



HAL
open science

Gaussian process interpolation with conformal prediction: methods and comparative analysis

Aurélien Pion, Emmanuel Vazquez

► **To cite this version:**

Aurélien Pion, Emmanuel Vazquez. Gaussian process interpolation with conformal prediction: methods and comparative analysis. LOD 2024, 10th International Conference on Machine Learning, Optimization, and Data Science, Sep 2024, Castiglione della Pescaia Grosseto Italy, Italy. hal-04634707

HAL Id: hal-04634707

<https://centralesupelec.hal.science/hal-04634707v1>

Submitted on 10 Jul 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.

Copyright

Gaussian process interpolation with conformal prediction: methods and comparative analysis

Aurélien Pion^{1,2} and Emmanuel Vazquez¹

Université Paris-Saclay, CNRS, CentraleSupélec,
Laboratoire des Signaux et Systèmes, Gif-sur-Yvette, France
`firstname.lastname@centralesupelec.fr`
Transvalor S.A., Nice, France

July 10, 2024

Abstract: This article advocates the use of conformal prediction (CP) methods for Gaussian process (GP) interpolation to enhance the calibration of prediction intervals. We begin by illustrating that using a GP model with parameters selected by maximum likelihood often results in predictions that are not optimally calibrated. CP methods can adjust the prediction intervals, leading to better uncertainty quantification while maintaining the accuracy of the underlying GP model. We compare different CP variants and introduce a novel variant based on an asymmetric score. Our numerical experiments demonstrate the effectiveness of CP methods in improving calibration without compromising accuracy. This work aims to facilitate the adoption of CP methods in the GP community.

keywords Gaussian processes; Prediction intervals; Calibration; Conformal Prediction

1 Introduction

Building an approximation—whether interpolation or regression—of a computer code represented by a function $f : \mathbb{X} \rightarrow \mathbb{R}$, where $\mathbb{X} \subseteq \mathbb{R}^d$ and $d \in \mathbb{N}^*$, allows for predicting the result of the possibly expensive evaluation of f at a given $x \in \mathbb{X}$. It is often necessary to estimate the uncertainty resulting from the approximation. Using the framework of Gaussian Processes (GP) is a standard Bayesian approach to build such approximations, together with predictive distributions that quantify uncertainty.

This work was supported by Transvalor S.A.

In the following, we focus on interpolation and we consider a prior on f under the form of a GP $Z \sim \text{GP}(m, k)$, where $m : \mathbb{X} \rightarrow \mathbb{R}$ is a mean function and $k : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}$ is a positive definite covariance function. Given data $\mathcal{D}_n = \{(x_1, Z_1), \dots, (x_n, Z_n)\}$, with $Z_i = Z(x_i)$, $i = 1, \dots, n$, the posterior mean function $m_n(\cdot) = \mathbb{E}(Z(\cdot) \mid \mathcal{D}_n)$ provides an approximation of the underlying function and the posterior variance $\sigma_n^2(\cdot) = \text{var}(Z(\cdot) \mid \mathcal{D}_n)$ offers a measure of confidence in the approximation (see, e.g., [Rasmussen and Williams, 2006](#), [Stein, 1999](#)). This allows for the construction of prediction intervals and the identification of regions where uncertainty is high, which is often used for guiding further data collection and model refinement, as in Bayesian optimization (see, e.g., [Feliot et al., 2017](#)).

However, in GP interpolation, the posterior variance does not depend on the Z_i s but only on m , k and the x_i s. To make the posterior variance depend on the data, which is obviously desirable, parameterized processes Z are considered, with a parameter θ that is commonly selected by maximum likelihood (ML) or cross-validation techniques, thus making it dependent on the observed data. Numerical findings by [Petit et al. \(2023\)](#) show that ML generally yields good predictive distributions when the covariance is an anisotropic Matérn function. Additionally, [Karvonen et al. \(2020\)](#) show that ML predictions are generally not too optimistic, in the sense that the posterior variance does not decrease too quickly as the number of observations increases over a fixed domain.

Consider prediction intervals $I_{n,\alpha}(x)$ at level $\alpha \in [0, 1]$, for the approximation of Z at $x \in \mathbb{X}$ from \mathcal{D}_n . Typically, such intervals can be derived from the posterior distributions within the Bayesian framework:

$$I_{n,\alpha}(x) = [m_n(x) + \Phi^{-1}((1 - \alpha)/2) \sigma_n(x), m_n(x) + \Phi^{-1}((1 + \alpha)/2) \sigma_n(x)], \quad (1)$$

where Φ^{-1} stands for the quantile function of the normal inverse distribution. For any $x \in \mathbb{X}$, we have $\mathbb{P}_n(Z(x) \in I_{n,\alpha}) = \alpha$, where \mathbb{P}_n stands for the conditional probability given \mathcal{D}_n . For any α , the empirical coverage

$$\delta_{n,\alpha} = \frac{1}{n^{\text{test}}} \sum_{i=1}^{n^{\text{test}}} \mathbb{1}_{Z_i^{\text{test}} \in I_{n,\alpha}(x_i^{\text{test}})}$$

computed on n^{test} points $(x_i^{\text{test}}, Z_i^{\text{test}})$, with $Z_i^{\text{test}} = Z(x_i^{\text{test}})$, should ideally be close to α , as this is a desirable property expected by users. In this case, we say that the prediction intervals are well calibrated.

However, particularly when the model is misspecified, ML may yield overly optimistic or overly pessimistic predictions (too small or too large empirical coverage). This is illustrated in Figure 1, which shows a scatter plot of the values of the integrated absolute error (IAE, [Marrel and Iooss, 2024](#))

$$J_{\text{IAE}}(\theta) = \int_0^1 |\delta_\alpha - \alpha| d\alpha,$$




Figure 1: IAE versus RMSE, computed by leave-one-out (LOO) on the left-hand side, and on a test set of 1500 points on the right-hand side. Each red point correspond to a different value for the variance and the range/lengthscale parameters of the covariance of Z . The blue square represents the parameter selected by REML, and the green star represents the metrics computed when a conformal prediction method (Jackknife+ for GP) is used to build the prediction intervals. Values, in LOO, above $J_{IAE} > 0.3$ and $RMSE > 5 \cdot 10^4$ are not shown.

that quantifies the calibration of the predictive distributions, versus the root-mean-square error (RMSE) for a given θ . Here, the data \mathcal{D}_n is obtained using one draw of a uniform distribution in \mathbb{X} , with $n = 150$ points, evaluated on the Goldstein-Price test function ($d = 2$; see, e.g., Surjanovic and Bingham, 2013). The GP model Z has an unknown constant mean and an anisotropic Matérn covariance function (Stein, 1999), with regularity parameter $\nu = 2.5$. Each point corresponds to a random value for the parameter θ of the covariance (variance at origin, range/lengthscale parameters). On the left-hand side, the metrics (RMSE, IAE) are computed using \mathcal{D}_n only and a leave-one-out (LOO) strategy is used, while on the right-hand side, they are computed on a test set of 1500 points.

The blue square corresponds to the parameter selected by restricted-ML (REML) (Stein, 1999). Notice that on the left hand-side, the REML point minimizes the RMSE as noted by Petit et al. (2023). On the test set, the REML point does not minimize the RMSE but stays nevertheless close to the Pareto front. No red points reach a value of zero for the IAE, meaning that no parameter of the GP model yields perfectly calibrated predictive distributions. While being a particular case, and a particular function, this situation is by no means exceptional. The code to reproduce this experiment is available to the reader¹.

As shown in Figure 1, it turns out that it is possible to improve the calibration of the prediction intervals using *conformal prediction* (CP). CP is a framework originally proposed by Vovk et al. (2005) to construct prediction intervals for any prediction algorithm, with known statistical properties, under certain assumptions. Recent developments have extended the original

¹See <https://github.com/gpmp-dev/lod2024-conformal>

method to make it more suitable for Gaussian process approximation (Jaber et al., 2024, Lei et al., 2018). Thus, in Figure 1, the green star point represents the metrics computed when the Jackknife+ for GP (J+GP) method of Jaber et al. (2024) is used to construct the prediction intervals. IAE is now closer to zero with the same RMSE, and the calibration has been improved significantly without reducing the accuracy.

The first objective of this article is to advocate for CP-type methods for GP approximation to encourage their adoption by the GP community. Compared to existing works, particularly the work of Jaber et al. (2024), this article specifically examines the noise-free case, that is, GP interpolation, whereas Jaber et al. (2024) only consider the case with observation noise, i.e., GP regression. Additionally, our numerical experiments cover a larger number of functions. We also include in our comparisons the full conformal prediction for GPs proposed by Papadopoulos (2023), as well as a new variant based on an asymmetric score.

The article is organized as follows. In Section 2, we recall the main ideas of CP. Section 3 shows how to apply CP to Gaussian process approximation. In Section 4, we conduct numerical experiments to compare different variants of CP in the case of GP interpolation, using several test functions.

2 Conformal prediction for regression

In this section, we briefly recall the principles of conformal prediction for regression. The reader is referred to Lei et al. (2018), Barber et al. (2021) for more details.

2.1 Non-conformity scores

Consider i.i.d. data $(X_i, Z_i), i = 1, 2, \dots$, from a common distribution on $\mathbb{X} \times \mathbb{R}$, where the X_i s are covariates and the Z_i are the corresponding responses, with mean $\mathbb{E}(Z_i | X_i) = f(X_i)$. We assume that a method is available for constructing regression functions, denoted by $s(\cdot; \mathcal{D})$, that estimates f from a finite dataset $\mathcal{D} = \{(X_i, Z_i), i = 1, 2, \dots\}$.

Conformal prediction aims at producing a prediction interval $I_{n,\alpha}(X)$ for a new response Z at X , given \mathcal{D} . The main idea of CP is to evaluate how well a potential value z for Z fits with the existing dataset when added as a new data point. This is done by introducing a non-conformity score, R , which measures the "distance" or non-conformity of an observation $(x, z) \in \mathbb{X} \times \mathbb{R}$ with respect to the dataset \mathcal{D} . A common choice for the non-conformity score is the residual error

$$R(x, z; \mathcal{D}) = |z - s(x; \mathcal{D})|.$$

2.2 Full conformal prediction

Let $\mathcal{D}_n = \{(X_i, Z_i), i = 1, \dots, n\}$. For a new random covariate X_{n+1} and a potential value $z \in \mathbb{R}$ for the response Z_{n+1} , define the augmented dataset

$$\mathcal{D}_{n+1,z} = \{(X_1, Z_1), \dots, (X_n, Z_n), (X_{n+1}, z)\}.$$

Consider the scores

$$R_{z,i} = R(X_i, Z_i; \mathcal{D}_{n+1,z}), \quad i = 1, \dots, n, \quad \text{and} \quad R_{z,n+1} = R(X_{n+1}, z; \mathcal{D}_{n+1,z}). \quad (2)$$

The original CP method of [Vovk et al. \(2005\)](#), also called full conformal prediction (FCP), consists in building the prediction interval at X_{n+1} defined by

$$I_{n,\alpha}^{\text{FCP}}(X_{n+1}) = \{z \in \mathbb{R}, \gamma(z) \leq \lceil \alpha(n+1) \rceil\}, \quad (3)$$

where

$$\gamma(z) = \sum_{i=1}^{n+1} \mathbb{1}_{R_{i,z} \leq R_{n+1,z}} \quad (4)$$

is the number of non-conformity scores less than or equal to $R_{n+1,z}$.

The interval $I_{n,\alpha}^{\text{FCP}}$ verifies the finite-sample coverage property: $\mathbb{P}(Z_{n+1} \in I_{n,\alpha}^{\text{FCP}}(X_{n+1})) \geq \alpha$ (see [Lei et al., 2018](#); Theorem 2.1), where the probability is taken over the $n+1$ i.i.d. draws $(X_1, Z_1), \dots, (X_{n+1}, Z_{n+1})$. While FCP provides statistically calibrated intervals, it can be computationally expensive since it requires testing several potential values z to get (an approximation of) $I_{n,\alpha}^{\text{FCP}}(X_{n+1})$. For each candidate z , one needs to recalculate the non-conformity scores, which requires fitting a new model. This process results in significant computational overhead, especially for large datasets or complex models. This limitation has led to the development of alternative approaches, such as split conformal prediction, as recalled next.

2.3 Split conformal prediction

An alternative approach is split conformal prediction (SCP), which partitions the data into two sets: a training set and a calibration set. The training set is used to fit the regression model, while the calibration set is used to compute the non-conformity scores: given a partition of the data $\mathcal{D}_n = \mathcal{D}^{\text{train}} \cup \mathcal{D}^{\text{cal}}$, fit the regression model on $\mathcal{D}^{\text{train}}$ and use \mathcal{D}^{cal} to compute the non-conformity scores.

The prediction interval for a new response Z_{n+1} at X_{n+1} is then given by

$$I_{n,\alpha}^{\text{SCP}}(X_{n+1}) = \{z \in \mathbb{R}, R(X_{n+1}, z; \mathcal{D}^{\text{train}}) \leq q_\alpha\}, \quad (5)$$

where q_α is the α quantile of the set $\{R(X_i, Z_i; \mathcal{D}^{\text{train}}), (X_i, Z_i) \in \mathcal{D}^{\text{cal}}\}$ of non-conformity scores from the calibration set.

This approach provides a computationally efficient way to construct prediction intervals while maintaining statistically valid coverage. However, the choice of the split between the training and calibration sets can significantly affect the tightness and coverage of the prediction intervals because each split uses fewer data for training and calibration. To address these limitations, alternative methods such as Jackknife conformal prediction have been developed. These methods aim to use the entire dataset more effectively, reducing the dependency on a single partition.

2.4 Jackknife conformal prediction

Jackknife conformal prediction (JCP) uses all data points both for training and calibration by employing a LOO strategy, where, for each data point (X_i, Z_i) in \mathcal{D}_n , the regression model is fitted on the dataset $\mathcal{D}_{n,-i} = \mathcal{D}_n \setminus \{(X_i, Z_i)\}$.

The prediction interval for a given point $x \in \mathbb{X}$ is then defined by

$$I_{n,\alpha}^{\text{JCP}}(x) = [s(x; \mathcal{D}_n) - q_\alpha, s(x; \mathcal{D}_n) + q_\alpha] \quad (6)$$

where q_α is the empirical α -quantile of the set of scores $R(X_i, Z_i; \mathcal{D}_{n,-i})$, for $i = 1, \dots, n$.

Contrarily to FCP and SCP, JCP does not provide strong theoretical guarantees, but has nevertheless the (weak) following finite-sample in-sample property: $\mathbb{P}(Z_i \in I_{n,\alpha}^{\text{JCP}}(X_i)) \geq \alpha$, for all $i = 1, \dots, n$. If the dataset is large enough, we can expect this property to extend to a new, unobserved point (see Barber et al., 2021; Section 4). Moreover, since JCP uses the entire dataset for both model fitting and calibration, JCP can often produce tighter intervals than those obtained with SCP.

2.5 Jackknife+ conformal prediction

While JCP method has the advantage of using the entire dataset for both training and calibration, it does not provide strong theoretical guarantees as mentioned above. To address this, Barber et al. (2021) introduced the Jackknife+ method (J+), which enhances JCP to provide prediction intervals with finite-sample coverage properties.

Barber et al. observe that regression algorithms can be sensitive to the specific dataset used for training, causing the predictions from models trained on slightly different datasets (e.g., $\mathcal{D}_{n,-i}$) to vary significantly. To mitigate this issue, J+ modifies JCP by considering two sequences,

$$\xi_i^+ = s(X_{n+1}; \mathcal{D}_{n,-i}) + R_i \text{ and } \xi_i^- = s(X_{n+1}; \mathcal{D}_{n,-i}) - R_i, \quad (7)$$

where $R_i = R(X_i, Z_i; \mathcal{D}_{n,-i})$, $i = 1, \dots, n$.

The sequences $\{\xi_i^-\}_{i=1}^n$ and $\{\xi_i^+\}_{i=1}^n$ are then ordered, and the J+ confidence interval at X_{n+1} is defined by

$$I_\alpha^{\text{J}+}(X_{n+1}) = \left[\xi_{(\lfloor (n+1)(1-\alpha) \rfloor)}^-, \xi_{(\lceil (n+1)\alpha \rceil)}^+ \right], \quad (8)$$

provided $\alpha \leq n/(n+1)$.

The J+ method verifies the finite-sample coverage property

$$\mathbb{P}(Z_{n+1} \in I_{n,\alpha}^{\text{J}+}(X_{n+1})) \geq 2\alpha - 1$$

under the i.i.d. assumption for $(X_1, Z_1), \dots, (X_{n+1}, Z_{n+1})$ (see Barber et al., 2021; Theorem 1), but typically achieves coverages close to α in practice.

3 Conformal prediction for Gaussian processes

In this section, we return to the case of GP interpolation, as presented earlier in this article. More precisely, we consider an unknown function f and aim to build an approximation using the framework of GP interpolation. To this end, we assume a GP prior $Z \sim \text{GP}(m, k)$ with known mean and covariance functions, typically determined through a model selection procedure such as maximum likelihood or cross-validation techniques. Given a dataset $\mathcal{D}_n = \{(x_1, Z_1), \dots, (x_n, Z_n)\}$, where $Z_i = Z(x_i)$ for $i = 1, \dots, n$, GP interpolation consists in computing the posterior distribution of $Z(x)$ for all $x \in \mathbb{X}$. We now briefly present the adaptation of CP to GP interpolation. Subsequently, we introduce a novel non-conformity score for the Jackknife+ method.

3.1 Adaptation of Full-Conformal Prediction and Jackknife+

3.1.1 Full-Conformal Prediction for GP (FCP-GP) —

FCP is adapted by Papadopoulos (2023) to GP interpolation, building on the earlier adaptation of CP to kernel ridge regression by Nouretdinov et al. (2001). The main idea involves rewriting the non-conformity scores (2).

Specifically, given a dataset \mathcal{D}_n , and a new point $x \in \mathbb{X}$, define the augmented dataset

$$\mathcal{D}_{n+1} = \mathcal{D}_n \cup \{(x, Z_{n+1})\}, \quad Z_{n+1} = Z(x),$$

assuming that all points x_i and x are distinct. Denote by $m_{n+1,-i}(x)$ and $\sigma_{n+1,-i}^2(x)$ the posterior mean and variance of $Z(x)$ from the LOO dataset $\mathcal{D}_{n+1,-i} = \mathcal{D}_{n+1} \setminus \{(x_i, Z_i)\}$, $i = 1, \dots, n+1$. Then, consider the scores

$$R_i = \frac{|Z_i - m_{n+1,-i}(x_i)|}{\max(\epsilon, \sigma_{n+1,-i}^\beta(x_i))}, \quad i = 1, \dots, n+1, \quad (9)$$

where $\beta > 0$ is a parameter that controls the sensitivity to changes in the variance $\sigma_{n+1,-i}^2(x_i)$, and $\epsilon \geq 0$ is a small constant introduced to ensure

numerical stability when $\sigma_{n+1,-i}^2(x_i)$ is small. Taking $\beta = 1$ is a sensible choice, since under our Bayesian setting, the random variables $(Z_i - m_{n+1,-i}(x_i))/\sigma_{n+1,-i}(x_i)$ are $\mathcal{N}(0, 1)$ and do not depend on the parameters of the GP model.

The FCP-GP method replaces the definition of γ in (4) by

$$\gamma(z) = \mathbb{E} \left[\sum_{i=1}^{n+1} \mathbb{1}_{R_i \leq R_{n+1}} \mid Z_{n+1} = z \right] \quad (10)$$

and computes the prediction intervals as in (3). Using the linearity of $m_{n+1,-i}$ with respect to Z_{n+1} , Papadopoulos (2023) shows that these intervals can be computed efficiently. Note also that Papadopoulos (2023) proposes other types of scores that we do not present here.

3.1.2 Jackknife+ for GP (J+GP) —

Using the scores R_i defined by (9) computed with the training dataset \mathcal{D}_n instead of the augmented dataset \mathcal{D}_{n+1} , Jaber et al. (2024) propose another CP method for GP approximation based on a modification of the J+ method. The sequences (7) are replaced by

$$\xi_i^+(x) = m_{n,-i}(x) + R_i \max \left(\epsilon, \sigma_{n,-i}^\beta(x) \right), \quad (11)$$

$$\xi_i^-(x) = m_{n,-i}(x) - R_i \max \left(\epsilon, \sigma_{n,-i}^\beta(x) \right), \quad (12)$$

and the prediction intervals are computed using (8). Intervals calculated in this way have the same coverage property as that of the J+ method for i.i.d. data $(X_1, Z_1), \dots, (X_{n+1}, Z_{n+1})$, conditional on Z —as when the X_i are uniformly distributed on \mathbb{X} .

Remark. Score normalization has been suggested by Vovk et al. (2005), Chapter 2.3, in their application of full-conformal to ridge regression to adapt to noisy observations. They choose a normalization coefficient such that all the non-conformity scores have relatively the same variance. Later, Lei et al. (2018) discussed the scaling of residual errors $|z - s(z, \mathcal{D})|$ in the case of heteroscedastic noise on the observations, by a measure of the local dispersion.

3.2 Asymmetric scores

Barber et al. (2021) advocate using signed scores when dealing with skewed data. If the data exhibit asymmetry, such as local excursions, the prediction intervals should ideally reflect this asymmetry. To this end, we introduce a modification of the scores (9) by removing the absolute value:

$$R_i = \frac{Z_i - m_{n,-i}(x_i)}{\max(\epsilon, \sigma_{n,-i}(x_i))}. \quad (13)$$

Then, define the sequence

$$\xi_i(x) = m_{n,-i}(x) + R_i \max(\epsilon, \sigma_{n,-i}(x)).$$

The prediction intervals for a given $x \in \mathbb{X}$ is obtained as

$$I_{n,\alpha}^{\text{asymJ+GP}}(x) = [\xi_{\lfloor(1-\alpha)/2(n+1)\rfloor}(x), \xi_{\lfloor(1+\alpha)/2(n+1)\rfloor}(x)]. \quad (14)$$

We refer to this method as `asymJ+GP`. Moreover, the finite-sample coverage property for i.i.d. data, $\mathbb{P}(Z_{n+1} \in I_{\alpha}^{\text{J+}}(X_{n+1})) \geq 2\alpha - 1$, remains valid, as demonstrated in Appendix A of Barber et al. (2021).

4 Numerical Comparison

In this section, we conduct a numerical comparison between the different CP methods. We use `GPmp`, the Python GP micro package (Vazquez, 2024), for the implementation.

We consider GP models Z with a constant mean and a Matérn covariance function with regularity $\nu = p + 1/2$, where $p \in \mathbb{N}^*$. The Matérn correlation structure κ_{ν} is defined in Chapter 2.7 of Stein (1999), and the corresponding anisotropic covariance function can be written as

$$k_{\sigma,\nu,\rho}(x, y) = \sigma^2 \kappa_{\nu} \left(\sqrt{\sum_{i=1}^d \frac{(x_{[i]} - y_{[i]})^2}{\rho_i^2}} \right) \quad x, y \in \mathbb{R}^d. \quad (15)$$

The parameter ν controls the regularity of the covariance, σ^2 is a variance parameter, and the parameters ρ_i are lengthscale parameters controlling the scale of variations along each dimension. The parameters σ^2 and the ρ_i s are selected by REML.

The goal is to understand the behavior of the prediction intervals computed either using (1) or a CP method when the GP model has different and not necessarily “optimal” values of regularity ν .

For a given value of ν , we sample $n^{\text{train}} = 20 \times d$ points in \mathbb{X} , $(x_1, \dots, x_{n^{\text{train}}})$, using the uniform distribution on \mathbb{X} , and compute the outputs $f(x_1), \dots, f(x_{n^{\text{train}}})$. We apply the same strategy to compute the test set with $n^{\text{test}} = 1100$ points. This procedure is repeated 40 times to compute on each repetition the empirical coverage δ_{α} , with $\alpha = 90\%$, the spatial average of the width of the intervals, and the IAE on the test set. All non-conformity scores are computed with $\beta = 1$.

The test functions used for experiments are the Goldstein-Price function ($d = 2$), the Hartmann4 function ($d = 4$), the Hartmann6 function ($d = 6$), the Park function ($d = 4$), the Branin function ($d = 2$), and a Becker function in dimension $d = 2$. Information about the Goldstein-Price, and Hartmann functions can be found in Surjanovic and Bingham (2013) and information

about the Branin function can be found in [Dixon and Szego \(1978\)](#). The Park function is defined in, e.g., [Cox et al. \(2001\)](#). The Becker functions are from [Becker \(2020\)](#).

Figure 2 displays the boxplots of the empirical coverage and the average width of the intervals (over the 40 repetitions) when the regularity parameter varies, in the case of the Goldstein-Price function. All CP methods give better coverage than the GP model when the parameters are selected by REML, which is generally overconfident when p increases. The FCP-GP method is more optimistic (smaller prediction intervals, smaller coverage) than J+GP or asymJ+GP. The improvement in coverage by the CP methods is achieved by increasing the size of the prediction interval, as shown by the right-hand side of Figure 2.

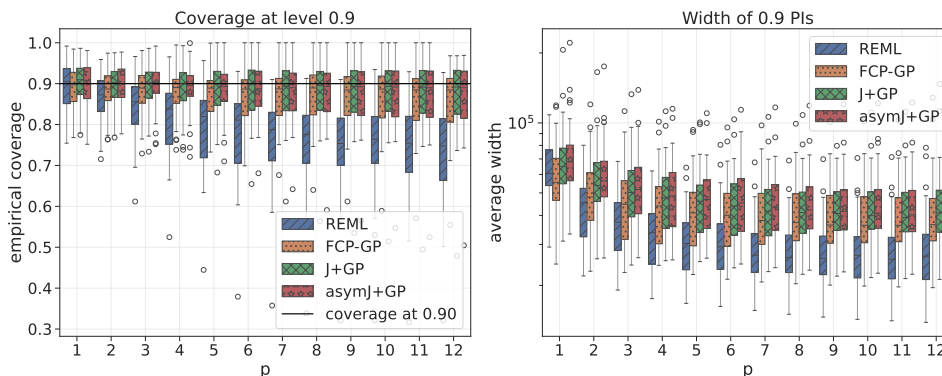


Figure 2: Coverage and average width of the prediction intervals at level 0.9 for the Goldstein-Price function with 40 training points. The GP model parameters σ and ρ are selected by REML. The intervals are computed using the posterior variance, the FCP-GP method, the J+GP method, and the asymJ+GP method.

Table 1 summarizes the performance of the methods for the other test functions. The average IAE and the average width of the prediction interval at the 90% level, computed on the test set, are reported for several values of model regularity $\nu = p + 1/2$. The IAE reflects the coverage at multiple levels, and all CP methods improve the IAE compared to REML. All CP methods produce relatively similar results. The asymJ+GP method gives very similar results to J+GP and has a better IAE for the Beck function. However, it should be noted that the asymJ+GP method produces wider prediction intervals on average than J+GP.

5 Discussion

CP methods can enhance prediction intervals, as shown in Figure 1. All CP methods provide relatively similar coverage, though J+GP tends to perform slightly better on average. For data with excursions, such as the Goldstein-Price function, using asymJ+GP should be preferable. In other scenarios,

Function	P	REML		FCP-GP		J+GP		asymJ+GP	
		90%W	IAE	90%W	IAE	90%W	IAE	90%W	IAE
Beck	1	0.045	0.21	0.035	0.06	0.043	0.06	0.047	0.06
	5	0.0059	0.1	0.009	0.09	0.01	0.08	0.011	0.07
	9	0.0064	0.24	0.012	0.12	0.013	0.11	0.014	0.11
Branin	1	$1.4 \cdot 10^1$	0.24	7.6	0.06	8.7	0.06	9.1	0.06
	5	0.52	0.12	0.5	0.1	0.51	0.1	0.54	0.1
	9	0.75	0.08	0.9	0.09	0.88	0.09	0.91	0.09
Goldstein Price	1	$6.4 \cdot 10^4$	0.19	$6 \cdot 10^4$	0.06	$7 \cdot 10^4$	0.06	$7.4 \cdot 10^4$	0.06
	5	$3.3 \cdot 10^4$	0.09	$4.4 \cdot 10^4$	0.08	$4.8 \cdot 10^4$	0.08	$4.9 \cdot 10^4$	0.08
	9	$2.9 \cdot 10^4$	0.11	$4.2 \cdot 10^4$	0.09	$4.5 \cdot 10^4$	0.08	$4.6 \cdot 10^4$	0.08
Hartmann 4	1	0.98	0.08	0.86	0.04	0.9	0.04	0.92	0.04
	5	0.84	0.05	0.82	0.04	0.83	0.04	0.85	0.04
	9	0.83	0.05	0.82	0.05	0.83	0.04	0.85	0.04
Hartmann 6	1	0.68	0.14	0.54	0.03	0.56	0.03	0.6	0.03
	5	0.64	0.11	0.54	0.04	0.56	0.04	0.59	0.04
	9	0.63	0.11	0.53	0.04	0.56	0.04	0.59	0.04
Park Function	1	0.072	0.25	0.032	0.04	0.034	0.04	0.034	0.05
	5	0.0015	0.08	0.0015	0.07	0.0015	0.07	0.0015	0.07
	9	0.0012	0.08	0.0013	0.08	0.0014	0.08	0.0014	0.08

Table 1: Average width of 90% interval and average IAE for multiple test functions. 90%W stands for the average width of 90% prediction interval.

J+GP is generally a good choice.

As a final remark, Figure 1 suggests that the RMSE of the GP model obtained by REML is near optimal, as already noted by [Petit et al. \(2023\)](#). Our findings indicate that it might be beneficial to decouple the objectives of achieving high prediction accuracy (low RMSE) and obtaining reliable prediction intervals (low IAE), but we did not explore the fully Bayesian approach (see, e.g., [Benassi et al., 2011](#)) in this study.

References

- R. F. Barber, E. J. Candès, A. Ramdas, and R. J. Tibshirani. Predictive inference with the jackknife+. *The Annals of Statistics*, 49(1), 2021.
- W. Becker. Metafunctions for benchmarking in sensitivity analysis. *Reliability Engineering & System Safety*, 204:107189, 2020. ISSN 0951-8320.
- R. Benassi, J. Bect, and E. Vazquez. Robust gaussian process-based global optimization using a fully bayesian expected improvement criterion. In Carlos A. Coello Coello, editor, *Learning and Intelligent Optimization*, pages 176–190, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg. ISBN 978-3-642-25566-3.

- D. Cox, J. S. Park, and C. E. Singer. A statistical method for tuning a computer code to a data base. *Computational Statistics and Data Analysis*, 37(1):77–92, 2001. ISSN 0167-9473.
- L. C W. Dixon and G. P. Szego. The global optimization problem: an introduction. *Towards global optimization*, (2):1–15, 1978.
- P. Feliot, J. Bect, and E. Vazquez. A Bayesian approach to constrained single-and multi-objective optimization. *Journal of Global Optimization*, 67(1):97–133, 2017.
- E. Jaber, V. Blot, N. Brunel, V. Chabridon, E. Remy, B. Iooss, D. Lucor, M. Mougeot, and A. Leite. Conformal approach to Gaussian process surrogate evaluation with coverage guarantees, 2024. hal-04389163 (preprint submitted on 11 January 2024).
- T. Karvonen, G. Wynne, F. Tronarp, C. Oates, and S. Särkkä. Maximum likelihood estimation and uncertainty quantification for Gaussian process approximation of deterministic functions. *SIAM/ASA Journal on Uncertainty Quantification*, 8(3):926–958, 2020. ISSN 2166-2525.
- J. Lei, M. G’Sell, A. Rinaldo, R. J. Tibshirani, and L. Wasserman. Distribution-free predictive inference for regression. *Journal of the American Statistical Association*, 113(523):1094–1111, 2018.
- A. Marrel and B. Iooss. Probabilistic surrogate modeling by Gaussian process: A new estimation algorithm for more robust prediction. *Reliability Engineering & System Safety*, 247:110120, 2024. ISSN 0951–8320.
- I. Nouretdinov, T. Melluish, and V. Vovk. Ridge regression confidence machine. In *Proceedings of the 18th International Conference on Machine Learning ICML*, pages 385–392. CA: Morgan Kaufmann, 2001.
- H. Papadopoulos. Guaranteed coverage prediction intervals with Gaussian process regression. *IEEE Transactions on Pattern Analysis and Machine Learning* (submitted for publication), 2023.
- S. J. Petit, J. Bect, P. Feliot, and E. Vazquez. Parameter selection in Gaussian process interpolation: An empirical study of selection criteria. *SIAM/ASA Journal on Uncertainty Quantification*, 11(4):1308–1328, 2023. ISSN 2166-2525.
- C. E. Rasmussen and C. K. I. Williams. *Gaussian processes for machine learning*. Adaptive computation and machine learning. MIT Press, 2006. ISBN 978-0-262-18253-9.
- M. L. Stein. *Interpolation of Spatial Data*. Springer Series in Statistics. Springer New York, 1999. ISBN 978-1-4612-7166-6 978-1-4612-1494-6.

- S. Surjanovic and D. Bingham. Virtual library of simulation experiments: Test functions and datasets, 2013. <http://www.sfu.ca/~ssurjano>.
- E. Vazquez. GPmp: the Gaussian process micro package, 2024. <https://github.com/gpmp-dev/gpmp>.
- V. Vovk, A. Gammerman, and G. Shafer. *Algorithmic Learning in a Random World*. Springer, 2005.