}{3}\right)\right) \right], \tag{53}$$

where

$$\phi(\mathbf{x}) = (2\pi)^{-d/2} \exp\left(-0.5\|\mathbf{x}\|^2\right). \tag{54}$$

In Example 3, we create an experimental design \mathbf{X} similar to Example 1, containing $n = 600$ and $d = 10$ variables where observations $\mathbf{x}^{(i)} = (x_1^{(i)}, \dots, x_d^{(i)})$ are sampled independently with uniform distribution over $[0, 1]^d$. As the Zhou function in (53) takes some high values, we generate the response \mathbf{y} by applying a logarithmic transformation:

$$y^{(i)} = \log f(\mathbf{x}^{(i)}) / (d \log 10). \tag{55}$$

Note that there is no measurement noise here. We will address two situations: In the first setting, we assume that we know that there is no measurement noise, we impose that there is no nugget effect in the model $\sigma_\epsilon^2 = 0$ and we consider the Exponential anisotropic geometric correlation model ($\nu = 1/2$) as covariance model. In the second setting, we assume that we do not know whether there is measurement noise and we estimate the nugget effect of the model. We consider consequently the Matérn 3/2 anisotropic geometric correlation model ($\nu = 3/2$), a reasonable choice for a smooth covariance model when assuming a nugget effect (see Appendix B for further discussion).

In Table 4, the models are good in terms of accuracy Q^2 with a small advantage for the Full-Bayesian approach, but none of them satisfies the required level of CP, especially the MSE-CV model with an extremely low CP. As we do not estimate the nugget effect in this setting, the computational time of the MLE method is low (a few seconds) where the RPIE still takes a couple of minutes, as in Example 2. We will notice (also in the industrial application) that the computational time after the RPIE method is generally twice to three times the computational time of MLE method when there is a nugget effect.

When proceeding similarly as Example 2 to build robust Prediction Intervals by the RPIE model, the result is striking in Table 4: The estimated Prediction Intervals for the MSE-CV solution $\hat{\theta}_{MSE}$ after RPIE are now four times larger, meaning

Table 5
Performances of methods for Zhou function (53) in the second setting ($\hat{\sigma}_\epsilon^2 = 1.71 \cdot 10^{-2}$); here $1 - \alpha = 90\%$.

	Before RPIE		After RPIE		Full-Bayesian
	MLE	MSE-CV	MLE	MSE-CV	
Q^2	0.941	0.944	n.c	n.c	0.941
$\tilde{\mathbb{P}}_{1-\alpha}$	99.4	100	90.0	90.0	99.3
$CP_{1-\alpha}$	99.3	100	92.0	85.3	99.6
$MPIW_{1-\alpha}$	$6.48 \cdot 10^{-1}$	1.19	$2.26 \cdot 10^{-1}$	$2.28 \cdot 10^{-1}$	$6.56 \cdot 10^{-1}$
$SdPIW_{1-\alpha}$	$6.88 \cdot 10^{-2}$	$2.56 \cdot 10^{-1}$	$4.73 \cdot 10^{-2}$	$5.27 \cdot 10^{-2}$	$6.97 \cdot 10^{-2}$
Ct	1 min 20 s	31 min 22 s	3 min 39 s	33 min 37 s	4 h 25 min 59 s

Q^2 : Accuracy; $\tilde{\mathbb{P}}_{1-\alpha}$: The Leave-One-Out CP in % on the training set; $CP_{1-\alpha}$: The CP in % on the testing set; MPIW: Mean of Prediction Interval widths; SdPIW: standard deviation of Prediction Interval widths and Ct: computational time.

that the amplitude $\hat{\sigma}_{MSE}^2$ was largely underestimated. Table 4 also shows that the CPs for the testing set are close to their desired value $1 - \alpha = 90\%$.

In the second setting, the nugget effect is estimated to $\hat{\sigma}_\epsilon^2 = 1.71 \cdot 10^{-2}$ by using looss and Marrel (2017). The results of MLE, MSE-CV and Full-Bayesian methods are shown in Table 5. The accuracy is still satisfying and similar to the previous setting, but the CP is close to 100%, meaning that the Prediction Intervals of all three methods are overestimated. Table 5 shows that, with the RPIE method, we reduce Prediction Intervals width, five times shorter than Prediction Intervals of the MSE-CV solution, and three shorter than Prediction Intervals of the MLE solution. The variances of the obtained Prediction Intervals are between MLE and MSE-CV Prediction Intervals variances. One can notice also a decrease of 50% of the MPIW compared to the first setting, while maintaining an optimal coverage of $1 - \alpha = 90\%$.

Example 3 illustrates a case of misspecified model without noise where the RPIE method adjusts Prediction Intervals width and improves the robustness of Prediction Intervals so that the CP is respected. One can also conclude that it is preferable to consider a nugget effect for shorter Prediction Intervals and optimal coverage.

5.2. Application to Gas production for future wells

In this section, we illustrate the interest of the RPIE method in energy production forecasting. It includes many industrial applications such as battery capacity, wind turbine, solar panel performance or, more specifically, unconventional gas wells where a decline in production may be observed. We show that the RPIE can estimate robust Prediction Intervals, covering the lower bounds of level $\alpha/2 = 10\%$ (pessimistic scenario) and the upper bounds of level $1 - \alpha/2 = 90\%$ (optimistic scenario).

Indeed, a fundamental challenge of Oil and Gas companies is to predict their assets and their production capacities in the future. It drives both their exploration and development strategy. However, forecasting a well future production is challenging because subsurface reservoirs properties are never fully known. This makes estimating well production with their associated uncertainty a crucial task. The agencies Securities and Exchange Commission and Society of Petroleum Engineers define specific rules **1P/2P/3P**, known as Petroleum Reserves and Resources Definitions (PRMS), for reserves estimates based on quantile estimates:

- **1P**: 90% of wells produce more than **1P** predictions (**proven**).
- **2P**: 50% of wells produce more than **2P** predictions (**probable**).
- **3P**: 10% of wells produce more than **3P** predictions (**possible**).

These rules are to be disclosed to security investors for publicly traded Oil and Gas companies and aim to provide investors with consistent information and associated value assessments. Many Machine Learning algorithms have shown their efficiency in estimating the median **2P** (e.g. using GP with MLE method, or MSE-CV if interested more in point-wise predictions) but failed to estimate **1P** and **3P**. Thus, the objective of this study is to build a proper estimation of the quantiles $p_{90\%}$ and $p_{10\%}$ by applying the RPIE method described in section 4.

Our dataset, *field data*, is derived from unconventional wells localized in the Utica shale reservoir, located in the north-east of the United States. It contains approximately $n = 1850$ wells and $d = 12$ variables, including localization, Cumulative Production of natural gas over 12 months in MCFE, completion design and exploitation conditions. The raw dataset can be found at the Ohio Oil & Gas well locator of the Ohio Department of Natural Resources (2022).

We standardized the data (\mathbf{X}, \mathbf{y}) and we divided into a 60% – 20% – 20% partition of three datasets: a training set and two validation sets. The response \mathbf{y} (Cumulative Production over 12 months in MCFE) is noisy due to the uncertainty of the reservoir parameters in the field. The nugget effect σ_ϵ^2 is unknown but estimated to $\hat{\sigma}_\epsilon^2 = 0.16$ using the method of looss and Marrel (2017).

Based on results drawn from the previous subsection and for practical reasons (particularly the computational cost of methods), we will present only the application of the RPIE method on the MLE solution. Table 6 shows the performances

Table 6
Results obtained for GP model, Random Forest and Gradient Boosting; here $1 - \alpha = 80\%$.

	MLE	Random Forest	XGBoost
Q^2	0.872	0.870	0.885
$CP_{1-\alpha}$	92.8	98.1	49.8
$MPIW_{1-\alpha}$	1.18	1.52	0.48
$SdPIW_{1-\alpha}$	0.21	0.29	0.22
Ct	14 min 37 s	2 s	1 min 36 s

Q^2 : Accuracy; CP: The CP in % on validation set I; MPIW: Mean of Prediction Interval widths; SdPIW: standard deviation of Prediction Interval widths and Ct: computational time.

Table 7
Obtained results before and after RPIE method; here $1 - \alpha = 80\%$.

	MLE before RPIE	MLE after RPIE
$\tilde{P}_{1-\alpha}$	90.9	79.9
$CP_{1-\alpha}^{Val,1}$	92.6	81.0
$MPIW_{1-\alpha}^{Val,1}$	1.18	1.06
$SdPIW_{1-\alpha}^{Val,1}$	$2.09 \cdot 10^{-1}$	$8.25 \cdot 10^{-3}$
$CP_{1-\alpha}^{Val,2}$	94.1	83.2
$MPIW_{1-\alpha}^{Val,2}$	1.17	1.06
$SdPIW_{1-\alpha}^{Val,2}$	$1.68 \cdot 10^{-1}$	$7.00 \cdot 10^{-3}$
Ct	14 min 37 s	59 min 25 s

$CP_{1-\alpha}^{Val,1}$ (resp. $CP_{1-\alpha}^{Val,2}$): The CP in % on Validation set I (resp. Validation set II); $MPIW_{1-\alpha}^{Val,1}$ (resp. $MPIW_{1-\alpha}^{Val,2}$): Mean of Prediction Interval widths on Validation set I (resp. Validation set II); $SdPIW_{1-\alpha}^{Val,1}$ (resp. $SdPIW_{1-\alpha}^{Val,2}$): standard deviation of Prediction Interval widths on Validation set I (resp. Validation set II) and Ct: computational time.

of the GP model trained by MLE compared with two other statistical models: Random Forest and Gradient Boosting whose Prediction Intervals are estimated using the Bootstrap method. Here we consider the Prediction Intervals of level $1 - \alpha = 80\%$: the lower bound is the 10% quantile ($p_{10\%}$) and the upper bound the 90% quantile ($p_{90\%}$) of the predictive distribution.

The accuracy of the MLE model is 0.873 and has approximately the same accuracy as other models like Random Forest or Gradient Boosting. Furthermore, the CP of the Prediction Intervals of $1 - \alpha = 80\%$ is not satisfactory but it is quite *reasonable* for MLE model compared to Random Forest (overestimated Prediction Intervals) or Gradient Boosting (underestimated Prediction Intervals). Finally, it appears that the GP model requires some computing resources to be built and to estimate its hyperparameters by MLE method.

In the following, we define the MLE's solution as reference $\theta_0 = \hat{\theta}_{ML}$ in the optimization problem (42) for the quantiles $\alpha/2 = 10\%$ and $1 - \alpha/2 = 90\%$ and we build robust Prediction Intervals confidence level $1 - \alpha = 80\%$ with the RPIE method. The results are presented in Table 7. When considering the estimated Prediction Intervals by the RPIE method, we can see the CP is optimal for the training set and is close to $1 - \alpha = 80\%$ for both validation sets. Therefore, we fulfill the objective of estimating the upper and lower bounds, the obtained quantiles $p_{90\%}$ and $p_{10\%}$ respect **1P** and **3P** rules as mentioned above. Finally, in Fig. 3a, we present the estimated Prediction Intervals defined by the upper bounds $P90$ and lower bounds $P10$ against the true values of \mathbf{y} on Validation set I. The x-axis designs well's indices ordered with respect to the barycenters of the Prediction Intervals (engineers choose this representation for interpretation purposes). We can see that the estimated Prediction Intervals by the MLE method are not homogeneous, and some of them are longer. The RPIE method makes them shorter and more homogeneous as it can be seen in Fig. 3b, and in the evolution of the standard deviation width SdPIW in Table 7.

In a second attempt and following the engineers' recommendation, we consider a logarithmic transformation to the raw response \mathbf{y} to avoid having non-positive lower bounds and integrate heterogeneity between performant and less performant well. The accuracy of the MLE method decreases now to $Q^2 = 0.615$, the MLE method still over estimates Prediction Intervals as it can be seen in Table 8. Most claims of the previous analysis remain true, in particular we can clearly see (also in Figs. 3c and 3d) that Prediction Intervals obtained by RPIE are shorter and have reduced standard-deviations.

6. Conclusion

In this paper, we have introduced a new approach for Prediction Intervals estimation based on the Cross-Validation method. We use the Gaussian Processes model because the predictive distribution at a new point is completely characterized by Gaussian distribution. We address an optimization problem for model's hyperparameters estimation by considering the

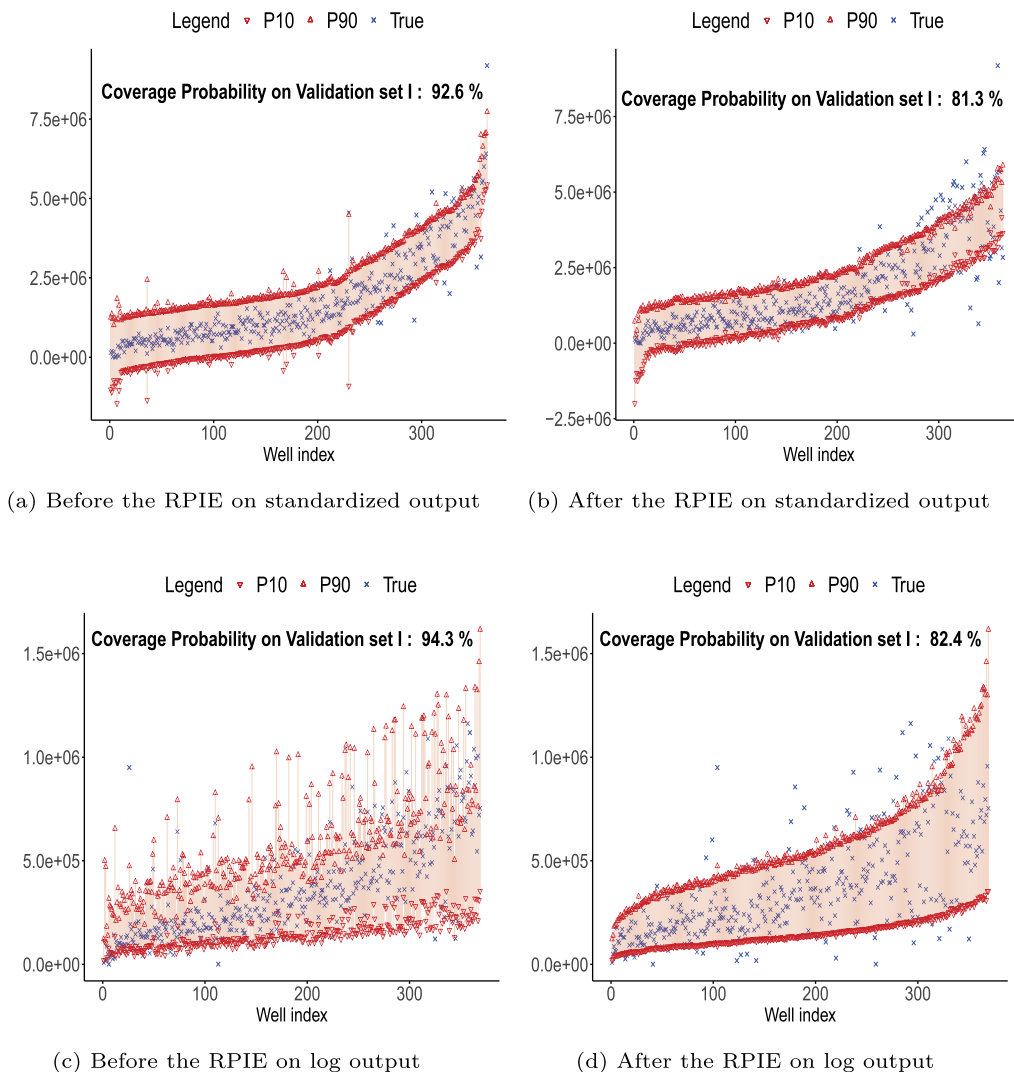


Fig. 3. Production data after re-scaling: True values vs 80% confidence Prediction Intervals.

Table 8

Obtained results before and after RPIE method; $1 - \alpha = 80\%$. Here the output data are log-transformed.

	MLE before RPIE	MLE after RPIE
$\hat{P}_{1-\alpha}$	91.1	79.9
$CP_{1-\alpha}^{Val,1}$	94.3	83.2
$MPIW_{1-\alpha}^{Val,1}$	1.53	1.40
$SdPIW_{1-\alpha}^{Val,1}$	$2.20 \cdot 10^{-1}$	$1.40 \cdot 10^{-2}$
$CP_{1-\alpha}^{Val,2}$	90.4	76.6
$MPIW_{1-\alpha}^{Val,2}$	1.54	1.40
$SdPIW_{1-\alpha}^{Val,2}$	$1.92 \cdot 10^{-1}$	$1.42 \cdot 10^{-2}$
Ct	17 min 47 s	53 min 21 s

$CP_{1-\alpha}^{Val,1}$ (resp. $CP_{1-\alpha}^{Val,2}$): The CP in % on Validation set I (resp. Validation set II); $MPIW_{1-\alpha}^{Val,1}$ (resp. $MPIW_{1-\alpha}^{Val,2}$): Mean of Prediction Interval widths on Validation set I (resp. Validation set II); $SdPIW_{1-\alpha}^{Val,1}$ (resp. $SdPIW_{1-\alpha}^{Val,2}$): standard deviation of Prediction Interval widths on Validation set I (resp. Validation set II) and Ct: computational time.

notion of Coverage Probability. The optimal hyperparameters are identified by minimizing the Wasserstein distance between the Gaussian distribution with the hyperparameters determined by Cross-Validation, and the Gaussian distribution with hyperparameters achieving the desired Coverage Probability. This method is relevant when the model is misspecified. It insures an optimal Leave-One-Out Coverage Probability for the training set. It also achieves a reasonable Coverage Probability for the validation set when it is available. It can be also extended to other statistical models with a predictive distribution but more detailed work is needed to consider the influence of hyperparameters on Prediction Interval's coverage and solve the optimization problem more efficiently in these cases. Finally, it should be possible to include categorical inputs in the covariance function by using group kernels (Roustant et al., 2020), which would extend the application range of the RPIE method.

Acknowledgements

The author⁽¹⁾ would like to thank Achraf Ourir, Zinyat Agharzayeva (TotalEnergies SE - La Défense, France), Daniel Busby (TotalEnergies SE, CSTJF - Pau, France) and the research laboratory SINCLAIR (IA Commun Lab - Saclay, France) for useful discussions and helpful suggestions. This work was supported by TotalEnergies and the French National Agency for Research and Technology 2019/0714 (ANRT) (Grant n°2019/0714).

Appendix A. Proofs of Propositions 1–3

A.1. Preliminary lemmas

Lemma 1. Let \mathbf{F} be a full rank matrix (hypothesis \mathcal{H}_1), let \mathbf{K} be a positive definite matrix and let $\bar{\mathbf{K}}$ defined by $\bar{\mathbf{K}} = \mathbf{K}^{-1}(\mathbf{I}_n - \mathbf{F}(\mathbf{F}^\top \mathbf{K}^{-1} \mathbf{F})^{-1} \mathbf{F}^\top \mathbf{K}^{-1})$ then $\text{Ker } \bar{\mathbf{K}} = \text{Im } \mathbf{F}$ and $\bar{\mathbf{K}}$ is singular.

Proof. Let $\bar{\mathbf{K}}$ be the matrix defined above. Suppose that $\mathbf{x} \in \text{Im } \mathbf{F}$, then there exists \mathbf{y} such that $\mathbf{x} = \mathbf{F}\mathbf{y}$, and $\bar{\mathbf{K}}\mathbf{x} = \mathbf{K}^{-1}(\mathbf{F}\mathbf{y} - \mathbf{F}(\mathbf{F}^\top \mathbf{K}^{-1} \mathbf{F})^{-1} \mathbf{F}^\top \mathbf{K}^{-1} \mathbf{F}\mathbf{y}) = \mathbf{K}^{-1}(\mathbf{F}\mathbf{y} - \mathbf{F}\mathbf{y}) = \mathbf{0}$. Thus $\mathbf{x} \in \text{Ker } \bar{\mathbf{K}}$.

If $\mathbf{x} \in \text{Ker } \bar{\mathbf{K}}$, then $\mathbf{K}\bar{\mathbf{K}}\mathbf{x} = \mathbf{0}$, and $\mathbf{x} = \mathbf{F}(\mathbf{F}^\top \mathbf{K}^{-1} \mathbf{F})^{-1} \mathbf{F}^\top \mathbf{K}^{-1} \mathbf{x} = \mathbf{F}\mathbf{x}' \in \text{Im } \mathbf{F}$.

In case of Ordinary or Universal kriging, $p = \text{rank}(\mathbf{F}) = \text{dim}(\text{Ker } \bar{\mathbf{K}}) \geq 1$ which means that $\bar{\mathbf{K}}$ is not invertible. \square

Lemma 2 (De Oliveira (2007)). Under the hypotheses of Lemma 1 and given the full rank regression matrix \mathbf{F} , there exists a matrix $\mathbf{W} \in \mathbb{R}^{n \times (n-p)}$ satisfying:

$$\mathbf{W}^\top \mathbf{W} = \mathbf{I}_{n-p}, \tag{A.1}$$

$$\mathbf{F}^\top \mathbf{W} = \mathbf{0}_{p \times (n-p)}, \tag{A.2}$$

and

$$\bar{\mathbf{K}} = \mathbf{W}(\mathbf{W}^\top \mathbf{K} \mathbf{W})^{-1} \mathbf{W}^\top. \tag{A.3}$$

Lemma 3. Under the hypotheses of Lemma 1, if additionally hypothesis \mathcal{H}_2 holds true, then $\bar{\mathbf{K}}_{ii} > 0$ for all $i \in \{1, \dots, n\}$.

Proof. $\bar{\mathbf{K}}$ is a positive semi-definite matrix by Lemma 2 and we can write

$$\bar{\mathbf{K}} = \sum_{j=1}^n \lambda_j \mathbf{u}_j \mathbf{u}_j^\top, \tag{A.4}$$

with $\lambda_j \geq 0$ the eigenvalues of $\bar{\mathbf{K}}$ and $(\mathbf{u}_j)_{j=1}^n$ the orthonormal basis of the eigenvectors. We have

$$\bar{\mathbf{K}}_{ii} = \mathbf{e}_i^\top \bar{\mathbf{K}} \mathbf{e}_i = \sum_{j=1}^n \lambda_j (\mathbf{u}_j^\top \mathbf{e}_i)^2. \tag{A.5}$$

If $\bar{\mathbf{K}}_{ii} = 0$, then $\mathbf{u}_j^\top \mathbf{e}_i = 0$ for all j such that $\lambda_j > 0$. Therefore

$$\bar{\mathbf{K}} \mathbf{e}_i = \sum_{j=1}^n \lambda_j (\mathbf{u}_j^\top \mathbf{e}_i) \mathbf{u}_j = \mathbf{0}, \tag{A.6}$$

which shows that $\mathbf{e}_i \in \text{Ker } \bar{\mathbf{K}}$, that is, $\mathbf{e}_i \in \text{Im } \mathbf{F}$ by Lemma 1. \square

Lemma 4. Let $\Pi = \mathbf{W}\mathbf{W}^\top = \mathbf{I}_n - \mathbf{F}(\mathbf{F}^\top\mathbf{F})^{-1}\mathbf{F}^\top$ the orthogonal projection matrix on $(\text{Im } \mathbf{F})^\perp$ then, with the hypothesis \mathcal{H}_2 , $(\Pi)_{i,i} \neq 0$ for all $i \in \{1, \dots, n\}$.

Proof. This lemma is a direct application of Lemma 3 by choosing $\mathbf{K} = \mathbf{I}_n$. \square

A.2. Proof of Proposition 1

From preliminary lemmas, we show now the stronger result (stronger than Proposition 1):

Lemma 5. Under the hypotheses $\mathcal{H}_1 - \mathcal{H}_3$, for any $\theta \in (0, +\infty)^d$, there exists $\sigma^2 \in (0, +\infty)$ such that $(\sigma^2, \theta) \in \mathcal{A}_{a,\delta}$.

Proof. Here $\sigma_\epsilon^2 > 0$. Let us assume that $a > 1/2$ (i.e. $q_a > 0$), then for θ fixed in $(0, +\infty)^d$, the limit of $\bar{\mathbf{K}}$ when $\sigma^2 \rightarrow 0$ is well defined and is equal to

$$\lim_{\sigma^2 \rightarrow 0} \bar{\mathbf{K}} = \sigma_\epsilon^{-2} \mathbf{W}\mathbf{W}^\top = \sigma_\epsilon^{-2} \Pi. \tag{A.7}$$

By the hypothesis \mathcal{H}_2 and from Lemma 4 we can write for all $i \in \{1, \dots, n\}$

$$\frac{(\bar{\mathbf{K}}\mathbf{y})_i}{\sqrt{(\bar{\mathbf{K}})_{i,i}}} \xrightarrow{\sigma^2 \rightarrow 0} \frac{1}{\sigma_\epsilon} \frac{(\Pi\mathbf{y})_i}{\sqrt{(\Pi)_{i,i}}}. \tag{A.8}$$

Since $h_\delta^+ \leq h$ for all $\delta > 0$, then

$$\lim_{\sigma^2 \rightarrow 0} \psi_a^{(\delta)}(\sigma^2, \theta) \leq \lim_{\sigma^2 \rightarrow 0} \psi_a(\sigma^2, \theta) = \frac{1}{n} \sum_{i=1}^n h \left(q_a - \frac{1}{\sigma_\epsilon} \frac{(\Pi\mathbf{y})_i}{\sqrt{(\Pi)_{i,i}}} \right) = \frac{k_\epsilon}{n} \tag{A.9}$$

When $\sigma^2 \rightarrow +\infty$, we have

$$\bar{\mathbf{K}} \xrightarrow{\sigma^2 \rightarrow +\infty} \sigma^{-2} \bar{\mathbf{R}}_\theta, \tag{A.10}$$

where

$$\bar{\mathbf{R}}_\theta = \mathbf{W} \left(\mathbf{W}^\top \mathbf{R}_\theta \mathbf{W} \right)^{-1} \mathbf{W}^\top. \tag{A.11}$$

By Lemma 3, we have $(\bar{\mathbf{R}}_\theta)_{i,i} > 0$ for all $i \in \{1, \dots, n\}$ and we obtain that

$$\frac{1}{\sigma} \frac{(\bar{\mathbf{R}}_\theta \mathbf{y})_i}{\sqrt{(\bar{\mathbf{R}}_\theta)_{i,i}}} \xrightarrow{\sigma^2 \rightarrow +\infty} 0. \tag{A.12}$$

With δ small enough satisfying $\delta < q_a$, we obtain

$$\psi_a^{(\delta)}(\sigma^2, \theta) \xrightarrow{\sigma^2 \rightarrow +\infty} \frac{1}{n} \sum_{i=1}^n h_\delta^+(q_a) = 1. \tag{A.13}$$

Since $k_\epsilon < an < n$ by hypothesis \mathcal{H}_3 and since $\psi_a^{(\delta)}$ is continuous, the Intermediate Value Theorem gives the existence of $\sigma_\delta^2 \in (0, +\infty)$ such that

$$\psi_a^{(\delta)}(\sigma_\delta^2, \theta) = a, \tag{A.14}$$

which gives the desired result.

Similarly, if $a < a/2$ then $q_a < 0$ and

$$\lim_{\sigma^2 \rightarrow 0} \psi_a^{(\delta)}(\sigma^2, \theta) \geq \lim_{\sigma^2 \rightarrow 0} \psi_a(\sigma^2, \theta) = \frac{1}{n} \sum_{i=1}^n h \left(q_a - \frac{1}{\sigma_\epsilon} \frac{(\Pi\mathbf{y})_i}{\sqrt{(\Pi)_{i,i}}} \right) = \frac{k_\epsilon}{n} > a. \tag{A.15}$$

When $\delta < q_{1-a}$, one obtains

$$\psi_a^{(\delta)}(\sigma^2, \theta) \xrightarrow{\sigma^2 \rightarrow +\infty} \frac{1}{n} \sum_{i=1}^n h_\delta^-(q_a) = 0. \tag{A.16}$$

By the hypothesis \mathcal{H}_3 , one has the existence of $\sigma_\delta^2 \in (0, +\infty)$ such that

$$\psi_a^{(\delta)}(\sigma_\delta^2, \boldsymbol{\theta}) = a, \tag{A.17}$$

which completes the proof of the lemma. \square

A.3. Proof of Proposition 2

The existence of $\sigma_{\text{opt}}^2(\lambda)$ for all $\lambda \in (0, +\infty)$ results directly from the following Lemma 6:

Lemma 6. For all $\lambda \in (0, +\infty)$, $H_\delta(\lambda)$ is a non-empty and compact subset of \mathbb{R}^+ i.e. H_δ is compact-valued.

Proof. By Lemma 5, $H_\delta(\lambda)$ is non-empty for all $\lambda \in (0, +\infty)$.

$H_\delta(\lambda)$ is closed since the functions h_δ^+, h_δ^- are continuous and the map $(\sigma^2, \boldsymbol{\theta}) \mapsto \bar{\mathbf{K}}$ is also continuous for all $(\sigma^2, \boldsymbol{\theta})$ by the continuity of the kernel function $\mathbf{k}_{\nu, \cdot}^v(\mathbf{x}, \mathbf{x}')$ for any $\nu > 0$ and $\mathbf{x}, \mathbf{x}' \in \mathcal{D}$.

We now prove that $H_\delta(\lambda)$ is bounded. Let us assume that $a \in (1/2, 1)$. If $H_\delta(\lambda)$ is not bounded then there exists a sequence $(\sigma_m^2)_{m \in \mathbb{N}}$ of $H_\delta(\lambda)$ such that $\lim_{m \rightarrow +\infty} \sigma_m^2 = +\infty$ and, by continuity of $\psi_a^{(\delta)}$

$$a = \lim_{m \rightarrow +\infty} \psi_a^{(\delta)}(\sigma_m^2, \lambda \boldsymbol{\theta}_0) = \frac{1}{n} \sum_{i=1}^n h_\delta^+(q_a) = 1, \tag{A.18}$$

which is a contradiction. Therefore, $H_\delta(\lambda)$ is closed and bounded, $H_\delta(\lambda)$ is compact. \square

$\sigma_{\text{opt}}^2(\lambda)$ can be seen the solution of a constrained maximization problem

$$\sigma_{\text{opt}}^2(\lambda) = - \max_{\sigma^2 \in H_\delta(\lambda)} u(\sigma^2, \lambda), \quad \lambda \in (0, +\infty), \tag{A.19}$$

where $u(\sigma^2, \lambda) = -\sigma^2$ is a continuous function. H_δ is non-empty-valued and compact-valued by Lemma 6, upper semi-continuous since $\psi_a^{(\delta)}$ is continuous on $[0, +\infty) \times (0, +\infty)^d$, and continuous if the hypothesis \mathcal{H}_4 is satisfied, the Maximum theorem (Berge (1963), p. 116) provides the continuity of σ_{opt}^2 on $(0, +\infty)$.

A.4. Proof of Proposition 3

Let $\boldsymbol{\theta}_0$ be a solution of one of the problems described in (13) or (18). The continuity of \mathcal{L} on $(0, +\infty)$ follows from the continuity of the trace function $\text{Tr}(\cdot)$, the continuity of the map $(\sigma^2, \boldsymbol{\theta}) \mapsto \bar{\mathbf{K}}$ and the continuity of σ_{opt}^2 by Proposition 2.

Assume that $\lim_{\lambda \rightarrow +\infty} \sigma_{\text{opt}}^2(\lambda) \neq +\infty$, then there exists $M > 0$ such that for all $\lambda > 0$ there exists $\lambda' \geq \lambda$ and $\sigma_{\text{opt}}^2(\lambda') \leq M$. Hence, we can recursively build a sequence $(\lambda_m)_{m \in \mathbb{N}}$ of integers such that $\lambda_{m+1} \geq \lambda_m + 1$ and $\sigma_{\text{opt}}^2(\lambda_m) \leq M$ for all $m \in \mathbb{N}$.

By the Bolzano-Weierstrass theorem, we extract a convergent sub-sequence $(\lambda_{\phi(m)})_{m \in \mathbb{N}}$ where $\phi : \mathbb{N} \rightarrow \mathbb{N}$ such that $\sigma_{\text{opt}}^2(\lambda_{\phi(m)}) \xrightarrow{m \rightarrow +\infty} \sigma_\infty^2 < +\infty$ and

$$\mathbf{K}_{\sigma_{\text{opt}}^2(\lambda_{\phi(m)}), \lambda_{\phi(m)} \boldsymbol{\theta}_0} \xrightarrow{m \rightarrow +\infty} \sigma_\infty^2 \mathbf{J} + \sigma_\epsilon^2 \mathbf{I}_n = \mathbf{K}_\infty. \tag{A.20}$$

When there is a nugget effect $\sigma_\epsilon^2 > 0$, the limit of $\bar{\mathbf{K}}_m := \bar{\mathbf{K}}_{\sigma_{\text{opt}}^2(\lambda_{\phi(m)}), \lambda_{\phi(m)} \boldsymbol{\theta}_0}$ when $m \rightarrow +\infty$ exists because the matrix \mathbf{K}_∞ is nonsingular by the auxiliary fact 1 of De Oliveira et al. (2001)

$$\det \mathbf{K}_\infty = \left(\frac{\sigma_\epsilon^2}{\sigma_\infty^2} \right)^n \left(1 + \frac{\sigma_\epsilon^2}{\sigma_\infty^2} \mathbf{e}^\top \mathbf{I}_n \mathbf{e} \right) = \left(\frac{\sigma_\epsilon^2}{\sigma_\infty^2} \right)^n \left(1 + n \frac{\sigma_\epsilon^2}{\sigma_\infty^2} \right) > 0. \tag{A.21}$$

From hypothesis \mathcal{H}_1 , \mathbf{e} is a column of \mathbf{F} and we can prove that

$$\begin{aligned} \bar{\mathbf{K}}_m \xrightarrow{m \rightarrow +\infty} \bar{\mathbf{K}}_\infty &:= \mathbf{W} \left(\mathbf{W}^\top \left(\sigma_\infty^2 \mathbf{J} + \sigma_\epsilon^2 \mathbf{I}_n \right) \mathbf{W} \right)^{-1} \mathbf{W}^\top \\ &= \sigma_\epsilon^{-2} \mathbf{W} \left(\mathbf{W}^\top \mathbf{W} \right)^{-1} \mathbf{W}^\top = \sigma_\epsilon^{-2} \boldsymbol{\Pi}. \end{aligned} \tag{A.22}$$

By hypothesis \mathcal{H}_2 , the Leave-One-Out formulas (16)-(17) give for all $i \in \{1, \dots, n\}$

$$\frac{(\bar{\mathbf{K}}_m \mathbf{y})_i}{\sqrt{(\bar{\mathbf{K}}_m)_{i,i}}} \xrightarrow{m \rightarrow +\infty} \frac{1}{\sigma_\epsilon} \frac{(\boldsymbol{\Pi} \mathbf{y})_i}{\sqrt{(\boldsymbol{\Pi})_{i,i}}}. \tag{A.23}$$

If $a > 1/2$ for example and by definition of $\sigma_{\text{opt}}^2(\lambda_{\phi(m)})$, one obtains

$$\begin{aligned}
 a &= \frac{1}{n} \sum_{i=1}^n h_{\delta}^+ \left(q_a - \frac{(\bar{\mathbf{K}}_m \mathbf{y})_i}{\sqrt{(\bar{\mathbf{K}}_m)_{i,i}}} \right) \\
 &\xrightarrow{m \rightarrow +\infty} \frac{1}{n} \sum_{i=1}^n h_{\delta}^+ \left(q_a - \frac{(\bar{\mathbf{K}}_{\infty} \mathbf{y})_i}{\sqrt{(\bar{\mathbf{K}}_{\infty})_{i,i}}} \right) \\
 &= \frac{1}{n} \sum_{i=1}^n h_{\delta}^+ \left(q_a - \frac{1}{\sigma_{\epsilon}} \frac{(\boldsymbol{\Pi} \mathbf{y})_i}{\sqrt{(\boldsymbol{\Pi})_{i,i}}} \right) = \frac{k_{\epsilon}}{n} < a,
 \end{aligned}
 \tag{A.24}$$

which is contradictory. Therefore, $\lim_{\lambda \rightarrow +\infty} \sigma_{\text{opt}}^2(\lambda) = +\infty$ and \mathcal{L} is coercive. The case $a < 1/2$ can be addressed in the same way.

Appendix B. The no-nugget case

B.1. Proof of the existence of a solution to problem (33)

In the absence of $\sigma_{\epsilon}^2 = 0$, it follows from the Leave-One-Out formulas that, for all $i \in \{1, \dots, n\}$

$$\frac{(\bar{\mathbf{K}} \mathbf{y})_i}{\sqrt{(\bar{\mathbf{K}})_{i,i}}} = \frac{1}{\sigma} \frac{(\bar{\mathbf{R}}_{\theta} \mathbf{y})_i}{\sqrt{(\bar{\mathbf{R}}_{\theta})_{i,i}}},
 \tag{B.1}$$

which is a monotonic function in σ^2 when θ is fixed in $(0, \infty)^d$.

Let θ be fixed in $(0, +\infty)^d$ and let $a > 1/2$. The proportion $\psi_a^{(\delta)}(\sigma^2, \theta)$ has the limit

$$\lim_{\sigma^2 \rightarrow +\infty} \psi_a^{(\delta)}(\sigma^2, \theta) = \frac{1}{n} \sum_{i=1}^n h_{\delta}^+(q_a) = 1,
 \tag{B.2}$$

and, if $\sigma^2 \rightarrow 0$, it has the limit

$$\lim_{\sigma^2 \rightarrow 0} \psi_a^{(\delta)}(\sigma^2, \theta) = \frac{1}{n} \text{Card} \{i \in \{1, \dots, n\}, (\bar{\mathbf{R}}_{\theta} \mathbf{y})_i \leq 0\} = \frac{k_{\theta}}{n}.
 \tag{B.3}$$

Let θ denote the norm of θ (i.e. $\theta = \|\theta\|$) and consider the set $\mathcal{J} = \{i \in \{1, \dots, n\}, (\boldsymbol{\Pi} \mathbf{y})_i \leq 0\}$. For $i \in \mathcal{J}^c$, one has $(\boldsymbol{\Pi} \mathbf{y})_i > 0$, and, since $\bar{\mathbf{R}}_{\theta}$ converges to $\boldsymbol{\Pi}$ when $\theta \rightarrow 0$

$$\forall i \in \mathcal{J}^c : (\bar{\mathbf{R}}_{\theta} \mathbf{y})_i \xrightarrow{\theta \rightarrow 0} (\boldsymbol{\Pi} \mathbf{y})_i > 0
 \tag{B.4}$$

It results that, there exists $\theta_c > 0$ such that if $\theta \in \mathcal{B}_r(\mathbf{0}, \theta_c)$ (the open ball of radius θ_c centered at $\mathbf{0}$) then $(\bar{\mathbf{R}}_{\theta} \mathbf{y})_i > 0$ for any $i \in \mathcal{J}^c$. Consequently, one gets for any $\theta \in \mathcal{B}_r(\mathbf{0}, \theta_c)$

$$\text{Card} \{i \in \{1, \dots, n\}, (\bar{\mathbf{R}}_{\theta} \mathbf{y})_i > 0\} \geq \text{Card}(\mathcal{J}^c) = n - k_{\epsilon}.
 \tag{B.5}$$

Hence

$$k_{\theta} = \text{Card} \{i \in \{1, \dots, n\}, (\bar{\mathbf{R}}_{\theta} \mathbf{y})_i \leq 0\} \leq k_{\epsilon}.
 \tag{B.6}$$

Therefore, if θ belongs to a neighborhood of $\mathbf{0}$, the condition $k_{\theta} \leq k_{\epsilon}$ is satisfied and, under the hypothesis \mathcal{H}_3 , the set of solutions $\mathcal{A}_{a,\delta}$ is also non-empty.

B.2. Proof of the coercivity

Let assume that, under some conditions on \mathbf{y} , $\lambda \mapsto \sigma_{\text{opt}}^2(\lambda)$ is well-defined for all $\lambda \in (0, +\infty)$. In the absence of nugget effect $\sigma_{\epsilon}^2 = 0$, the limit of $\bar{\mathbf{R}}_{\lambda, \theta_0}$ does not exist when $\lambda \rightarrow +\infty$. Still, we can assume that the correlation matrix $\mathbf{R}_{\lambda, \theta_0}$ satisfies (De Oliveira et al., 2001)

$$\mathbf{R}_{\lambda, \theta_0} = \mathbf{J} + g_{\lambda} (\mathbf{D}_0 + o(1)),
 \tag{B.7}$$

where

- $\lambda \mapsto g_\lambda$ is a continuous function such that $\lim_{\lambda \rightarrow +\infty} g_\lambda = 0$.
- \mathbf{D}_0 and $\mathbf{J} = \mathbf{e}\mathbf{e}^\top$ are fixed symmetric matrices.

\mathbf{D}_0 can be singular or nonsingular depending on the chosen kernel \mathbf{k} . A review of Yagloom’s book (Rosenblatt, 1989) shows that \mathbf{D}_0 is nonsingular only for Power-Exponential ($q < 2$) and Matérn kernels with smoothness parameter $\nu < 1$ like the Exponential kernel ($\nu = 1/2$ in (3)). For the rest of Matérn kernels with smoothness parameter $\nu \geq 1$ \mathbf{D}_0 becomes singular.

Case 1 : \mathbf{D}_0 is nonsingular.

In this case, let $\mathbf{D}_\lambda = g_\lambda \mathbf{D}_0 (1 + o(1))$ such that

$$\mathbf{R}_{\lambda, \theta_0} = \mathbf{J} + \mathbf{D}_\lambda. \tag{B.8}$$

We consider the matrix $\bar{\mathbf{R}}_{\lambda, \theta_0}$ in $\bar{\mathbf{K}} = \sigma^{-2} \bar{\mathbf{R}}_{\lambda, \theta_0}$, we have

$$\bar{\mathbf{R}}_{\lambda, \theta_0} = \mathbf{R}_{\lambda, \theta_0}^{-1} \left[\mathbf{I}_n - \mathbf{F} \left(\mathbf{F}^\top \mathbf{R}_{\lambda, \theta_0}^{-1} \mathbf{F} \right)^{-1} \mathbf{F}^\top \mathbf{R}_{\lambda, \theta_0}^{-1} \right]. \tag{B.9}$$

By using Lemma 4, Appendix B3 in De Oliveira et al. (2001) and under assumption that $\mathbf{e} \in \text{Im } \mathbf{F}$ (hypothesis \mathcal{H}_1), we have

$$\bar{\mathbf{R}}_{\lambda, \theta_0} = \mathbf{D}_\lambda^{-1} \left[\mathbf{I}_n - \mathbf{F} \left(\mathbf{F}^\top \mathbf{D}_\lambda^{-1} \mathbf{F} \right)^{-1} \mathbf{F}^\top \mathbf{D}_\lambda^{-1} \right]. \tag{B.10}$$

Then we get

$$\bar{\mathbf{R}}_{\lambda, \theta_0} = g_\lambda^{-1} \left[\mathbf{D}_0^{-1} \left(\mathbf{I}_n - \mathbf{F} \left(\mathbf{F}^\top \mathbf{D}_0^{-1} \mathbf{F} \right)^{-1} \mathbf{F}^\top \mathbf{D}_0^{-1} \right) + o(1) \right]. \tag{B.11}$$

Finally

$$\bar{\mathbf{R}}_{\lambda, \theta_0} \stackrel{\lambda \rightarrow +\infty}{\sim} g_\lambda^{-1} \mathbf{A}, \tag{B.12}$$

where

$$\mathbf{A} = \mathbf{D}_0^{-1} \left(\mathbf{I}_n - \mathbf{F} \left(\mathbf{F}^\top \mathbf{D}_0^{-1} \mathbf{F} \right)^{-1} \mathbf{F}^\top \mathbf{D}_0^{-1} \right). \tag{B.13}$$

Hypothesis \mathcal{H}_5 : Let \mathbf{A} be the matrix defined in (B.13). We assume that \mathbf{y} does not belong to a family of vectors such that $(\mathbf{A}\mathbf{y})_i = 0$ for all $i \in \{1, \dots, n\}$ and that $\text{Card} \{i \in \{1, \dots, n\}, (\mathbf{A}\mathbf{y})_i \leq 0\} \neq na$.

By applying Lemmas 1 and 2 on \mathbf{D}_0 , we show that $(\mathbf{A}\mathbf{y})_{ii} \neq 0$ and we can write for all i in $\{1, \dots, n\}$

$$\frac{(\bar{\mathbf{R}}_{\lambda, \theta_0} \mathbf{y})_i}{\sqrt{(\bar{\mathbf{R}}_{\lambda, \theta_0})_{ii}}} \stackrel{\lambda \rightarrow +\infty}{\sim} g_\lambda^{-1/2} \frac{(\mathbf{A}\mathbf{y})_i}{\sqrt{(\mathbf{A})_{ii}}}. \tag{B.14}$$

Analogously to the proof of Proposition 3, if we assume that $\lim_{\lambda \rightarrow +\infty} \sigma_{\text{opt}}^2(\lambda) \neq +\infty$ and by taking a sub-sequence $(\sigma_{\text{opt}}^2(\lambda_{\psi(m)}))_{m \in \mathbb{N}}$ converging to σ_∞^2

$$\frac{1}{\sigma_\infty} g_{\lambda_{\psi(m)}}^{-1/2} \frac{(\mathbf{A}\mathbf{y})_i}{\sqrt{(\mathbf{A})_{ii}}} \xrightarrow{m \rightarrow +\infty} \begin{cases} +\infty & \text{if } (\mathbf{A}\mathbf{y})_i > 0 \\ -\infty & \text{otherwise} \end{cases} \tag{B.15}$$

The limit $\psi_a^{(\delta)}(\sigma_{\text{opt}}^2(\lambda_{\psi(m)}), \lambda_{\psi(m)} \theta_0)$ when $m \rightarrow +\infty$ exists and is equal to

$$a = \lim_{m \rightarrow +\infty} \psi_a^{(\delta)}(\sigma_{\text{opt}}^2(\lambda_{\psi(m)}), \lambda_{\psi(m)} \theta_0) = \frac{1}{n} \text{Card} \{i \in \{1, \dots, n\}, (\mathbf{A}\mathbf{y})_i \leq 0\}, \tag{B.16}$$

which is contradictory and completes the proof.

Case 2 : \mathbf{D}_0 is singular.

In this case, one needs to go further in the Taylor expansion of $\bar{\mathbf{R}}_{\lambda, \theta_0}$. We consider the matrix \mathbf{W} in Lemma 3, by Lemma 6 of Ren et al. (2012)

$$\bar{\mathbf{R}}_{\lambda, \theta_0} = \mathbf{W} \left(\mathbf{W}^\top \mathbf{R}_{\lambda, \theta_0} \mathbf{W} \right)^{-1} \mathbf{W}^\top. \tag{B.17}$$

By setting $\Sigma_\lambda = \mathbf{W}^\top \mathbf{R}_{\lambda, \theta_0} \mathbf{W}$, the asymptotic study of $\bar{\mathbf{R}}_{\lambda, \theta_0}$ is equivalent to the asymptotic study of Σ_λ . In case of Matérn kernel with noninteger smoothness $\nu \geq 1$, the matrix Σ_λ can be written as (Muré, 2021)

$$\Sigma_\lambda = g_\lambda \left(\mathbf{W}^\top \mathbf{D}_1 \mathbf{W} + g_\lambda^* \mathbf{W}^\top \mathbf{D}_1^* \mathbf{W} + \mathbf{R}_g(\lambda) \right), \tag{B.18}$$

where

- Either $g_\lambda = c\lambda^{-2k_1}$ with k_1 a nonnegative integer, or $g_\lambda = c\lambda^{-2\nu}$.
- $g_\lambda^* = c^*\lambda^{-2l}$ with $l \in (0, +\infty)$.
- \mathbf{R}_g is a differentiable mapping from $[0, +\infty)$ to \mathcal{M}_n such that $\|\mathbf{R}_g(\lambda)\| = o(\lambda^{-2l})$.
- \mathbf{D}_1 and \mathbf{D}_1^* are both fixed symmetric matrices with elements $\|x_i - x_j\|^{2k}$ where $k \in k_1 \cup \nu$ for \mathbf{D}_1 and $k = l$ for \mathbf{D}_1^* .

The matrix $\mathbf{W}^\top \mathbf{D}_1 \mathbf{W} + g_\lambda^* \mathbf{W}^\top \mathbf{D}_1^* \mathbf{W}$ is nonsingular when $\lambda \rightarrow +\infty$, whether if $\mathbf{W}^\top \mathbf{D}_1 \mathbf{W}$ is nonsingular or if it is singular.

The case where $\mathbf{W}^\top \mathbf{D}_1 \mathbf{W}$ is nonsingular happens for Matérn kernels with smoothness $1 \leq \nu < 2$ (Muré, 2021), whereas the other case occurs for regular and smooth Matérn kernels with $\nu \geq 2$. These kernels are however less robust in uncertainty quantification so we will give only the proof for less smooth kernels with $1 \leq \nu < 2$ in particular the Matérn 3/2 kernel.

In this case, we write Σ_λ in (B.18) as

$$\Sigma_\lambda = g_\lambda \mathbf{W}^\top \mathbf{D}_1 \mathbf{W} \left(\mathbf{I}_n + g_\lambda^* \left(\mathbf{W}^\top \mathbf{D}_1 \mathbf{W} \right)^{-1} \left(\mathbf{W}^\top \mathbf{D}_1^* \mathbf{W} + \mathbf{R}_g(\lambda) \right) \right). \tag{B.19}$$

As \mathbf{W} is full rank matrix, Σ_λ is non-singular and

$$\Sigma_\lambda^{-1} = g_\lambda^{-1} \left(\mathbf{I}_n + g_\lambda^* \left(\mathbf{W}^\top \mathbf{D}_1 \mathbf{W} \right)^{-1} \left(\mathbf{W}^\top \mathbf{D}_1^* \mathbf{W} + \mathbf{R}_g(\lambda) \right) \right)^{-1} \left(\mathbf{W}^\top \mathbf{D}_1 \mathbf{W} \right)^{-1}. \tag{B.20}$$

Let $\mathbf{M}_\lambda = g_\lambda^* \left(\mathbf{W}^\top \mathbf{D}_1 \mathbf{W} \right)^{-1} \left(\mathbf{W}^\top \mathbf{D}_1^* \mathbf{W} + \mathbf{R}_g(\lambda) \right)$, since $\|\mathbf{M}_\lambda\| \xrightarrow{\lambda \rightarrow +\infty} 0$, we can assume that $\|\mathbf{M}_\lambda\| < 1$ when λ is large enough and apply the Taylor series expansion at order 1

$$\begin{aligned} \left[\mathbf{I}_n + g_\lambda^* \left(\mathbf{W}^\top \mathbf{D}_1 \mathbf{W} \right)^{-1} \left(\mathbf{W}^\top \mathbf{D}_1^* \mathbf{W} + \mathbf{R}_g(\lambda) \right) \right]^{-1} &= \mathbf{I}_n - g_\lambda^* \left(\mathbf{W}^\top \mathbf{D}_1 \mathbf{W} \right)^{-1} \\ &\times \left(\mathbf{W}^\top \mathbf{D}_1^* \mathbf{W} + \mathbf{R}_g(\lambda) + o(g_\lambda^*) \right). \end{aligned} \tag{B.21}$$

Then, we plug this quantity into the equation (B.20)

$$\begin{aligned} \Sigma_\lambda^{-1} &= g_\lambda^{-1} \left(\mathbf{I}_n - g_\lambda^* \left(\mathbf{W}^\top \mathbf{D}_1 \mathbf{W} \right)^{-1} \left(\mathbf{W}^\top \mathbf{D}_1^* \mathbf{W} + \mathbf{R}_g(\lambda) \right) \right) \left(\mathbf{W}^\top \mathbf{D}_1 \mathbf{W} \right)^{-1} \\ &= g_\lambda^{-1} \left[\left(\mathbf{W}^\top \mathbf{D}_1 \mathbf{W} \right)^{-1} - g_\lambda^* \left(\mathbf{W}^\top \mathbf{D}_1 \mathbf{W} \right)^{-1} \left(\mathbf{W}^\top \mathbf{D}_1^* \mathbf{W} + \mathbf{R}_g(\lambda) \right) \left(\mathbf{W}^\top \mathbf{D}_1 \mathbf{W} \right)^{-1} \right]. \end{aligned} \tag{B.22}$$

Finally, we can write the matrix $\bar{\mathbf{R}}_{\lambda, \theta_0}$ as

$$\bar{\mathbf{R}}_{\lambda, \theta_0} = g_\lambda^{-1} \mathbf{W} \left[\left(\mathbf{W}^\top \mathbf{D}_1 \mathbf{W} \right)^{-1} - g_\lambda^* \left(\mathbf{W}^\top \mathbf{D}_1 \mathbf{W} \right)^{-1} \left(\mathbf{W}^\top \mathbf{D}_1^* \mathbf{W} + \mathbf{R}_g(\lambda) \right) \left(\mathbf{W}^\top \mathbf{D}_1 \mathbf{W} \right)^{-1} \right] \mathbf{W}^\top. \tag{B.23}$$

We can also simplify the previous expression into

$$\bar{\mathbf{R}}_{\lambda, \theta_0} = g_\lambda^{-1} (\mathbf{A} - \mathbf{B}_\lambda), \tag{B.24}$$

where \mathbf{A} is a fixed matrix and $\mathbf{B}_\lambda \stackrel{\lambda \rightarrow +\infty}{\equiv} o(1)$ such that

$$\mathbf{A} = \mathbf{W} \left(\mathbf{W}^\top \mathbf{D}_1 \mathbf{W} \right)^{-1} \mathbf{W}^\top \tag{B.25}$$

$$\mathbf{B}_\lambda = g_\lambda^* \mathbf{W} \left(\mathbf{W}^\top \mathbf{D}_1 \mathbf{W} \right)^{-1} \left(\mathbf{W}^\top \mathbf{D}_1^* \mathbf{W} + \mathbf{R}_g(\lambda) \right) \left(\mathbf{W}^\top \mathbf{D}_1 \mathbf{W} \right)^{-1} \mathbf{W}^\top. \tag{B.26}$$

Or, equivalently,

$$\bar{\mathbf{R}}_{\lambda, \theta_0} \stackrel{\lambda \rightarrow +\infty}{\sim} g_\lambda^{-1} \mathbf{A}. \tag{B.27}$$

Lemma 7. Let \mathbf{A} be the matrix defined in (B.25), then $\mathbf{A}_{ii} \neq 0$ for all $i \in \{1, \dots, n\}$.

Proof. \mathbf{A} is non-singular because

$$\det \mathbf{A} = \det \mathbf{W} \left(\mathbf{W}^\top \mathbf{D}_1 \mathbf{W} \right)^{-1} \mathbf{W}^\top = \det \left(\mathbf{W}^\top \mathbf{D}_1 \mathbf{W} \right)^{-1} \neq 0. \quad (\text{B.28})$$

\mathbf{A} is then a positive definite matrix

$$\mathbf{A}_{ii} = \mathbf{e}_i^\top \mathbf{A} \mathbf{e}_i > 0. \quad \square \quad (\text{B.29})$$

Hypothesis \mathcal{H}_6 : Let \mathbf{A} be the matrix defined in (B.25). We assume that \mathbf{y} does not belong to a family of vectors such that $(\mathbf{A}\mathbf{y})_i = 0$ for all $i \in \{1, \dots, n\}$ and that $\text{Card} \{i \in \{1, \dots, n\}, (\mathbf{A}\mathbf{y})_i \leq 0\} \neq na$.

With Lemma 6 and Hypothesis \mathcal{H}_6 , the proof of the divergence of $\sigma_{\text{opt}}^2(\lambda)$ when $\lambda \rightarrow +\infty$ is similar to the previous case when \mathbf{D}_0 is nonsingular.

Remark 5. The hypotheses \mathcal{H}_5 and \mathcal{H}_6 are not restrictive, one can verify numerically, that each component of $\mathbf{A}\mathbf{y}$ is not null where \mathbf{A} is one of the matrices defined in (B.13) or (B.25).

References

- Bachoc, F., 2013a. Cross validation and maximum likelihood estimations of hyper-parameters of Gaussian processes with model misspecification. *Comput. Stat. Data Anal.* 66, 55–69.
- Bachoc, F., 2013b. Estimation paramétrique de la fonction de covariance dans le modèle de Krigeage par processus Gaussiens: application à la quantification des incertitudes en simulation numérique. Ph.D. thesis. University Paris 7. <http://www.theses.fr/2013PA077111>.
- Barber, F.R., Candès, E.J., Ramdas, A., Tibshirani, R.J., 2020. The limits of distribution-free conditional predictive inference. *Inf. Inference* 10, 455–482. <https://doi.org/10.1093/imaiai/iaaa017>.
- Berge, C., 1963. *Topological Spaces: Including a Treatment of Multi-Valued Functions, Vector Spaces and Convexity*. Oliver & Boyd.
- Curran, C., Mitchell, T.J., Morris, M.D., Ylvisaker, D., 1991. Bayesian prediction of deterministic functions, with applications to the design and analysis of computer experiments. *J. Am. Stat. Assoc.* 86, 953–963.
- De Oliveira, V., 2007. Objective Bayesian analysis of spatial data with measurement error. *Can. J. Stat. (La Revue Canadienne de Statistique)* 35, 283–301.
- De Oliveira, V., Berger, J.O., Sansó, B., 2001. Objective Bayesian analysis of spatially correlated data. *J. Am. Stat. Assoc.* 96, 1361–1374.
- Dubrule, O., 1983. Cross validation of kriging in a unique neighborhood. *J. Int. Assoc. Math. Geol.* 15, 687–699.
- Efron, B., 1992. Jackknife-after-bootstrap standard errors and influence functions. *J. R. Stat. Soc., Ser. B, Methodol.* 54, 83–127. <http://www.jstor.org/stable/2345949>.
- Efron, B., Tibshirani, R., 1994. *An Introduction to the Bootstrap*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis. <https://books.google.fr/books?id=gLlPlUxRntoC>.
- Filippone, M., Zhong, M., Girolami, M.A., 2013. A comparative evaluation of stochastic-based inference methods for Gaussian process models. *Mach. Learn.* 93, 93–114. <https://doi.org/10.1007/s10994-013-5388-x>.
- Forrester, A.I.J., Söbester, A., Keane, A.J., 2008. *Engineering Design Via Surrogate Modelling: A Practical Guide*. Progress in Astronautics and Aeronautics. American Institute of Aeronautics and Astronautics.
- Gal, Y., Ghahramani, Z., 2016. Dropout as a Bayesian approximation: representing model uncertainty in deep learning. In: *Proceedings of the 33rd International Conference on Machine Learning*. New York. In: *Proceedings of Machine Learning Research*, vol. 48. PMLR, New York, USA, pp. 1050–1059.
- Hastie, T., Tibshirani, R., Friedman, J., 2009. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, second edition. Springer Series in Statistics. Springer, New York.
- Heskes, T., 1997. Practical confidence and prediction intervals. In: *Advances in Neural Information Processing Systems 9*. MIT Press, pp. 176–182.
- Hong, Y., Meeker, W.Q., McCalley, J.D., 2009. Prediction of remaining life of power transformers based on left truncated and right censored lifetime data. *Ann. Appl. Stat.* 3, 857–879. <https://doi.org/10.1214/00-AOAS231>.
- Hwang, J.T.G., Ding, A.A., 1997. Prediction intervals for artificial neural networks. *J. Am. Stat. Assoc.* 92, 748–757.
- Iooss, B., Marrel, A., 2017. An efficient methodology for the analysis and modeling of computer experiments with large number of inputs. In: *UNCECOMP 2017 2nd ECCOMAS Thematic Conference on Uncertainty Quantification in Computational Sciences and Engineering*. Rhodes Island, Greece, pp. 187–197.
- Khosravi, A., Nahavandi, S., Creighton, D., 2010. A prediction interval-based approach to determine optimal structures of neural network metamodels. *Expert Syst. Appl.* 37, 2377–2387. <https://doi.org/10.1016/j.eswa.2009.07.059>.
- Khosravi, A., Nahavandi, S., Creighton, D., Atiya, A.F., 2011. Lower upper bound estimation method for construction of neural network-based prediction intervals. *IEEE Trans. Neural Netw.* 22, 337–346.
- Kleijnen, J.P.C., Sargent, R.G., 2000. A methodology for fitting and validating metamodels in simulation. *Eur. J. Oper. Res.* 120, 14–29.
- Landon, J., Singpurwalla, N., 2008. Choosing a coverage probability for prediction intervals. *Am. Stat.* 62, 120–124. <https://doi.org/10.1198/000313008X304062>.
- Lawless, J.F., Fredette, M., 2005. Frequentist prediction intervals and predictive distributions. *Biometrika* 92, 529–542.
- Lei, J., G'Sell, M., Rinaldo, A., Tibshirani, R.J., Wasserman, L., 2018. Distribution-free predictive inference for regression. *J. Am. Stat. Assoc.* 113, 1094–1111. <https://doi.org/10.1080/01621459.2017.1307116>.
- Li, K., Wang, R., Lei, H., Zhang, T., Liu, Y., Zheng, X., 2018. Interval prediction of solar power using an improved bootstrap method. *Sol. Energy* 159, 97–112. <https://doi.org/10.1016/j.solener.2017.10.051>. <http://www.sciencedirect.com/science/article/pii/S0038092X17309313>.
- MacKay, D.J.C., 1992. A practical Bayesian framework for backpropagation networks. *Neural Comput.* 4, 448–472. <https://doi.org/10.1162/neco.1992.4.3.448>.
- Mardia, K.V., Marshall, R.J., 1984. Maximum likelihood estimation of models for residual covariance in spatial regression. *Biometrika* 71, 135–146.
- Meinshausen, N., 2006. Quantile regression forests. *J. Mach. Learn. Res.* 7, 983–999.
- Moon, H., 2010. *Design and Analysis of Computer Experiments for Screening Input Variables*. Ph.D. thesis. The Ohio State University.
- Morokoff, W.J., Caffisch, R.E., 1995. Quasi-Monte Carlo integration. *J. Comput. Phys.* 122, 218–230.
- Muré, J., 2018. Objective Bayesian analysis of Kriging models with anisotropic correlation kernel. Ph.D. thesis. Sorbonne Paris Cité. <http://www.theses.fr/2018USPCC069>.
- Muré, J., 2021. Propriety of the reference posterior distribution in Gaussian process modeling. *Ann. Stat.* 49, 2356–2377. <https://doi.org/10.1214/20-AOS2040>.
- Neal, R.M., 1993. *Probabilistic Inference Using Markov Chain Monte Carlo Methods*. Technical Report. Dept. of Computer Science, University of Toronto.
- Neal, R.M., 1996. *Bayesian Learning for Neural Networks*. Springer, New York.

- Nix, D.A., Weigend, A.S., 1994. Estimating the mean and variance of the target probability distribution. In: Proceedings of 1994 IEEE International Conference on Neural Networks (ICNN'94), vol. 1, pp. 55–60.
- Oakley, J.E., O'Hagan, A., 2004. Probabilistic sensitivity analysis of complex models: a Bayesian approach. *J. R. Stat. Soc., Ser. B, Stat. Methodol.* 66, 751–769.
- Ohio Department of Natural Resources, 2022. The Ohio oil and gas well locator. <https://ohiodnr.gov/wps/portal/gov/odnr/discover-and-learn/safety-conservation/about-odnr/oil-gas/oil-gas-resources/well-locator>.
- Pang, J., Liu, D., Peng, Y., Peng, X., 2018. Optimize the coverage probability of prediction interval for anomaly detection of sensor-based monitoring series. *Sensors (Basel, Switzerland)* 18. <https://doi.org/10.3390/s18040967>.
- Pearce, T., Brintrup, A., Zaki, M., Neely, A., 2018. High-quality prediction intervals for deep learning: a distribution-free, ensembled approach. In: Proceedings of the 35th International Conference on Machine Learning. Stockholm Sweden. In: Proceedings of Machine Learning Research, vol. 80. PMLR, Stockholm, Sweden, pp. 4075–4084.
- Rasmussen, C.E., Williams, C.K.I., 2005. Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning). The MIT Press.
- Ren, C., Sun, D., He, C.Z., 2012. Objective bayesian analysis for a spatial model with nugget effects. *J. Stat. Plan. Inference* 142, 1933–1946.
- Robert, C.P., Casella, G., 2004. Monte Carlo Statistical Methods. Springer, New York.
- Romano, Y., Patterson, E., Candès, E.J., 2019. Conformalized quantile regression. In: In Advances in Neural Information Processing Systems 32 (NIPS 2019). Curran Associates, Inc., pp. 3538–3548. papers.nips.cc/paper/8613-conformalized-quantile-regression.
- Rosenblatt, M., 1989. Review: A.M. Yaglom, correlation theory of stationary and random functions vol. i; basic results, vol. ii, supplementary notes and references. *Bull., New Ser., Am. Math. Soc.* 20, 207–211. <https://projecteuclid.org:443/euclid.bams/1183555023>.
- Roustant, O., Padonou, E., Deville, Y., Clément, A., Perrin, G., Giorla, J., Wynn, H., 2020. Group kernels for Gaussian process metamodels with categorical inputs. *SIAM/ASA J. Uncertain. Quantificat.* 8, 775–806. <https://doi.org/10.1137/18M1209386>.
- Santner, T.J., Williams, B.J., Notz, W.I., 2003. The Design and Analysis of Computer Experiments. Springer, New York.
- Securities and Exchange Commission, 2010. Modernization of oil and gas reporting, revisions and additions to the definition section in rule 4-10 of regulation s-x. <https://www.sec.gov/rules/final/2008/33-8995.pdf>.
- Shapiro, S.S., Wilk, M.B., 1965. An analysis of variance test for normality (complete samples). *Biometrika* 52, 591–611. <http://www.jstor.org/stable/2333709>.
- Society of Petroleum Engineers, 2022. Petroleum Reserves and Resources Definitions. <https://www.spe.org/en/industry/reserves/>.
- Stein, M.L., 1999. Interpolation of Spatial Data. Springer, New York.
- Tipping, M.E., 2004. Bayesian inference: an introduction to principles and practice in machine learning. In: Advanced Lectures on Machine Learning: ML Summer Schools 2003, Canberra, Australia, February 2 - 14, 2003, Tübingen, Germany, August 4 - 16, 2003, Revised Lectures. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Villani, C., 2009. The Wasserstein distances. In: Optimal Transport: Old and New. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 93–111.
- Wager, S., Hastie, T., Efron, B., 2014. Confidence intervals for random forests: the jackknife and the infinitesimal jackknife. *J. Mach. Learn. Res.* 15 (1), 1625–1651.
- Wallach, D., Goffinet, B., 1989. Mean squared error of prediction as a criterion for evaluating and comparing system models. *Ecol. Model.* 44, 299–306.
- Williams, C., Barber, D., 1998. Bayesian classification with Gaussian processes. *IEEE Trans. Pattern Anal. Mach. Intell.* 20, 1342–1351. <https://doi.org/10.1109/34.735807>.
- Zhang, H., Wang, Y., 2010. Kriging and cross-validation for massive spatial data. *Environmetrics* 21, 290–304. <https://doi.org/10.1002/env.1023>.
- Zhang, H., Zimmerman, J., Nettleton, D., Nordman, D., 2020. Random forest prediction intervals. *Am. Stat.* 74, 392–406.
- Zhao, J., 1997. The lower semicontinuity of optimal solution sets. *J. Math. Anal. Appl.* 207, 240–254.
- Zhou, Y., 1998. Adaptive Importance Sampling for Integration. Ph.D. thesis. Stanford University.