

Practical and Rigorous Uncertainty Bounds for Gaussian Process Regression

Christian Fiedler,^{1, 2, 3} Carsten W. Scherer,² Sebastian Trimpe^{1, 3}

¹Intelligent Control Systems Group, Max Planck Institute for Intelligent Systems

²Mathematical Systems Theory, University of Stuttgart

³Institute for Data Science in Mechanical Engineering, RWTH Aachen University

christian.fiedler@dsme.rwth-aachen.de, carsten.scherer@imng.uni-stuttgart.de, trimpe@dsme.rwth-aachen.de

Abstract

Gaussian Process Regression is a popular nonparametric regression method based on Bayesian principles that provides uncertainty estimates for its predictions. However, these estimates are of a Bayesian nature, whereas for some important applications, like learning-based control with safety guarantees, frequentist uncertainty bounds are required. Although such rigorous bounds are available for Gaussian Processes, they are too conservative to be useful in applications. This often leads practitioners to replacing these bounds by heuristics, thus breaking all theoretical guarantees. To address this problem, we introduce new uncertainty bounds that are rigorous, yet practically useful at the same time. In particular, the bounds can be explicitly evaluated and are much less conservative than state of the art results. Furthermore, we show that certain model misspecifications lead to only graceful degradation. We demonstrate these advantages and the usefulness of our results for learning-based control with numerical examples.

Introduction

Gaussian Processes Regression (GPR) is an established and successful nonparametric regression method based on Bayesian principles (Rasmussen and Williams 2006) which has recently become popular in learning-based control (Liu et al. 2018; Kocijan 2016). In this context, safety and performance guarantees are important aspects (Åström and Murray 2010; Skogestad and Postlethwaite 2007). In fact, the lack of rigorous guarantees has been identified as one of the major obstacles preventing the usage of learning-based control methodologies in safety-critical areas like autonomous driving, human-robot interaction or medical devices, see e.g. (Berkenkamp 2019). One approach to tackle this challenge is to use the posterior variance of GPR to derive frequentist uncertainty bounds and combine these with robust control methods that can deal with the remaining uncertainty. This strategy has been successfully applied in a number of works that also provide control-theoretical guarantees, cf. Section Bayesian and Frequentist Bounds.

These approaches rely on rigorous frequentist uncertainty bounds for GPR. Although such results are available (Srinivas et al. 2010, Theorem 6), (Chowdhury and Gopalan 2017,

Theorem 2), they turn out to be very conservative and difficult to evaluate numerically in practice and are, hence, replaced by heuristics. That is, instead of bounds obtained from theory, much smaller ones are assumed, sometimes without any practical justification, cf. Related Work for more discussion. Unfortunately, using heuristic approximations leads to a breakdown of the theoretical guarantees of these control approaches, as already observed for example in (Lederer, Umlauf, and Hirche 2019). Furthermore, when deriving theoretical guarantees based on frequentist results like (Srinivas et al. 2010, Theorem 6) or (Chowdhury and Gopalan 2017, Theorem 2), model misspecifications (like wrong hyperparameters of the underlying GPR model or approximations such as the usage of sparse GPs) are ignored. Since model misspecifications are to be expected in any realistic setting, the validity of such theoretical guarantees based on idealized assumptions is unclear.

In summary, rigorous and practical frequentist uncertainty bounds for GPR are currently not available. By practical we mean that concrete, not excessively conservative numerical bounds can be computed based on reasonable and established assumptions and that these are robust against model misspecifications at least to some extent. We note that such bounds are of independent interest, for example, for Bayesian Optimization (Shahriari et al. 2015). In this work, we improve previous frequentist uncertainty bounds for GPR leading to practical, yet theoretically sound results. In particular, our bounds can be directly used in algorithms and are sharp enough to avoid the use of unjustified heuristics. Furthermore, we provide robustness results that can handle moderate model misspecifications. Numerical experiments support our theoretical findings and illustrate the practical use of the bounds.

Background

Gaussian Process Regression and Reproducing Kernel Hilbert Spaces

We briefly recall the basics of GPR, for more details we refer to (Rasmussen and Williams 2006). A Gaussian Process (GP) over an (input or index) set D is a collection of random variables, such that any finite subset is normally distributed. A GP f is uniquely determined by its mean function $m(x) = \mathbb{E}[f(x)]$ and covariance function

$k(x, x') = \mathbb{E}[(f(x) - m(x))(f(x') - m(x'))]$ and we write $f \sim \mathcal{G}_D(m, k)$. Without loss of generality we focus on the case $m \equiv 0$. Common covariance functions include the Squared Exponential (SE) and Matern kernel. Consider a Gaussian Process prior $f \sim \mathcal{G}_D(0, k)$ and noisy data $(x_i, y_i)_{i=1, \dots, N}$, where $x_i \in D$ and $y_i = f(x_i) + \epsilon_i$ with i.i.d. $\mathcal{N}(0, \sigma^2)$ noise, then the posterior is also a GP, with the *posterior mean* μ_N , *posterior covariance* k_N and *posterior variance* given by

$$\begin{aligned}\mu_N(x) &= \mu(x) + \mathbf{k}_N(x)^T (\mathbf{K}_N + \sigma^2 \mathcal{I}_N)^{-1} \mathbf{y}_N \\ k_N(x, x') &= k(x, x') - \mathbf{k}_N(x)^T (\mathbf{K}_N + \sigma^2 \mathcal{I}_N)^{-1} \mathbf{k}_N(x') \\ \sigma_N^2(x) &= k_N(x, x),\end{aligned}$$

where we defined the kernel matrix $\mathbf{K}_N = (k(x_j, x_i))_{i, j=1, \dots, N}$ and the column vectors $\mathbf{k}_N(x) = (k(x_i, x))_{i=1, \dots, N}$ and $\mathbf{y}_N = (y_i)_{i=1, \dots, N}$.

Later on we follow (Srinivas et al. 2010; Chowdhury and Gopalan 2017) and assume that the ground truth is a function from a Reproducing Kernel Hilbert Space (RKHS). For an introduction to the RKHS framework we refer to (Steinwart and Christmann 2008, Chapter 4) or (Berlinet and Thomas-Agnan 2011) as well as (Kanagawa et al. 2018) for connections between Gaussian Processes and the RKHS framework.

Bayesian and Frequentist Bounds

GPR is based on Bayesian principles: A prior distribution is chosen (here a GP with given mean and covariance function) and then updated with the available data using Bayes rule, assuming a certain likelihood or noise model (here independent Gaussian noise). The updated distribution is called the posterior distribution and can be interpreted as a tradeoff between prior belief (encoded in the prior distribution together with the likelihood model) and evidence (the data) (Murphy 2012, Chapter 5). In contrast to the Bayesian approach, in frequentist statistics the existence of a ground truth is assumed and noisy data about this ground truth is acquired (Murphy 2012, Chapter 6).

For many applications which require safety guarantees it is important to get reliable frequentist uncertainty bounds. A concrete setting and major motivation for this work is learning-based robust control. Here, the ground truth is an only partially known dynamical system and the goal is to solve a certain control task, like stabilization or tracking, by finding a suitable controller. A machine learning method like GPR is used to learn more about the unknown dynamical system from data. Thereafter a set of possible models is derived from the learning method that contains the ground truth with a given (high) probability. For this reason such a set is often called an uncertainty set. We then use a robust method on this set, i.e., a method that leads to a controller that works for every element of this set. Since the ground truth is contained in this set with a given (high) probability, the task is solved with this (high) probability. Examples of such an approach are (Koller et al. 2018) (using robust model predictive control with state constraints), (Umlauf et al. 2017) (using feedback linearization to achieve ultimate bounded-

ness) and (Helwa, Heins, and Schoellig 2019) (considering tracking of Lagrangian systems).

The control performance typically degrades with larger uncertainty sets, and it might even be impossible to achieve the control goal if the uncertainty sets are too large (Skogstad and Postlethwaite 2007). Therefore it is desirable to extract uncertainty sets that are as small as possible, but still include the ground truth with a given high probability. In particular, when using GPR, this necessitates frequentist uncertainty bounds that are not too conservative, so that the uncertainty sets are not too large. Furthermore, we also have to be able to evaluate the uncertainty bounds numerically, since robust control methods typically need an explicit uncertainty set.

For learning-based control applications using GPR together with robust control methods, bounds of the following form have been identified as the most useful ones: Let X be an arbitrary input set and assume that $f : X \rightarrow \mathbb{R}$ is the unknown ground truth. Let μ_N be the posterior mean using a dataset of size N , generated from the ground truth, and let $\delta \in (0, 1)$ be given. We need a function $\nu_N(x)$ that can be explicitly evaluated, such that with probability at least $1 - \delta$ with respect to the data generating process, we have

$$|f(x) - \mu_N(x)| \leq \nu_N(x) \quad \forall x \in X. \quad (1)$$

We emphasize that the probability statement is with respect to the noise generating process, and that the underlying ground truth is just some deterministic function. For a more thorough discussion we refer to (Berkenkamp 2019, Section 2.5.3).

Related Work

Frequentist uncertainty bounds for GPR as considered in this work were originally developed in the literature on bandits. To the best of our knowledge, the first result in this direction was (Srinivas et al. 2010, Theorem 6), which is of the form (1) with $\nu_N(x) = \beta_N \sigma_N(x)$. Here β_N is a constant that depends on the *maximum information gain*, an information theoretic quantity. For common settings, there exist upper bounds on the latter quantity, though these increase with sample size. The result from (Srinivas et al. 2010) has been significantly improved in (Chowdhury and Gopalan 2017, Theorem 2), though the latter still uses an upper bound depending on the maximum information gain. The importance of (Srinivas et al. 2010, Theorem 6) and (Chowdhury and Gopalan 2017, Theorem 2) for other applications outside the bandit setting has been recognized early on, in particular in the control community, for example in (Berkenkamp, Schoellig, and Krause 2016; Berkenkamp et al. 2016; Koller et al. 2018; Berkenkamp 2019; Umlauf et al. 2017; Helwa, Heins, and Schoellig 2019).

Unfortunately, both (Srinivas et al. 2010, Theorem 6) and (Chowdhury and Gopalan 2017, Theorem 2) tend to be very conservative, especially for control applications (Berkenkamp, Schoellig, and Krause 2016; Berkenkamp 2019). Furthermore, both results rely on an upper bound of the maximum information gain, which can be difficult to evaluate (Srinivas et al. 2010), though asymptotic bounds are available for standard kernels like the linear, SE or Matern

kernel. However, for most control applications, these asymptotic bounds are not useful since one requires a concrete numerical bound in the nonasymptotic setting.

For these reasons, previous work utilizing (Srinivas et al. 2010, Theorem 6) or (Chowdhury and Gopalan 2017, Theorem 2) used heuristics. In the control community, it is common to choose a constant value $\beta_N \equiv \beta$ in an ad-hoc manner, see for example (Berkenkamp et al. 2017, 2016; Koller et al. 2018; Berkenkamp, Schoellig, and Krause 2016; Helwa, Heins, and Schoellig 2019). This choice does not reflect the asymptotic behaviour of the scaling since the maximum information gain grows with the number of samples (Srinivas et al. 2010)¹. The problem with such heuristics is that they might work in practice, but *all* theoretical guarantees that are based on results like (Chowdhury and Gopalan 2017, Theorem 2) are lost, in particular, safety guarantees like constraint satisfaction or certain stability notions. This is especially problematic since one of the major incentives to use GPR together with such results is to *provably* ensure properties like stability of a closed loop system.

Note that in concrete applications one has to make some assumptions on the ground truth at some point. However, it is not clear at all how a constant scaling β as used in these heuristics can be derived from real-world properties in a principled manner. In contrast to such heuristics, in the original bounds (Srinivas et al. 2010, Theorem 6), (Chowdhury and Gopalan 2017, Theorem 2) every ingredient (i.e., desired probability, size of noise, bound on RKHS norm) has a clear interpretation.

It seems to be well-known in the multiarm-bandit literature that, in the present situation, it is possible to use more empirical bounds than (Chowdhury and Gopalan 2017, Theorem 2). Such approaches are conceptually similar to the results that we will present below, cf. (Abbasi-Yadkori 2013) and the recent work (Calandriello et al. 2019). However, to the best of our knowledge, results like Theorem 1 are rarely used in applications. In particular, it seems that no attempts have been made in the control community to use such a-posteriori bounds in the GPR context.

Model misspecification in the context of GPR has been discussed before in some works. In (Beckers, Umlauf, and Hirche 2018) a numerical method for bounding the mean-squared-error of a misspecified GP is considered, but it relies on a probabilistic setting. The recent article (Wang, Tuo, and Jeff Wu 2019) provides uniform error bounds and deals with misspecified kernels, but again uses a probabilistic setting and focuses on the noise-free case.

A work with goals similar to ours is (Lederer, Umlauf, and Hirche 2019). The authors recognize and explicitly discuss some of the problems of (Chowdhury and Gopalan 2017, Theorem 2) in the context of control. However, (Lederer, Umlauf, and Hirche 2019) uses a probabilistic setting, while our work is concerned with a fixed, but unknown underlying target function, and hence our results are of a worst-case nature. As we argued above, this is the setting required for robust approaches. Finally, the very recent work

¹Sometimes other heuristics are used that ensure that β_N grows with N , cf. e.g. (Kandasamy, Schneider, and Póczos 2015)

(Maddalena, Scharnhorst, and Jones 2020) requires bounded noise and does not deal with model misspecification.

Practical and Rigorous Frequentist Uncertainty Bounds

We now present our main technical contributions. The key observation is that for many applications relying on frequentist uncertainty bounds, in particular, learning-based control methods, only *a-posteriori* bounds are needed. This means that the frequentist uncertainty set can be derived *after* the learning process and hence the concrete realization of the dataset can be used. We take advantage of this fact and modify existing uncertainty results so that they explicitly depend on the dataset used for learning. In general no *a-priori* guarantees can be derived from the results we present here, but this does not play a role in the present setting.

The following result, which is a modified version of (Chowdhury and Gopalan 2017, Theorem 2), is fundamental for the rest of the paper.

Theorem 1. *Let $D \neq \emptyset$ be a set and $k : D \times D \rightarrow \mathbb{R}$ a positive definite kernel with corresponding RKHS $(H_k, \|\cdot\|_k)$ and let $f \in H_k$ with $\|f\|_k \leq B$ for some $B \geq 0$. Let $\mathbb{F} = (\mathcal{F}_n)_{n \in \mathbb{N}}$ be a filtration and $(x_n)_{n \in \mathbb{N}}$ a D -valued discrete-time stochastic process that is predictable w.r.t. \mathbb{F} and let $(\epsilon_n)_n$ be an \mathbb{R} -valued stochastic process adapted to \mathbb{F} , such that ϵ_n conditioned on \mathcal{F}_{n-1} is R -subgaussian with $R \geq 0$. Furthermore, define $y_n = f(x_n) + \epsilon_n$ for all $n \geq 1$.*

Consider a Gaussian process $g \sim \mathcal{G}_D(0, k)$ and denote its posterior mean function by μ_N , its posterior covariance function by k_N and its posterior variance by $\sigma_N^2(x) := k_N(x, x)$, w.r.t. to data $(x_1, y_1), \dots, (x_N, y_N)$, assuming in the likelihood independent Gaussian noise with mean zero and variance $\lambda > 0$. Then for any $\delta \in (0, 1)$ with $\bar{\lambda} = \max\{1, \lambda\}$ and

$$\begin{aligned} \beta_n &= \beta_n(\delta, B, R, \lambda) \\ &= B + R \sqrt{\log(\det(\mathbf{K}_n + \bar{\lambda} \mathcal{I}_n)) - 2 \log(\delta)} \end{aligned} \quad (2)$$

one has

$$\mathbb{P}[|\mu_N(x) - f(x)| \leq \beta_N \sigma_N(x) \forall N \in \mathbb{N}, x \in D] \geq 1 - \delta. \quad (3)$$

Proof. (Idea) Essentially identical to the proof of (Chowdhury and Gopalan 2017, Theorem 2), however, we do not upper-bound (2) by the maximum information gain. Details are given in the supplementary material. \square

The key insight is that by appropriate modifications of the proof of (Chowdhury and Gopalan 2017, Theorem 2) we obtain a frequentist uncertainty bound for GPR that fulfills all desiderata from Section Bayesian and Frequentist Bounds. In the numerical examples below we show that the bound is often tight enough for practical purposes.

Independent inputs and noise (e.g. deterministic inputs and i.i.d. noise) is a common situation that is simpler than the setting of Theorem 1. In this case the following a-posteriori bound can be derived that does not depend anymore on $\log \det$.

Proposition 2. Let $D \neq \emptyset$ be a set and $k : D \times D \rightarrow \mathbb{R}$ a positive definite kernel with corresponding RKHS $(H_k, \|\cdot\|_k)$ and let $f \in H_k$ with $\|f\|_k \leq B$ for some $B \geq 0$. Let $x_1, \dots, x_N \in D$ be given and $\epsilon_1, \dots, \epsilon_N$ be \mathbb{R} -valued independent R -subgaussian random variables. Furthermore, define $y_n = f(x_n) + \epsilon_n$ for all $n = 1, \dots, N$.

Consider a Gaussian process $g \sim \mathcal{G}_D(0, k)$ and denote its posterior mean function by μ_N , its posterior covariance function by k_N and its posterior variance by $\sigma_N^2(x) := k_N(x, x)$, w.r.t. to data $(x_1, y_1), \dots, (x_N, y_N)$, assuming in the likelihood independent Gaussian noise with mean zero and variance $\lambda \geq 0$. Then for any $\delta \in (0, 1)$ with

$$\eta_N(x) = R \|(\mathbf{K}_N + \lambda \mathcal{I}_N)^{-1} \mathbf{k}_N(x)\| \times \sqrt{N + 2\sqrt{N} \sqrt{\log \left[\frac{1}{\delta} \right]} + 2 \log \left[\frac{1}{\delta} \right]}$$

one has

$$\mathbb{P} \left[|\mu_N(x) - f(x)| \leq B \sigma_N(x) + \eta_N(x) \forall x \in D \right] \geq 1 - \delta.$$

Proof. (Idea) Similar to the proof of Theorem 1, but use (Hsu et al. 2012, Theorem 2.1) instead of (Chowdhury and Gopalan 2017, Theorem 1). Details are provided in the supplementary material. \square

Using the Nominal Bound

We will now discuss how the previous bounds can be applied. In particular, we will carefully examine potential difficulties that have been identified in similar settings in the literature before.

Kernel Choice and RKHS Norm Bound Theorem 1 and Proposition 2 require an upper bound on the RKHS norm of the target function, i.e., $\|f\|_k \leq B$. In particular, the target function has to be in the RKHS corresponding to the covariance function used in GPR. Since we do not rely on bounds on the maximum information gain, very general kernels can be used together with Theorem 1 and Proposition 2. For example, highly customized kernels can be directly used and no derivation of additional bounds is necessary, as would be the case for (Srinivas et al. 2010, Theorem 6) or (Chowdhury and Gopalan 2017, Theorem 2). In particular, our results support the usage of linearly constrained Gaussian Processes (Jidling et al. 2017; Lange-Hegemann 2018) and related approaches like (Geist and Trimpe 2020). Getting a large enough, yet not too conservative bound on the RKHS norm of the target function can be difficult in general. However, since the kernel encodes prior knowledge, domain knowledge could be used to arrive at such upper bounds. Developing general and systematic methods to transform established domain knowledge into bounds on the RKHS norm is left for future work.

Hyperparameters As with other bounds for GPR the inference of hyperparameters is critical. First of all, an inspection of the proof of (Chowdhury and Gopalan 2017, Theorem 2) shows that the nominal noise variance λ is inde-

pendent of the actual subgaussian noise² (with subgaussian constant R). In particular, Theorem 1 and Proposition 2 hold true for any admissible noise variance, though it is not clear what the optimal choice would be.

In GPR the precise specification of the covariance function is usually not given. In the most common situation a certain kernel class is selected, e.g. SE kernels, and the remaining hyperparameters, e.g. the length-scale of an SE kernel, is inferred by likelihood optimization or in a fully Bayesian setting by using hyperpriors. In both cases the results from above do not directly apply since they rely on *given* correct kernels. Incorporating the hyperparameter inference in results like Theorem 1 in a principled manner is an important open question for future work. As a first step into this direction, we provide robustness results in next section.

Computational Complexity For many interesting applications the sample sizes are small enough, so that standard GPR together with Theorem 1 can be used directly. Furthermore, for control applications the GP model is often used offline and only derived information is used in the actual controller which has restricted computational resources and real-time constraints.

There might be cases where standard GPR can be used, but the calculation of log det is prohibitive. In such a situation it is possible to use powerful approximation methods for the latter quantity. Since quantitative approximation error bounds are available in this situation, e.g. (Han, Malioutov, and Shin 2015; Dong et al. 2017), one can simply add a corresponding margin in the definition of β_N . Additionally, in some cases the log det can be efficiently calculated, in particular for semi-separable kernels (Andersen and Chen 2020), which play an important role in system identification (Chen and Andersen 2020). Note also that the application of Proposition 2 does not require the computation of log det terms.

Additionally, due to the very general nature of Theorem 1 *any* approximation approach for GPR based on subsets of the training data that does not update the kernel parameters can be used and no modifications of the uncertainty bounds are necessary. In particular, the maximum variance selection procedure (Jain et al. 2018; Koller et al. 2018) is compatible with Theorem 1 and Proposition 2.

Finally, any GPR approximation method that does not update the kernel hyperparameters and gives quantitative approximation error estimates for the posterior mean and variance can be used with Theorem 1, where the respective quantities have to be adapted according to the method used. An example of such an approximation method with error estimates is (Huggins et al. 2019).

Bounds under Model Misspecification

As usual in the literature, Theorem 1 and Proposition 2 use the same kernel k for generating the RKHS containing the

²In fact λ in the corresponding (Chowdhury and Gopalan 2017, Theorem 2) has been used as a tuning parameter in (Chowdhury and Gopalan 2017).

ground truth and as a covariance function in GPR. Obviously, in practice it is unlikely that one gets the kernel of the ground truth exactly right. However, it turns out that this is not a big problem. We first note a simple, yet interesting fact that follows immediately from the proof of (Chowdhury and Gopalan 2017, Theorem 2).

Proposition 3. *Consider the situation of Theorem 1, but this time assume that the ground truth f is from another RKHS \tilde{H} . If $\tilde{H} \subseteq H$, and the inclusion $\text{id} : \tilde{H} \rightarrow H$ is continuous with operator norm at most 1, then the result holds true without any modification.*

This simple result can easily be used to verify that for many common situations misspecification of the kernel is not a problem. As an example, we consider the case of the popular isotropic SE kernel.

Proposition 4. *Consider the situation of Theorem 1, but let $f \in \tilde{H}$, where \tilde{H} is the RKHS corresponding to the SE kernel \tilde{k} (on $\emptyset \neq D \subseteq \mathbb{R}^d$) with length scale $0 < \tilde{\gamma}$. Use for the Gaussian Process Regression the SE kernel k with length-scale $0 < \gamma \leq \tilde{\gamma}$. Then Theorem 1 holds true without change.*

Proof. Follows immediately from Proposition 3 and (Steinwart and Christmann 2008, Proposition 4.46). \square

These results tell us that it is not a problem if we do not get the hyperparameter of the isotropic SE kernel right, as long as we *underestimate* the length-scale. This is intuitively clear since a smaller length-scale corresponds to more complex functions. Similar results are known in related contexts, for example (Szabó et al. 2015). Note that Proposition 4 does not imply that one should choose a very small length-scale. The result makes a statement on the validity of the frequentist uncertainty bound from Theorem 1, but not on the size of the uncertainty set. For a similar discussion in a slightly different setting see (Wang, Tuo, and Jeff Wu 2019).

Next, we consider the question what happens when the ground truth is from a different kernel \tilde{k} , such that the corresponding RKHS \tilde{H} is *not* included in the RKHS H corresponding to the covariance function k used in the GPR. Intuitively, since the functions in an RKHS are built from the corresponding reproducing kernel, cf. e.g. (Steinwart and Christmann 2008, Theorem 4.21), Theorem 1 should hold approximately true if the kernels \tilde{k} and k are not too different. This is made precise in the following result. Its proof is provided in the supplementary material.

Theorem 5. *Consider the situation of Theorem 1 but this time assume that the target function f is from the RKHS $(\tilde{H}, \|\cdot\|_{\tilde{k}})$ of a different kernel \tilde{k} such that still $\|f\|_{\tilde{k}} \leq B$ and $\sup_{x,x' \in D} |k(x,x') - \tilde{k}(x,x')| \leq \tilde{\epsilon}$ for some $\tilde{\epsilon} \geq 0$. We then have for any $\delta \in (0, 1)$ that*

$$\mathbb{P}[\|\mu_N(x) - f(x)\| \leq \nu_N(x) \forall N \in \mathbb{N}, x \in D] \geq 1 - \delta$$

where

$$\nu_N(x) = \tilde{\beta}_N(\delta) \sqrt{\sigma_N^2(x) + S_N^2(x) + C_N(x) \|\mathbf{y}_N\|},$$

and $\tilde{\beta}_N(\delta) = \beta_N(\delta, B, R, \lambda + N\tilde{\epsilon})$ as in Theorem 1 and

$$S_N^2(x) = \tilde{\epsilon} + \sqrt{N}\tilde{\epsilon} \|(\mathbf{K}_N + \lambda\mathcal{I}_N)^{-1} \mathbf{k}_N(x)\| + (\sqrt{N}\tilde{\epsilon} + \|\mathbf{k}_N(x)\|) C_N(x) \quad (4)$$

$$C_N(x) = \left(\frac{1}{\lambda} + \|(\mathbf{K}_N + \lambda\mathcal{I}_N)^{-1}\| \right) (\|\mathbf{k}_N(x)\| + \sqrt{N}\tilde{\epsilon}) + \|(\mathbf{K}_N - \lambda\mathcal{I}_N)^{-1}\| \sqrt{N}\tilde{\epsilon} \quad (5)$$

We can interpret Theorem 5 as a robust version (w.r.t. disturbance of the kernel) of Theorem 1. By using the bounds from Theorem 5 we can ensure that the resulting uncertainty set contains the ground truth with prescribed probability despite using the wrong kernel.

In Theorem 1 we have a tube around $\mu_N(x)$ of width $\beta_N(\delta, B, R, \lambda)\sigma_N(x)$, whereas in Theorem 5 we have a tube around $\mu_N(x)$ of width

$$\beta_N(\delta, B, R, \lambda + N\tilde{\epsilon}) \sqrt{\sigma_N^2(x) + S_N^2(x) + C_N(x) \|\mathbf{y}_N\|}$$

Because of the uncertainty or disturbance in the kernel we have to increase the nominal noise variance (used in the nominal noise model in GPR) from λ to $\lambda + N\tilde{\epsilon}$, increase the nominal posterior standard deviation from $\sigma_N(x)$ to $\sqrt{\sigma_N^2(x) + S_N^2(x)}$ and add an offset to the width of the tube of $C_N(x) \|\mathbf{y}_N\|$. In particular, the uncertainty set now depends on the measured values \mathbf{y}_N . Note that $C_N(x)$ and $S_N(x)$ depend on the input, but if necessary this dependence can be easily removed by finding an upper bound on $\|\mathbf{k}_N(x)\|$. An interesting observation is that even if $\sigma_N(x) = 0$, the width of the tube around $\mu_N(x)$ has nonzero width. Intuitively this is clear since in general it can happen that $f \notin H$, but $\mu_N \in H$ by construction. Finally, a robust version of Proposition 2 can be derived similarly to Theorem 5, see the supplementary material.

Numerical Experiments

We now test the theoretical results in numerical experiments using synthetic data, where the ground truth is known. First, we investigate the frequentist behaviour of the uncertainty bounds. The general setup consists of generating a function from an RKHS with known RKHS norm (the ground truth), sampling a finite data set, running GPR on this data set and then evaluating our uncertainty bounds. This learning process is then repeated on many sampled data sets and we check how often the uncertainty bounds are violated. To the best of our knowledge, these types of experiments have actually not been done before, since usually only the final algorithms using the uncertainty bounds have been tested, e.g. in (Srinivas et al. 2010; Chowdhury and Gopalan 2017), or no ground truth from an RKHS has been used, cf. (Berkenkamp 2019). Furthermore, we would like to remark that the method for generating the ground truth can result in misleading judgements of the theoretical result. The latter has to hold for all ground truths, but the generating method can introduce a bias, e.g. only generating functions of a certain shape. To the best of our knowledge, this issue has not been pointed out before.

Second, we demonstrate the practicality and usefulness of our uncertainty bounds by applying them to concrete example from robust control.

Frequentist Behavior of Uncertainty Bounds

Setup of the Experiments Unless noted otherwise, for each of the following experiments we use $D = [-1, 1]$ as input space and generate randomly 50 functions (the ground truths) from an RKHS and evaluate these on a grid 1000 equidistant points from D . We generate the functions by randomly selecting a number of center points $x \in D$ and form a linear combination of kernel functions $k(\cdot, x)$, where k is the kernel of the RKHS. Additionally, for the SE kernel we use the orthonormal basis (ONB) from (Steinwart and Christmann 2008, Section 4.4) with random, normally distributed coefficients to generate functions from the corresponding RKHS. For each function we repeat the following learning instance 10000 times: We sample uniformly 50 inputs, evaluate the ground truth on these inputs and adding normal zero-mean i.i.d. noise with standard deviation (SD) 0.5. We then run GPR on each of the training sets, compute the uncertainty sets (for Theorem 1 this corresponds to the scaling β_{50}) and check (on the equidistant grid of D) whether the resulting uncertainty set contains the ground truth.

Nominal Setting Consider the case that the covariance function used in GPR and the kernel corresponding to the RKHS of the target function coincide. First, we test the nominal bound from Theorem 1 with SE and Matern kernels, respectively, for $\delta = 0.1, 0.01, 0.001, 0.0001$. The resulting scalings β_{50} are shown in Table 1. As can be seen there, the scalings are reasonably small, roughly in the range of heuristics used in the literature, cf. (Berkenkamp 2019). Furthermore, the graph of the target function is fully included in the uncertainty set in all repetitions. This is illustrated in Figure 1 (LEFT), where an example instance of this experiment is shown. As can be clearly seen there, the posterior mean (blue solid line) is well within the uncertainty set from Theorem 1 (with $\delta = 0.01$), which is not overly conservative.

Misspecified Setting We now consider misspecification of the kernel, i.e. different kernels are used for generating the ground truth and as prior covariance function in GPR. As an example we use the SE kernel with different length-scales and use the ONB from (Steinwart and Christmann 2008, Section 4.4) to generate the RKHS functions. We start with the *benign* setting from Proposition 3, where the RKHS corresponding to the covariance function used in GPR contains the RKHS of the target function. For this, identical settings as in our first experiment above are used, but now we generate functions from the SE kernel with length-scale 0.5 and use the SE kernel with length-scale 0.2 in GPR. This corresponds to the benign setting according to Proposition 4. As expected, we find the same results as above, i.e. the uncertainty set fully contains the ground truth in all repetitions. Furthermore, the scalings β_{50} are roughly of the same

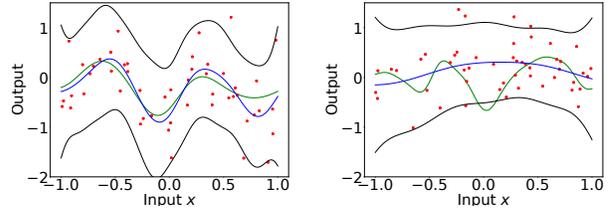


Figure 1: LEFT (Nominal setting): Example function (green) from SE kernel with length-scale 0.5 and RKHS norm 2, learned from 50 samples (red). Shown is the posterior mean (blue) and the uncertainty set from Theorem 1 for $\delta = 0.01$. RIGHT (Misspecified setting): Example function from SE kernel with length-scale 0.2, learned with GPR using SE covariance function with length-scale 0.5. The violation of the uncertainty set is clearly visible.

δ	0.1	0.01	0.001	0.0001
SE	$4.20 \pm .02$	$4.45 \pm .01$	$4.67 \pm .01$	$4.88 \pm .01$
Matern	$4.33 \pm .02$	$4.57 \pm .02$	$4.78 \pm .01$	$4.98 \pm .01$

Table 1: β_{50} in nominal setting (mean \pm standard deviation over all repetitions)

size as in the nominal case, cf. Table 2 (upper row).

Next, we investigate the *problematic* setting where the RKHS corresponding to the covariance function used in GPR does not contain the RKHS of the target function anymore. As an example we use again the SE kernel, but now with length-scale 0.2 for generating RKHS functions and the SE kernel with length-scale 0.5 for GPR. We found a considerable number of function instances where the bounds from Theorem 1 were violated with higher frequency than δ . More precisely, for a given function the tube of width $\beta_{50}\sigma_{50}(x)$ around $\mu_{50}(x)$ does not fully contain the function in more than $\delta \times 10000$ of the learning instances. This happened for 2, 6, 12, 13 out of 50 functions for $\delta = 0.1, 0.01, 0.001, 0.0001$, respectively.

Interestingly, when performing this experiment using the standard approach of generating functions from the RKHS based on linear combinations of kernels, we did not find functions that violated the uncertainty bounds more often than prescribed. This reaffirms our introductory remark that the method generating the test targets can lead to wrong judgements of the theoretical results.

The results of the previous two experiments indicate that a model misspecification of the kernel can be a problem and a robust result like Theorem 5 is necessary. Indeed, a repetition of the last experiment with the uncertainty set from Theorem 5 instead of Theorem 1 resulted in all functions being contained in the uncertainty set in all repetitions. An inspection of the average uncertainty set widths in Table 3 indicates some conservatism.

Control Example

We now show the usefulness of our results for robust control by applying it to a concrete, existing learning-based control method. Due to space constraints only a brief description of

δ	0.1	0.01	0.001	0.0001
Benign	4.20 \pm .02	4.45 \pm .01	4.67 \pm .01	4.88 \pm .01
Problematic	3.88 \pm .01	4.16 \pm .01	4.41 \pm .01	4.64 \pm .01

Table 2: β_{50} in the misspecified setting (mean \pm standard deviation over all repetitions)

δ	0.1	0.01	0.001
Mean	94.99 \pm 5.61	95.84 \pm 5.61	96.67 \pm 5.61
SD	8.38 \pm 2.16	8.46 \pm 2.17	8.53 \pm 2.19

Table 3: Width of robust uncertainty set (mean \pm SD of average width)

the example can be given here. For more details and discussions we refer to the supplementary material.

As an example, we choose the algorithm from (Soloperto et al. 2018) which is a learning-based Robust Model Predictive Control (RMPC) approach that comes with rigorous control-theoretic guarantees. We refer to (Rawlings, Mayne, and Diehl 2017, Chapter 3) for background on RMPC and to (Hewing et al. 2019) for a recent survey on related learning-based control methods. We follow (Soloperto et al. 2018) and consider the discrete-time system

$$\begin{bmatrix} x_1^+ \\ x_2^+ \end{bmatrix} = \begin{bmatrix} 0.995 & 0.095 \\ -0.095 & 0.900 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + \begin{bmatrix} 0.048 \\ 0.95 \end{bmatrix} u + \begin{bmatrix} 0 \\ -r(x_2) \end{bmatrix} \quad (6)$$

modelling a mass-spring-damper system with some nonlinearity r (this could be interpreted as a friction term). The goal is the stabilization of the origin subject to the state and control constraints $\mathbb{X} = [-10, 10] \times [-10, 10]$ and $\mathbb{U} = [-3, 3]$, as well as minimizing a quadratic cost.

The approach from (Soloperto et al. 2018) performs this task by interpreting (6) as a linear system with disturbance, given by the nonlinearity r , whose graph is a-priori known to lie in the set $\mathbb{W}_0 = [-10, 10] \times [-7, 7]$. The nonlinearity is assumed to be unknown and has to be learned from data. The RMPC algorithm requires as an input disturbance sets $\mathbb{W}(x)$ such that $(0 \quad -r(x_2))^\top \in \mathbb{W}(x)$ for all $x \in \mathbb{X}$, which are in turn used to generate tightened nominal constraints ensuring robust constraint satisfaction. Furthermore, the tighter the sets $\mathbb{W}(x)$ are, the better is the performance of the algorithm. We now randomly generate the function r (which will be our ground truth) from the RKHS with kernel $k(x, x') = 4 \exp\left(-\frac{(x-x')^2}{2 \times 0.8^2}\right)$ with RKHS norm 2. Following (Soloperto et al. 2018), we uniformly sample 100 partial states $x_2 \in [-10, 10]$, evaluate r at these and add i.i.d. Gaussian noise with a standard deviation of 0.01 to it. The unknown function is then learned using GPR from this data set. Our results then lead to an uncertainty set of the form $\mathbb{W}(x) = [\mu_{100}(x_2) - \beta_{100}\sigma_{100}(x_2), \mu_{100}(x_2) + \beta_{100}\sigma_{100}(x_2)]$, with β_{100} from Theorem 1 for $\delta := 0.001$. In particular, with probability at least $1 - \delta$ we can guarantee that $r(x_2) \in \mathbb{W}(x)$ holds for all $x \in \mathbb{X}$. Figure 2 shows the resulting tightened state constraints for an MPC horizon of 9. Clearly, the state constraint sets from the learned uncertainty sets are much larger. Furthermore, in contrast to previous work, we can guarantee that the RMPC controller using these tightened state constraints retains all control-theoretic

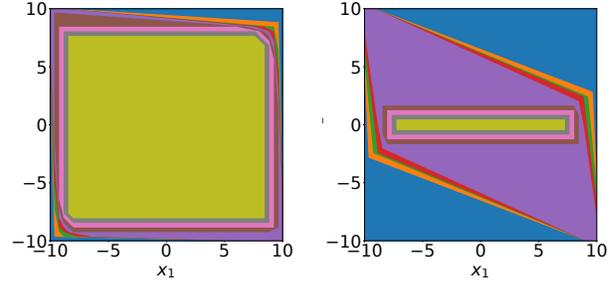


Figure 2: Tightened state constraint sets \mathbb{Z}_k for $k = 0, \dots, 9$. Computed from a-priori uncertainty set \mathbb{W}_0 (LEFT) and learned uncertainty sets $\mathbb{W}(x)$ (RIGHT).

guarantees with probability at least $1 - \delta$. In the present case we can ensure state and control constraint satisfaction, input-to-state stability and convergence to a neighborhood of the origin, with probability at least $1 - \delta$. This follows immediately from (Soloperto et al. 2018, Theorem 1), since the ground truth r is covered by the uncertainty sets $\mathbb{W}(x)$ with this probability. Note that no changes to the existing RMPC scheme were necessary and the control-theoretic guarantees were retained with prescribed high probability.

Conclusion

We discussed the importance of frequentist uncertainty bounds for Gaussian Process Regression and improved existing results in this kind. By aiming only at a-posteriori bounds we were able to provide rigorous *and* practical uncertainty results. Our bounds can be explicitly evaluated and are sharp enough to be useful for concrete applications, as demonstrated with numerical experiments. Furthermore, we also introduced robust versions that work despite certain model mismatches, which is an important concern in real-world applications. We see the present work as a starting point for further developments, in particular, domain-specific versions of our results and more specific and less conservative robustness results.

Acknowledgments. We would like to thank Sayak Ray Chowdhury and Tobias Holicki for helpful discussions, Steve Heim for useful comments on a draft of this work, and Raffaele Soloperto for providing the Matlab code from (Soloperto et al. 2018). Furthermore, we would like to thank the reviewers for their helpful and constructive feedback. Funded by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy - EXC 2075 – 390740016 and in part by the Cyber Valley Initiative. We acknowledge the support by the Stuttgart Center for Simulation Science (SimTech).

References

- Abbasi-Yadkori, Y. 2013. *Online learning for linearly parametrized control problems*. Ph.D. thesis, University of Alberta.
- Andersen, M. S.; and Chen, T. 2020. Smoothing Splines and Rank Structured Matrices: Revisiting the Spline Kernel. *SIAM Journal on Matrix Analysis and Applications* 41(2): 389–412.
- Åström, K. J.; and Murray, R. M. 2010. *Feedback systems: an introduction for scientists and engineers*. Princeton university press.
- Beckers, T.; Umlauft, J.; and Hirche, S. 2018. Mean square prediction error of misspecified Gaussian process models. In *2018 IEEE Conference on Decision and Control (CDC)*, 1162–1167. IEEE.
- Berkenkamp, F. 2019. *Safe Exploration in Reinforcement Learning: Theory and Applications in Robotics*. Ph.D. thesis, ETH Zurich.
- Berkenkamp, F.; Moriconi, R.; Schoellig, A. P.; and Krause, A. 2016. Safe learning of regions of attraction for uncertain, nonlinear systems with gaussian processes. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, 4661–4666. IEEE.
- Berkenkamp, F.; Schoellig, A. P.; and Krause, A. 2016. Safe controller optimization for quadrotors with Gaussian processes. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 491–496. IEEE.
- Berkenkamp, F.; Turchetta, M.; Schoellig, A.; and Krause, A. 2017. Safe model-based reinforcement learning with stability guarantees. In *Advances in neural information processing systems*, 908–918.
- Berlinet, A.; and Thomas-Agnan, C. 2011. *Reproducing kernel Hilbert spaces in probability and statistics*. Springer Science & Business Media.
- Calandriello, D.; Carratino, L.; Lazaric, A.; Valko, M.; and Rosasco, L. 2019. Gaussian Process Optimization with Adaptive Sketching: Scalable and No Regret. In *Conference on Learning Theory*, 533–557.
- Chen, T.; and Andersen, M. 2020. On Semiseparable Kernels and Efficient Computation of Regularized System Identification and Function Estimation. *IFAC-V 2020*.
- Chowdhury, S. R.; and Gopalan, A. 2017. On Kernelized Multi-armed Bandits. In *International Conference on Machine Learning*, 844–853.
- Dong, K.; Eriksson, D.; Nickisch, H.; Bindel, D.; and Wilson, A. G. 2017. Scalable log determinants for Gaussian process kernel learning. In *Advances in Neural Information Processing Systems*, 6327–6337.
- Geist, A. R.; and Trimpe, S. 2020. Learning Constrained Dynamics with Gauss Principle adhering Gaussian Processes. In *Proceedings of the 2nd Conference on Learning for Dynamics and Control*, 225–234.
- Han, I.; Malioutov, D.; and Shin, J. 2015. Large-scale log-determinant computation through stochastic Chebyshev expansions. In *International Conference on Machine Learning*, 908–917.
- Helwa, M. K.; Heins, A.; and Schoellig, A. P. 2019. Provably robust learning-based approach for high-accuracy tracking control of lagrangian systems. *IEEE Robotics and Automation Letters* 4(2): 1587–1594.
- Hewing, L.; Wabersich, K. P.; Menner, M.; and Zeilinger, M. N. 2019. Learning-Based Model Predictive Control: Toward Safe Learning in Control. *Annual Review of Control, Robotics, and Autonomous Systems* 3.
- Hsu, D.; Kakade, S.; Zhang, T.; et al. 2012. A tail inequality for quadratic forms of subgaussian random vectors. *Electronic Communications in Probability* 17.
- Huggins, J. H.; Campbell, T.; Kasprzak, M.; and Broderick, T. 2019. Scalable Gaussian Process Inference with Finite-data Mean and Variance Guarantees. In *The 22nd International Conference on Artificial Intelligence and Statistics*, 796–805.
- Jain, A.; Nghiem, T.; Morari, M.; and Mangharam, R. 2018. Learning and control using Gaussian processes. In *2018 ACM/IEEE 9th International Conference on Cyber-Physical Systems (ICCPS)*, 140–149. IEEE.
- Jidling, C.; Wahlström, N.; Wills, A.; and Schön, T. B. 2017. Linearly constrained Gaussian processes. In *Advances in Neural Information Processing Systems*, 1215–1224.
- Kanagawa, M.; Hennig, P.; Sejdinovic, D.; and Sriperumbudur, B. K. 2018. Gaussian processes and kernel methods: A review on connections and equivalences. *arXiv preprint arXiv:1807.02582*.
- Kandasamy, K.; Schneider, J.; and Póczos, B. 2015. High dimensional Bayesian optimisation and bandits via additive models. In *International Conference on Machine Learning*, 295–304.
- Kocijan, J. 2016. *Modelling and control of dynamic systems using Gaussian process models*. Springer.
- Koller, T.; Berkenkamp, F.; Turchetta, M.; and Krause, A. 2018. Learning-based model predictive control for safe exploration. In *2018 IEEE Conference on Decision and Control (CDC)*, 6059–6066. IEEE.
- Lange-Hegermann, M. 2018. Algorithmic linearly constrained Gaussian processes. In *Advances in Neural Information Processing Systems*, 2137–2148.
- Lederer, A.; Umlauft, J.; and Hirche, S. 2019. Uniform Error Bounds for Gaussian Process Regression with Application to Safe Control. In *Advances in Neural Information Processing Systems*, 657–667.
- Liu, M.; Chowdhury, G.; Da Silva, B. C.; Liu, S.-Y.; and How, J. P. 2018. Gaussian processes for learning and control: A tutorial with examples. *IEEE Control Systems Magazine* 38(5): 53–86.

- Maddalena, E. T.; Scharnhorst, P.; and Jones, C. N. 2020. Deterministic error bounds for kernel-based learning techniques under bounded noise. *arXiv preprint arXiv:2008.04005*.
- Murphy, K. P. 2012. *Machine learning: a probabilistic perspective*. MIT press.
- Rasmussen, C. E.; and Williams, C. K. 2006. *Gaussian Processes for Machine Learning*. The MIT Press.
- Rawlings, J. B.; Mayne, D. Q.; and Diehl, M. 2017. *Model predictive control: theory, computation, and design*, volume 2. Nob Hill Publishing Madison, WI.
- Shahriari, B.; Swersky, K.; Wang, Z.; Adams, R. P.; and De Freitas, N. 2015. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE* 104(1): 148–175.
- Skogestad, S.; and Postlethwaite, I. 2007. *Multivariable feedback control: analysis and design*, volume 2. Wiley New York.
- Soloperto, R.; Müller, M. A.; Trimpe, S.; and Allgöwer, F. 2018. Learning-based robust model predictive control with state-dependent uncertainty. *IFAC-PapersOnLine* 51(20): 442–447.
- Srinivas, N.; Krause, A.; Kakade, S.; and Seeger, M. 2010. Gaussian process optimization in the bandit setting: no regret and experimental design. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, 1015–1022.
- Steinwart, I.; and Christmann, A. 2008. *Support vector machines*. Springer Science & Business Media.
- Szabó, B.; Van Der Vaart, A. W.; van Zanten, J.; et al. 2015. Frequentist coverage of adaptive nonparametric Bayesian credible sets. *The Annals of Statistics* 43(4): 1391–1428.
- Umlauft, J.; Beckers, T.; Kimmel, M.; and Hirche, S. 2017. Feedback linearization using Gaussian processes. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, 5249–5255. IEEE.
- Wang, W.; Tuo, R.; and Jeff Wu, C. 2019. On prediction properties of kriging: Uniform error bounds and robustness. *Journal of the American Statistical Association* 1–27.