Bayesian optimisation under uncertain inputs

Rafael Oliveira[1] **Lionel Ott**[1] **Fabio Ramos**[1,2]
[1]School of Computer Science, The University of Sydney, Australia. [2]NVIDIA, USA.

Abstract

Bayesian optimisation (BO) has been a successful approach to optimise functions which are expensive to evaluate and whose observations are noisy. Classical BO algorithms, however, do not account for errors about the location where observations are taken, which is a common issue in problems with physical components. In these cases, the estimation of the actual query location is also subject to uncertainty. In this context, we propose an upper confidence bound (UCB) algorithm for BO problems where both the outcome of a query and the true query location are uncertain. The algorithm employs a Gaussian process model that takes probability distributions as inputs. Theoretical results are provided for both the proposed algorithm and a conventional UCB approach within the uncertain-inputs setting. Finally, we evaluate each method's performance experimentally, comparing them to other input noise aware BO approaches on simulated scenarios involving synthetic and real data.

1 Introduction

Bayesian optimisation (BO) (Shahriari et al., 2016) is a technique to find the global optimum of functions that are unknown, expensive to evaluate, and whose output observations are possibly noisy. In this sense, BO has been applied across different fields to a wide class of problems, including hyper-parameter tuning (Snoek et al., 2012), policy search (Wilson et al., 2014), environmental monitoring (Marchant and Ramos, 2012), robotic grasping (Nogueira et al., 2016), etc. Although taking into account that we might have a noisy observation of the function's output, conventional BO approaches assume that the function has been sampled precisely at a specified query location. While this is true for many applications of BO, there are certain problems, especially in areas of robotics and process control, in which this assumption typically does not hold.

As an illustration, consider a problem where we are interested in finding the peak of an environmental process $f(\mathbf{x})$ over a region $\mathcal{S} \subset \mathbb{R}^d$. To this end, we send a mobile robot to different target locations $\mathbf{x}_t \in \mathcal{S}$ to observe the process. Unfortunately, due to localisation uncertainty and motion control errors, *execution noise* prevents the robot from reaching the planned target location exactly. Instead, after each query, the robot provides us with an estimate of its actual location $\tilde{\mathbf{x}}_t$ via a probability distribution P_t^L, which takes into account *localisation noise*, as depicted in Figure 1. From each query, we obtain a noisy observation of the environmental process $y_t = f(\tilde{\mathbf{x}}_t) + \zeta_t$, where ζ_t is an independent noise term. In this scenario, both the function inputs $\tilde{\mathbf{x}}_t$, i.e. query locations, and outputs $f(\tilde{\mathbf{x}}_t)$ are not directly observable.

This paper investigates optimisation problems where input noise affects both the execution of a query and the estimation of its true location. In particular, we analyse the standard BO approach when employing the improved Gaussian process upper-confidence bound (IGP-UCB) (Chowdhury and Gopalan, 2017) algorithm under input noise, and we propose the *uncertain-inputs Gaussian process upper confidence bound* (uGP-UCB) algorithm. The latter is equipped with a GP model that takes probability distributions as inputs in a similar framework to Oliveira et al. (2017). We apply kernel embeddings techniques (Muandet et al., 2016) to obtain the first theoretical results for BO under uncertain inputs, bounding the regret of both uGP-UCB and IGP-UCB. In addition, experiments provide empirical performance evaluations of different BO approaches to problems involving input noise.

Proceedings of the 22[nd] International Conference on Artificial Intelligence and Statistics (AISTATS) 2019, Naha, Okinawa, Japan. PMLR: Volume 89.

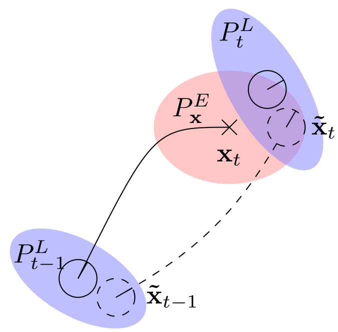

Figure 1: At time $t-1$, the robot is estimated to be at some $\tilde{\mathbf{x}}_{t-1} \sim P_{t-1}^L$. It is then sent to target location $\mathbf{x}_t$. However, due to uncertainty in the query execution, represented by $P_{\mathbf{x}}^E$, the robot actually ends up at another location $\tilde{\mathbf{x}}_t$, whose belief distribution, according to the localisation system, is represented by P_t^L. The robot's true locations and true path are indicated by the dashed lines.

2 Related work

Recently several BO approaches that deal with problems where the execution of queries to an objective function is affected by uncertainty have been proposed. Nogueira et al. (2016) presented a method that applies the unscented transform (Wan and van der Merwe, 2000) to query BO's acquisition function. By considering a stochastic query execution process, the method is able to find robust solutions to robotics problems such as grasping. Another approach to handle query uncertainty is presented in Pearce and Branke (2017) to optimise stochastic simulations. In that case, query uncertainty refers to imperfect knowledge about input variates for a simulation model (Lam, 2016). Pearce and Branke apply Monte Carlo integration to marginalise out input variates that are unknown when querying BO's acquisition function. In broader terms, all of these problems can be described as optimising an integrated cost function, where one may instead use a GP prior over the integrated function (Beland and Nair, 2017; Toscano-Palmerin and Frazier, 2018). Contrasted to uGP-UCB, however, the approaches mentioned above only deal with independent and identically distributed input noise and mostly offer no known theoretical guarantees. In addition, the data points in their GP datasets are only point estimates, instead of distributions as used in this paper.

Another BO approach is presented in Oliveira et al. (2017), which employed a Gaussian process (GP) model that takes probability distributions directly as inputs (Girard, 2004; Dallaire et al., 2011). However, Oliveira et al.'s method intent is to learn a model of the objective function with a robot, while minimising travelled distance, not as an optimisation framework.

Problems like the one illustrated in Figure 1 can also be related to partially-observable Markov decision processes (POMDPs) (Marchant and Ramos, 2014; Ling et al., 2016). This paper, however, is concerned with a general optimisation setup.

3 Problem formulation

We consider an optimisation problem where an algorithm sequentially selects target locations $\mathbf{x}_t$ within a compact search space $\mathcal{S} \subset \mathcal{X}$ at which to query a function $f : \mathcal{X} \to \mathbb{R}$, seeking its global optimum. In addition, the query execution itself is a stochastic process, leading the query to be made at some $\tilde{\mathbf{x}}_t | \mathbf{x}_t \sim P_{\mathbf{x}}^E$, instead.

How close the algorithm is to the global optimum can be measured in terms of regret. In a bandits optimisation setting, the *instantaneous regret* suffered by a maximisation algorithm for a choice of target $\mathbf{x}_t$ in our problem is given by:

$$\tilde{r}_t = \max_{\mathbf{x} \in \mathcal{S}} f(\mathbf{x}) - f(\tilde{\mathbf{x}}_t) \ . \quad (1)$$

In the deterministic-inputs case, the algorithmic design goal is to minimise cumulative regret, ensuring that the algorithm eventually hits the global optimum of f (Srinivas et al., 2010; Bull, 2011). However, as $\tilde{\mathbf{x}}_t$ is subject to noise, one can attempt to minimise the expected regret, which is such that:

$$\mathbb{E}[\tilde{r}_t | \mathbf{x}_t] = \max_{\mathbf{x} \in \mathcal{S}} f(\mathbf{x}) - \mathbb{E}[f(\tilde{\mathbf{x}}_t) | \mathbf{x}_t] = \rho_E + \hat{r}_t \ , \quad (2)$$

where:

$$\rho_E := \max_{\mathbf{x} \in \mathcal{S}} f(\mathbf{x}) - \max_{\mathbf{x} \in \mathcal{S}} \mathbb{E}[f(\tilde{\mathbf{x}}) | \mathbf{x}] \quad (3)$$

$$\hat{r}_t := \max_{\mathbf{x} \in \mathcal{S}} \mathbb{E}[f(\tilde{\mathbf{x}}) | \mathbf{x}] - \mathbb{E}[f(\tilde{\mathbf{x}}_t) | \mathbf{x}_t] \ . \quad (4)$$

Here ρ_E is a constant, representing the difference between the maximum of the function and the maximum value any algorithm is expected to reach under the query execution uncertainty. However, $\hat{r}_t$ is controllable via the algorithm's choices of $\mathbf{x}_t$ and is associated with the goal of finding:

$$\mathbf{x}^* \in \operatorname*{argmax}_{\mathbf{x} \in \mathcal{S}} \mathbb{E}[f(\tilde{\mathbf{x}}) | \mathbf{x}] \ , \quad (5)$$

which defines a target location that maximises the expected value of the function f under the querying process noise. As defined, $\mathbf{x}^*$ minimises the expected regret to a lower bound given by ρ_E and defines an optimum location which is robust to execution noise. Therefore, we call $\hat{r}_t$ the *uncertain-inputs* *regret*. Similarly, we also define the *uncertain-inputs cumulative regret* $\hat{R}_n = \sum_{t=1}^{n} \hat{r}_t$. With these definitions, an algorithm whose uncertain-inputs cumulative regret $\hat{R}_n$ grows sub-linearly achieves a minimum on the expected regret:

$$\begin{aligned} \lim_{n\to\infty} \min_{t\leq n} \mathbb{E}[\tilde{r}_t | \mathbf{x}_t] &= \rho_E + \lim_{n\to\infty} \min_{t\leq n} \hat{r}_t \\ &\leq \rho_E + \limsup_{n\to\infty} \frac{\hat{R}_n}{n} = \rho_E \ . \end{aligned} \tag{6}$$

Distribution assumptions: We are assuming that the query location distribution $P^E_{\mathbf{x}}$ marginalises over all other variables that could affect the querying process, such as starting points and effects from the environment that the agent is in. In addition, the true $P^E_{\mathbf{x}}$ might be unknown. However, *after* each query, we assume that a distribution P^L_t estimating the true query location is available. These probability distributions are illustrated by the example in Figure 1 for a robotics case.

For each $\mathbf{x}_t$, the algorithm is provided with observations $y_t = f(\tilde{\mathbf{x}}_t) + \zeta_t$, where ζ_t is σ_ζ-sub-Gaussian observation noise, for some $\sigma_\zeta \geq 0$. Sub-Gaussian random variables can be thought of as any random variable whose tail distribution decays at least as fast as a Gaussian. Both Gaussian and bounded random variables fall in this category (Boucheron et al., 2013).

Regularity assumptions: We assume $f : \mathcal{X} \to \mathbb{R}$ to be an element of $\mathcal{H}_k$, which is a reproducing kernel Hilbert space (RKHS) (Schölkopf and Smola, 2002). For a given positive-definite kernel $k : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$, a RKHS $\mathcal{H}_k$ is a Hilbert space of functions with inner product $\langle \cdot, \cdot \rangle_k$ and norm $\|\cdot\|_k = \sqrt{\langle \cdot, \cdot \rangle_k}$ such that $f(\mathbf{x}) = \langle f, k(\cdot, \mathbf{x}) \rangle_k$, for any $f \in \mathcal{H}_k$ and any $\mathbf{x} \in \mathcal{X}$. We assume k is continuous and bounded on $\mathcal{X} \times \mathcal{X}$, with $k(\mathbf{x}, \mathbf{x}) \leq 1, \forall \mathbf{x} \in \mathcal{X}$, and that $\|f\|_k \leq b$ for the objective function in Equation 5, where $b > 0$ is known.[1] When not explicitly mentioned, assume an Euclidean domain for f, i.e. $\mathcal{X} \subseteq \mathbb{R}^d$, $d \in \mathbb{N}$.

4 The uGP-UCB algorithm

This section describes a method for Bayesian optimisation under uncertain inputs. The section starts by presenting a Gaussian process that allows direct modelling of objectives defined in terms of expectations. This GP approach is then applied to derive a BO algorithm named *uncertain-inputs Gaussian process upper confidence bound* (uGP-UCB), presented in the second part of this section.

[1] These assumptions are met by most of the popular kernels in BO and are common in the regret bounds literature.

4.1 Gaussian process priors with uncertain inputs

To extend BO to the case where query locations $\mathbf{x}$ are uncertain, we can redefine the objective in Equation 5 as a function of the query probability distributions. Let $\mathcal{P}$ denote the set containing all probability measures on $\mathcal{X} \subseteq \mathbb{R}^d$. With $f \in \mathcal{H}_k$, we can define the map:

$$\begin{aligned} \psi : \mathcal{P} &\to \mathcal{H}_k \\ P &\mapsto \int_{\mathcal{X}} k(\cdot, \mathbf{x}) \, \mathrm{d}P(\mathbf{x}) \ . \end{aligned} \tag{7}$$

For any $\mathcal{X}$-valued random variable $\tilde{\mathbf{x}}$ distributed according to $P \in \mathcal{P}$, we then have that:

$$\mathbb{E}_P[f] := \mathbb{E}[f(\tilde{\mathbf{x}})] = \langle \psi_P, f \rangle_k, \quad \forall f \in \mathcal{H}_k \ , \tag{8}$$

where $\psi_P := \psi(P)$. If the kernel k is characteristic, such as radial kernels (Sriperumbudur et al., 2011), ψ is injective, defining a one-to-one relationship between measures in $\mathcal{P}$ and elements of $\mathcal{H}_k$. Therefore, ψ is referred to as the mean map, and ψ_P as the kernel mean embedding of P (Muandet et al., 2016).

Using ψ as defined in Equation 7, one can construct kernels over the set of probability measures $\mathcal{P}$. In particular, for any $P, P' \in \mathcal{P}$, we have that:

$$\hat{k}(P, P') := \langle \psi_P, \psi_{P'} \rangle_k = \int_{\mathcal{X}} \int_{\mathcal{X}} k(\mathbf{x}, \mathbf{x}') \, \mathrm{d}P(\mathbf{x}) \, \mathrm{d}P'(\mathbf{x}') \tag{9}$$

defines a positive-definite kernel over $\mathcal{P}$ (Muandet et al., 2012). Notice that in this formulation, even if we have inputs representing the same random variable $\tilde{\mathbf{x}} \sim P$, we have $\hat{k}(P, P) = \langle \psi_P, \psi_P \rangle \neq \mathbb{E}[k(\tilde{\mathbf{x}}, \tilde{\mathbf{x}})]$, which is then different from other kernel formulations for models with uncertain inputs (Dallaire et al., 2011).

The kernel in Equation 9 is associated with a RKHS $\mathcal{H}_{\hat{k}}$ containing functions over the space of probability measures $\mathcal{P}$. Besides the linear kernel in Equation 9, many other kernels on $\mathcal{P}$ can be defined via ψ, e.g. radial kernels using $\|\psi_P - \psi_{P'}\|_k$ as a metric on $\mathcal{P}$ (Muandet et al., 2012). However, the simple kernel in Equation 9 provides us with a useful property to model the objective in Equation 5, as presented next.

Lemma 1 (Expected function)**.** *Any $f \in \mathcal{H}_k$ is continuously mapped to a corresponding $\hat{f} \in \mathcal{H}_{\hat{k}}$, which is such that:*

$$\begin{aligned} \forall P \in \mathcal{P}, \quad \hat{f}(P) &= \mathbb{E}_P[f] \\ \|\hat{f}\|_{\hat{k}} &= \|f\|_k \ . \end{aligned} \tag{10}$$

The mapping $f \mapsto \hat{f}$ constitutes an isometric isomorphism between $\mathcal{H}_k$ and $\mathcal{H}_{\hat{k}}$.

Proof sketch. The proof follows from the fact that Dirac measures $D_{\mathbf{x}}$, for $\mathbf{x} \in \mathcal{X}$, are also probability measures in $\mathcal{P}$. Since $k(\mathbf{x}, \mathbf{x}') = \hat{k}(D_{\mathbf{x}}, D_{\mathbf{x}'})$, $\forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}$, we can define a bijective linear map between $\mathcal{H}_k$ and $\mathcal{H}_{\hat{k}}$ that preserves norms. A complete proof is presented in the appendix. □

As a positive-definite kernel, $\hat{k}$ defines the covariance function of a Gaussian process $\mathrm{GP}(0, \hat{k})$ modelling functions over $\mathcal{P}$. This GP model can then be applied to learn $\hat{f}$ from a given set of observations $\mathcal{D}_n = \{(P_i, y_i)\}_{i=1}^n$, as in Girard (2004). Under a zero-mean GP assumption, the value of $\hat{f}(P_*)$ for a given $P_* \in \mathcal{P}$ follows a Gaussian posterior distribution with mean and variance given by:

$$\hat{\mu}_n(P_*) = \hat{\mathbf{k}}_n(P_*)^\mathsf{T}(\hat{\mathbf{K}}_n + \lambda \mathbf{I})^{-1}\mathbf{y}_n \ , \tag{11}$$

$$\hat{k}_n(P, P') = \hat{k}(P, P') - \hat{\mathbf{k}}_n(P)^\mathsf{T}(\hat{\mathbf{K}}_n + \lambda \mathbf{I})^{-1}\hat{\mathbf{k}}_n(P') \tag{12}$$

$$\hat{\sigma}_n^2(P_*) = \hat{k}_n(P_*, P_*) \ , \tag{13}$$

where $\mathbf{y}_n := [y_1, \dots, y_n]^\mathsf{T}$, $[\hat{\mathbf{K}}_n]_{ij} = \hat{k}(P_i, P_j)$ and $\hat{\mathbf{k}}_n(P_*) := [\hat{k}(P_*, P_1), \dots, \hat{k}(P_*, P_n)]^\mathsf{T}$. For a $\hat{f} \in \mathcal{H}_{\hat{k}}$, we have that $\hat{f}$ is generally not a sample from the GP (Rasmussen and Williams, 2006, p. 131). However, we always have $\hat{\mu}_n \in \mathcal{H}_{\hat{k}}$ by definition, allowing the GP to learn an approximation for $\hat{f}$. Therefore, in these equations, $\lambda \geq 0$ is simply a parameter that is not necessarily related to the true observation noise as in usual GP modelling assumptions (Rasmussen and Williams, 2006).

4.2 Upper-confidence bound

Coming back to the problem definition in Equation 5, we consider a function $\hat{f} : \mathcal{P} \to \mathbb{R}$, such that for any $\tilde{\mathbf{x}} \sim P$, $\hat{f}(P) = \mathbb{E}[f(\tilde{\mathbf{x}})]$. The GP model proposed in the previous section allows deriving a BO algorithm to solve the problem in Equation 5. Given a set of past observations $\mathcal{D}_{t-1} = \{(P_i, y_i)\}_{i=1}^{t-1}$, the following defines an upper confidence bound (UCB) acquisition function:

$$h(P|\mathcal{D}_{t-1}) = \hat{\mu}_{t-1}(P) + \beta_t \hat{\sigma}_{t-1}(P) \ , \tag{14}$$

where β_t is a parameter controlling the exploration-exploitation trade-off. The theoretical results in the next section will show that β_t can be set accordingly for $h(P|\mathcal{D}_{t-1})$ to maintain a high-probability upper bound on $\hat{f}$.

Querying the GP model with $\mathbf{x} \mapsto P_{\mathbf{x}}^E$ would allow selecting points $\mathbf{x}_t$ based on an estimate for $\mathbb{E}_{P_{\mathbf{x}_t}^E}[f] := \mathbb{E}[f(\tilde{\mathbf{x}}_t)|\mathbf{x}_t]$. However, in general, the true mapping $\mathbf{x} \mapsto P_{\mathbf{x}}^E$ is unknown. Instead, we use a model $\mathbf{x} \mapsto \hat{P}_{\mathbf{x}}$ whose approximation error $|\mathbb{E}_{P_{\mathbf{x}}^E}[f] - \mathbb{E}_{\hat{P}_{\mathbf{x}}}[f]|$ is small.

Algorithm 1: uGP-UCB

Input: $\mathcal{S}$: search space
n: total number of iterations

1 **for** $t \in \{1, \dots, n\}$ **do**
2 $\quad \mathbf{x}_t = \operatorname{argmax}_{\mathbf{x} \in \mathcal{S}} \hat{\mu}_{t-1}(\hat{P}_{\mathbf{x}}) + \beta_t \hat{\sigma}_{t-1}(\hat{P}_{\mathbf{x}})$
3 $\quad (P_t^L, y_t) \leftarrow$ Sample f at $\tilde{\mathbf{x}}_t | \mathbf{x}_t \sim P_{\mathbf{x}_t}^E$
4 $\quad \mathcal{D}_t = \mathcal{D}_{t-1} \cup \{(P_t^L, y_t)\}$
5 $\mathbf{x}_n^* = \operatorname{argmax}_{\mathbf{x}_t \in \mathcal{D}_n} \hat{\mu}_n(\hat{P}_{\mathbf{x}_t})$

Result: $\mathbf{x}_n^*$

Algorithm 1 presents the uGP-UCB algorithm. Equipped with the acquisition function in Equation 14, at each iteration t, the algorithm selects the target location $\mathbf{x}_t$ that maximises $h(\hat{P}_{\mathbf{x}}|\mathcal{D}_{t-1})$ (line 2). In line 3, the function f is queried at some location $\tilde{\mathbf{x}}_t | \mathbf{x}_t \sim P_{\mathbf{x}_t}^E$. After the query is done, the algorithm is provided with an observation $y_t = f(\tilde{\mathbf{x}}_t) + \zeta_t$ and an independent estimate for $\tilde{\mathbf{x}}_t$ given by P_t^L, as described earlier. In line 4, the GP model is updated with the new observation pair (P_t^L, y_t). This process then repeats for a given number of iterations n. As a result, the algorithm finishes with an estimate of the optimum location $\mathbf{x}^*$ given as the target location with the best estimated outcome $\mathbf{x}_n^*$ (line 5).

5 Theoretical results

This section presents theoretical results bounding the uncertain-inputs regret of the uGP-UCB algorithm and a standard BO approach, IGP-UCB (Chowdhury and Gopalan, 2017), which was not originally designed to handle input noise. The theoretical analysis presented in this paper is mainly based on Chowdhury and Gopalan's results, which are advantageous in the uncertain-inputs setting due to mild assumptions on the observation noise. However, the results in this section also bring new insights into BO methods for problems with uncertain inputs. We refer the reader to the appendix for complete proofs of the next results.

5.1 The uncertain-inputs regret of IGP-UCB

In the uncertain-inputs setting, IGP-UCB selects target locations $\mathbf{x}_t$ by maximising $\mu_{t-1}(\mathbf{x}) + \beta_t \sigma_{t-1}(\mathbf{x})$, where μ_{t-1} and σ_{t-1}^2 are respectively the posterior mean and variance of the deterministic-inputs $\mathrm{GP}(0, k)$ given observations $\{(\mathbf{x}_i, y_i)\}_{i=1}^{t-1}$. For an asymptotic analysis, both the targets $\{\mathbf{x}_t\}_{t=1}^\infty$ and the equivalent observation noise $\{\nu_t\}_{t=1}^\infty$, where $\nu_t := y_t - \mathbb{E}[f(\tilde{\mathbf{x}}_t)|\mathbf{x}_t] \neq \zeta_t$, can be treated as sequences of random variables. At a given round $t \geq 1$, the history $\{\mathbf{x}_i, \nu_i\}_{i=1}^t$ generates a σ-algebra $\mathfrak{F}_t$, and the sequence

$\{\mathfrak{F}_t\}_{t=0}^{\infty}$ defines a filtration (Bauer, 1981). The sub-Gaussian condition on the sequence $\{\nu_t\}_{t=1}^{\infty}$ is then formally defined as:

$$\forall t \geq 1,\ \forall \lambda \in \mathbb{R}, \quad \mathbb{E}[e^{\lambda \nu_t} | \mathfrak{F}_{t-1}] \leq e^{\lambda^2 \sigma_\nu^2 / 2} \text{ (a.s.)} , \quad (15)$$

which denotes an upper bound on a conditional expectation (Bauer, 1981), so that the inequality above is defined as holding *almost surely* (a.s.).

The results in Chowdhury and Gopalan (2017) bound the cumulative regret of IGP-UCB in terms of the maximum information gain:

$$\gamma_n := \max_{\mathcal{Q} \subset \mathcal{S} : |\mathcal{Q}| = n} I(\mathbf{y}_n, \mathbf{g}_n | \mathcal{Q}) , \quad (16)$$

where $I(\mathbf{y}_n, \mathbf{g}_n | \mathcal{Q})$ represents the mutual information between $\mathbf{y}_n = \mathbf{g}_n + \boldsymbol{\nu}'_n$ and $\mathbf{g}_n \sim N(\mathbf{0}, \mathbf{K}_n)$, with $[\mathbf{K}_n]_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$, $\mathbf{x}_i, \mathbf{x}_j \in \mathcal{Q}$ and $\boldsymbol{\nu}'_n \sim N(\mathbf{0}, \lambda \mathbf{I})$. Here $\lambda > 0$ is the same parameter in Equation 11. Considering these definitions, we derive the following.

Theorem 2 (IGP-UCB uncertain-inputs regret). *For any $f \in \mathcal{H}_k$, assume that:*

1. *the mapping $\mathbf{x} \mapsto \mathbb{E}_{P^E_\mathbf{x}}[f]$ defines a function $g \in \mathcal{H}_k(\mathcal{S})$ and $\|g\|_k \leq b$;*
2. *$\forall \mathbf{x} \in \mathcal{S}, \Delta f_{P^E_\mathbf{x}} := f(\tilde{\mathbf{x}}^E) - \mathbb{E}_{P^E_\mathbf{x}}[f]$ is σ_E-sub-Gaussian, for a given $\sigma_E > 0$, where $\tilde{\mathbf{x}}^E \sim P^E_\mathbf{x}$;*
3. *and ζ_t is conditionally σ_ζ-sub-Gaussian.*

Then running IGP-UCB with $\sigma_\nu := \sqrt{\sigma_E^2 + \sigma_\zeta^2}$ and $\beta_t := b + \sigma_\nu \sqrt{2(\gamma_{t-1} + 1 + \log(1/\delta))}$ leads to the same bounds as Theorem 3 in Chowdhury and Gopalan (2017) for the uncertain-inputs cumulative regret of the algorithm. Namely, we have that:

$$\mathbb{P}\left\{\hat{R}_n \in \mathcal{O}\left(b\sqrt{n\gamma_n} + \sigma_\nu \sqrt{n(\gamma_n + \log(1/\delta))}\right)\right\} \geq 1 - \delta . \quad (17)$$

Proof sketch. Considering Theorem 3 in Chowdhury and Gopalan (2017), the proof follows almost immediately from the assumptions above. The only detail to notice is that $\nu_t := y_t - g(\mathbf{x}_t) = \zeta_t + f(\tilde{\mathbf{x}}_t) - \mathbb{E}[f(\tilde{\mathbf{x}}_t)|\mathbf{x}_t] = \zeta_t + \Delta f_{P^E_{\mathbf{x}_t}}$, which is a σ_ν-sub-Gaussian random variable for $\sigma_\nu^2 = \sigma_\zeta^2 + \sigma_E^2$. ☐

The result above states that, as long as σ_ν is large enough to accommodate for the additional variance in the observations due to noisy-inputs, IGP-UCB maintains bounded regret. Theoretical results bounding the growth of γ_n are available in the literature. For the squared-exponential kernel on $\mathbb{R}^d$, for example, $\gamma_n \in \mathcal{O}((\log n)^{d+1})$ (Srinivas et al., 2010, Thr. 5), so that IGP-UCB obtains asymptotically vanishing uncertain-inputs regret in this case. However, it is possible that the resulting σ_ν makes β_t impractically large, leading to excessive exploration. The following result addresses these issues.

Proposition 3. *Let $k : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$ be an at least twice-differentiable positive-definite kernel with finite $\ell_k^2 \geq \sup_{\mathbf{x} \in \mathbb{R}^d} \sup_{i \in [d]} \frac{\partial^2 k(\mathbf{x}, \mathbf{x}')}{\partial x_i \partial x'_i}\big|_{\mathbf{x}=\mathbf{x}'}$. Then, for $P \in \mathcal{P}$ and $\tilde{\mathbf{x}} \sim P$, we have that $\Delta f_P := f(\tilde{\mathbf{x}}) - \mathbb{E}_P[f(\tilde{\mathbf{x}})]$ is σ_F-sub-Gaussian with:*

1. *$\sigma_F = \|f\|_k \ell_k \mathrm{tr}(\boldsymbol{\Sigma})^{1/2}$, if P is Gaussian with covariance matrix $\boldsymbol{\Sigma}$;*
2. *$\sigma_F = \frac{1}{2}\|f\|_k \ell_k \sqrt{\sum_{i=1}^d \sigma_i^2}$, if P has compact support, with $|\tilde{x}_i - \hat{x}_i| \leq \frac{1}{2}\sigma_i$ for each coordinate i, where $\hat{\mathbf{x}} = \mathbb{E}_P[\tilde{\mathbf{x}}]$.*

Proof sketch. These results are derived from concentration inequalities for Lipschitz-continuous functions of Gaussian or bounded random variables. For the given kernel, any $f \in \mathcal{H}_k$ is $\|f\|_k \ell_k$-Lipschitz continuous (Steinwart and Christmann, 2008). ☐

Proposition 3 says that the second condition in Theorem 2 is met if the execution noise is uniformly bounded or Gaussian. What remains is to verify whether the first assumption in Theorem 2 can be met.

When working with kernel embeddings of conditional distributions, the assumption that $\mathbf{x} \mapsto \mathbb{E}[f(\tilde{\mathbf{x}})|\mathbf{x}]$ is an element of $\mathcal{H}_k$ is known to be met when the domain $\mathcal{X}$ is discrete, while not necessarily holding for continuous domains (Muandet et al., 2016). As most interesting problems involving uncertain inputs have continuous domains, the following result presents a case where Theorem 2's first assumption holds.

Proposition 4. *Let $\mathbf{x} \mapsto P_\mathbf{x}$ be a mapping such that, for any $\mathbf{x} \in \mathcal{S} \subset \mathcal{X}$, $\tilde{\mathbf{x}} \sim P_\mathbf{x} \in \mathcal{P}$ is decomposable as $\tilde{\mathbf{x}} = \mathbf{x} + \boldsymbol{\epsilon}$, where $\boldsymbol{\epsilon}$ is independent and identically distributed, i.e. $\boldsymbol{\epsilon} \sim P_{\boldsymbol{\epsilon}} \in \mathcal{P}$. Assume that k is translation invariant. Then we have that, for any $f \in \mathcal{H}_k$, the mapping $\mathbf{x} \mapsto \mathbb{E}_{P_\mathbf{x}}[f]$ defines a function $g \in \mathcal{H}_k(\mathcal{S})$, and $\|g\|_k \leq \|f\|_k$.*

Proof sketch. The proof follows by interpreting $\boldsymbol{\epsilon}$ as a random translation on f's inputs, for any $f \in \mathcal{H}_k$. Since k is translation invariant, the norm of any $\boldsymbol{\epsilon}$-shifted function $f^{\boldsymbol{\epsilon}}$ is equivalent to the norm of the original f. Then picking g as the restriction of $\mathbb{E}_{P_{\boldsymbol{\epsilon}}}[f^{\boldsymbol{\epsilon}}] \in \mathcal{H}_k$ to $\mathcal{S} \subset \mathcal{X}$ leads to the conclusion. ☐

Proposition 4 implies that Theorem 2 is applicable whenever the execution noise is independent and identically distributed and k is translation-invariant, such as the squared exponential and other popular kernels. However, in cases where the execution noise distribution changes significantly from target to target, algorithms such as uGP-UCB can yield better results.

5.2 Bounding the regret of uGP-UCB

In this section, we analyse the case when uGP-UCB has no access to location estimates P_t^L and uses instead $\hat{P}_{\mathbf{x}_t}$ with the observations $\mathcal{D}_n = \{\hat{P}_{\mathbf{x}_t}, y_t\}_{t=1}^n$. We will firstly consider an ideal setting, where $\hat{P}_{\mathbf{x}} = P_{\mathbf{x}}^E$, $\forall \mathbf{x} \in \mathcal{S}$, and then a non-ideal scenario. Recall that the regret bounds presented so far depend on the maximum information gain γ_n. As an analogy, in the case of uGP-UCB, given any $\{P_t\}_{t=1}^n \subset \mathcal{P}$, we have:

$$I(\mathbf{y}_n; \hat{\mathbf{f}}_n | \{P_t\}_{t=1}^n) = \frac{1}{2} \log |\mathbf{I} + \lambda^{-1} \hat{\mathbf{K}}_n| \ , \tag{18}$$

where $[\hat{\mathbf{K}}_n]_{ij} = \hat{k}(P_i, P_j),\ i, j \in \{1, \dots, n\}$. Let's assume an arbitrary set $\mathcal{P}_s \subset \mathcal{P}$ containing either the query model or the estimated location distributions. As the set $\mathcal{P}_s$ is not necessarily compact, a maximum for $I(\mathbf{y}_n; \hat{\mathbf{f}}_n | \mathcal{R})$ may not correspond to a given set $\mathcal{R} \subset \mathcal{P}_s$. However, we can always define:

$$\hat{\gamma}_n(\mathcal{P}_s) := \sup_{\mathcal{R} \subset \mathcal{P}_s : |\mathcal{R}| = n} I(\mathbf{y}_n; \hat{\mathbf{f}}_n \mid \mathcal{R}) \ , \tag{19}$$

The results presented next will use $\hat{\gamma}_n^E := \hat{\gamma}_n(\mathcal{P}_E)$, where $\mathcal{P}_E \subset \mathcal{P}$ is the image of $\mathcal{S}$ under the mapping $\mathbf{x} \mapsto P_{\mathbf{x}}^E$. Considering these definitions, the following bounds the uncertain-inputs regret of uGP-UCB.

Theorem 5 (uGP-UCB regret)**.** *Let $\delta \in (0, 1)$, $f \in \mathcal{H}_k$, and $b \geq \|f\|_k$. Consider ζ_t as σ_ζ-sub-Gaussian noise. Assume that both k and $P_{\mathbf{x}}^E$ satisfy the conditions for $\Delta f_{P_{\mathbf{x}}^E}$ to be σ_E-sub-Gaussian, for a given $\sigma_E > 0$, for all $t \geq 1$. Then, running uGP-UCB with:*

$$\beta_t = b + \sigma_\nu \sqrt{2(I(\mathbf{y}_{t-1}; \hat{\mathbf{f}}_{t-1} | \{P_{\mathbf{x}_i}^E\}_{i=1}^{t-1}) + 1 + \log(1/\delta))} \ , \tag{20}$$

where $\sigma_\nu := \sqrt{\sigma_E^2 + \sigma_\zeta^2}$, the uncertain-inputs cumulative regret satisfies:

$$\hat{R}_n \in \mathcal{O}\left(\sqrt{n \hat{\gamma}_n^E} \left(b + \sqrt{\hat{\gamma}_n^E + \log(1/\delta)} \right) \right) \tag{21}$$

with probability at least $1 - \delta$.

Proof sketch. This theorem applies the fact that $\hat{k}_E(\mathbf{x}, \mathbf{x}') := \hat{k}(P_{\mathbf{x}}^E, P_{\mathbf{x}'}^E)$, for $\mathbf{x}, \mathbf{x}' \in \mathcal{S}$, defines a positive-definite kernel on $\mathcal{S}$ (Steinwart and Christmann, 2008, Lem. 4.3). By Lemma 1, we have that exists a $\hat{f} \in \mathcal{H}_{\hat{k}}$ with $\hat{f}(P_{\mathbf{x}}^E) = \mathbb{E}[f(\tilde{\mathbf{x}})|\mathbf{x}]$, for $\tilde{\mathbf{x}}|\mathbf{x} \sim P_{\mathbf{x}}^E$. Then it follows that $g : \mathbf{x} \mapsto \hat{f}(P_{\mathbf{x}}^E)$ is in $\mathcal{H}_{\hat{k}_E}$. As $\hat{\gamma}_n^E$ is the maximum information gain of a model $\mathrm{GP}(0, \hat{k}_E)$, the rest follows similarly to Theorem 2. □

Theorem 5 states that uGP-UCB obtains similar bounds for the uncertain-inputs regret to those of IGP-UCB. However, notice that $\hat{\gamma}_n^E$, instead of γ_n, appears in Equation 21. The next result shows that $\hat{\gamma}_n^E \leq \gamma_n$, which means smaller regret bounds, in the i.i.d. execution noise case considered previously (Proposition 4).

Proposition 6. *Consider a compact set $\mathcal{S} \subset \mathcal{X}$, a distribution $P_{\boldsymbol{\epsilon}} \in \mathcal{P}$, with $\mathbb{E}_{P_{\boldsymbol{\epsilon}}}[\boldsymbol{\epsilon}] = 0$, and a set:*

$$\mathcal{P}_{\boldsymbol{\epsilon}} := \{ P \in \mathcal{P} \mid \tilde{\mathbf{x}} = \hat{\mathbf{x}} + \boldsymbol{\epsilon},\, \hat{\mathbf{x}} \in \mathcal{S}, \boldsymbol{\epsilon} \sim P_{\boldsymbol{\epsilon}},\, \tilde{\mathbf{x}} \sim P \} \ , \tag{22}$$

which is the set of location distributions with mean in $\mathcal{S}$ and affected by i.i.d. $P_{\boldsymbol{\epsilon}}$-noise. Assume that $k : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ is translation invariant, and let $\hat{k} : \mathcal{P} \times \mathcal{P} \to \mathbb{R}$ be defined according to Equation 9. Then we have that:

$$\forall n \geq 1, \quad \hat{\gamma}_n(\mathcal{P}_{\boldsymbol{\epsilon}}) \leq \gamma_n \ , \tag{23}$$

where $\hat{\gamma}_n$ is defined by Equation 19, and γ_n is the maximum information gain for $\mathrm{GP}(0, k)$.

Proof sketch. One can prove that $\mathbf{K}_n - \hat{\mathbf{K}}_n$ is positive definite for a $\mathbf{K}_n$ built with $\{\hat{\mathbf{x}}_t\}_{t=1}^n \subset \mathcal{S}$. The information gain is a function of the determinant of these matrices, so that the inequality above follows. □

The result above indicates that the uncertain-inputs information gain shrinks as the input noise variance grows. While that might indicate that the optimisation problem becomes easier, if one recalls Equation 2, the constant ρ_E grows, making the problem harder.

What remains to verify is the effect of the approximation error between the model $\hat{P}_{\mathbf{x}}$ and the actual $P_{\mathbf{x}}^E$. To minimise $\hat{r}_t$, using uGP-UCB with a model $\hat{P}_{\mathbf{x}} \approx P_{\mathbf{x}}^E$ is worth if the approximation error $\hat{\rho}_t := \max_{\mathbf{x} \in \mathcal{S}} \left| \mathbb{E}_{P_{\mathbf{x}}^E}[f] - \mathbb{E}_{\hat{P}_{\mathbf{x}}}[f] \right|$ is small. Ideally $\hat{P}_{\mathbf{x}}$ should be an adaptive model $\hat{P}_{\mathbf{x}}^t$ that can be learnt from past data in $\mathcal{D}_{t-1}$ so that $\hat{\rho}_t \to 0$ as $t \to \infty$. However, considering execution noise as marginally i.i.d. and Gaussian has been a popular approach when dealing with problems involving uncertain inputs (Mchutchon and Rasmussen, 2011; Nogueira et al., 2016). In this case, we provide an upper bound on $\hat{\rho}_t$.

Proposition 7. *Let $\mathcal{X} = \mathbb{R}^d$, $f \in \mathcal{H}_k$ and $\|f\|_k \leq b$. Assume that, for any $\mathbf{x} \in \mathcal{S} \subset \mathcal{X}$, the query distribution $P_{\mathbf{x}}^E$ is Gaussian with mean $\mathbf{x}$ and positive-definite covariance $\boldsymbol{\Sigma}^E$. Then, using a Gaussian model $\hat{P}_{\mathbf{x}}$ with same mean and a given constant positive-definite covariance matrix $\hat{\boldsymbol{\Sigma}}$, we have that for any $\mathbf{x} \in \mathcal{S}$:*

$$\left| \mathbb{E}_{P_{\mathbf{x}}^E}[f] - \mathbb{E}_{\hat{P}_{\mathbf{x}}}[f] \right| \leq \frac{b}{2} \sqrt{\mathrm{tr}(\hat{\boldsymbol{\Sigma}}^{-1} \boldsymbol{\Sigma}^E) - d + \log \frac{|\hat{\boldsymbol{\Sigma}}|}{|\boldsymbol{\Sigma}^E|}} \ .$$

Proof sketch. This result follows by applying Pinsker's inequality (Boucheron et al., 2013) to $\hat{P}_{\mathbf{x}}$ and $P_{\mathbf{x}}^E$. □

6 Experiments

In this section, we present experimental results obtained in simulation with the proposed uGP-UCB algorithm comparing it against other Bayesian optimisation methods: IGP-UCB, with adapted noise model (as in Theorem 2), and the unscented expected improvement (UEI) heuristic (Nogueira et al., 2016), which applies the unscented transform to the expected improvement over a conventional GP model. Our aim in this section is to evaluate the performance of these methods in optimisation problems where both the sampling of the objective function and the location at which the sample is taken are significantly noisy.

6.1 Objective functions in the same RKHS

In this experiment, for each trial a different objective function $f \in \mathcal{H}_k$ was generated, where k is the squared-exponential kernel with length-scale set to 0.1. The search space was set to the unit square $\mathcal{S} = [0,1]^2 \subset \mathbb{R}^2$. Each $f = \sum_{i=1}^{m} \alpha_i k(\cdot, \mathbf{x}_i)$ was generated by uniformly sampling $\alpha_i \in [-1, 1]$ and support points $\mathbf{x}_i \in \mathcal{S}$, for $i \in \{1, \dots, m\}$, with $m = 30$. Observation noise was set as $\zeta_t \sim N(0, \sigma_\zeta^2)$ with $\sigma_\zeta = 0.1$.

As parameters to verify the theoretical results for the UCB algorithms, we set $\delta = 0.4$, and computed $b = \|f\|_k$ directly. The querying execution noise in $P_{\mathbf{x}}^E$ was i.i.d. sampled from $N(\mathbf{0}, \sigma_{\mathbf{x}}^2 \mathbf{I})$ with $\sigma_{\mathbf{x}} = 0.1$. The output noise parameters for the GP model were computed according to Proposition 3, with each method assuming execution noise coming from $N(\mathbf{0}, \hat{\sigma}_{\mathbf{x}}^2 \mathbf{I})$. To verify robustness to noise-misspecification, we tested $\hat{\sigma}_{\mathbf{x}}$ set according to different ratios with respect to the true $\sigma_{\mathbf{x}}$. Noise on the localisation estimates P_t^L was set at half the standard deviation of the true execution noise. We directly computed the current information gain $I(\mathbf{y}_t; \hat{\mathbf{f}}_t \mid \{P_i^L\}_{i=1}^t)$ to set β_t. For all methods, the GP covariance function was set as the RKHS kernel k.

Results: Figure 2 presents performance results, in terms of mean uncertain-inputs regret, i.e. $\hat{r}_t^{\text{avg}} = \frac{1}{t} \sum_{i=1}^{t} \hat{r}_i$. This performance metric is an upper bound on the simple regret, since $\min_{i \le t} \hat{r}_i \le \hat{r}_t^{\text{avg}}$, and allows verifying how close each method gets to the global optimum within t iterations. As the plots show, when the execution noise model is correct, with $\hat{\sigma}_{\mathbf{x}} = \sigma_{\mathbf{x}}$, uGP-UCB is able to outperform both IGP-UCB and UEI, while every method's performance degrades under mismatch in the execution noise assumption. A larger than needed execution noise variance leads to a large β_t for the UCB methods, promoting exploration. Querying with a very noisy model $\hat{P}_{\mathbf{x}}$ also excessively smoothes the GP prior and the acquisition function for uGP-UCB and UEI, respectively. Consequently, each

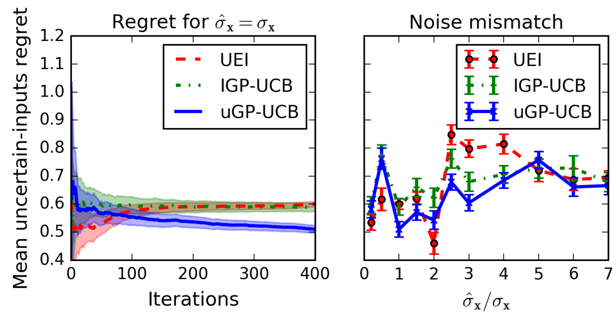

Figure 2: Optimisation of functions in the same RKHS. On the left, the UCB confidence-bound parameter β_t was set according to the theoretical results. The plot on the right shows the effect of execution noise model mismatch on each method's regret after running for a total of 400 iterations. Results were averaged over 10 trials, and the shaded areas and error bars correspond to one standard deviation.

method's model on f tends to a flat function, and none of them is able to make significant improvements after large mismatches, such as $\hat{\sigma}_{\mathbf{x}} \ge 5\sigma_{\mathbf{x}}$, as Figure 2 shows. Despite the loss of performance, uGP-UCB remains as a general lower bound in terms of regret, showing that the proposed method is relatively robust to the effects of mismatch in the execution noise model.

In practice, the convergence rate in Figure 2 can be improved by setting the UCB parameter β_t at a fixed low value. As the $\mathcal{O}$ notation indicates, cumulative regret bounds are valid only up to a constant factor. Their main focus is on guaranteeing asymptotic convergence, i.e. $\lim_{n \to \infty} \frac{\hat{R}_n}{n} = 0$, as most theoretical results in the UCB literature (Srinivas et al., 2010; Chowdhury and Gopalan, 2017). To achieve no regret, the value of the UCB parameter β_t monotonically increases over iterations, ensuring that the search space is fully explored. The drawback, however, is that excessive exploration decreases performance in the short term. In the next section, we present results where β_t is fixed.

6.2 Objective function in different RKHS

To verify uGP-UCB's performance under incorrect kernel assumptions, the next experiment performed tests with an objective function in a space not matching the GP squared-exponential kernel's RKHS. In particular, we chose the 4-dimensional Michalewicz function, which is a classic benchmark function for global optimisation algorithms (Vanaret et al., 2014), over the domain $\mathcal{S} = [0, \pi]^4$.

Figure 3 presents performance results for fixed $\beta_t = 3$ and a comparison of each algorithm's sensitivity to the

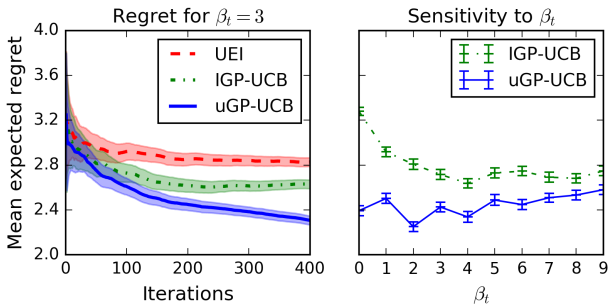

Figure 3: Optimisation of the Michalewicz function. The plot on the left presents the mean expected regret observed for each algorithm with $\beta_t = 3$ for UCB methods. On the right, we see how different settings for β_t affect each UCB method's mean expected regret after 300 iterations. Results were averaged over 10 (left) and 5 (right) trials with shaded areas and error bars corresponding to two standard deviations.

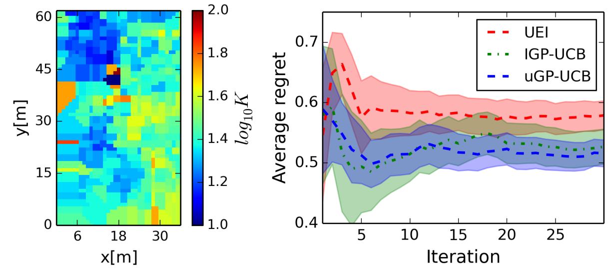

(a) Broom's barn data (b) Robotics problem

Figure 4: Robotics exploration experiment: (a) presents the Broom's barn data as distributed over the search space; and (b) shows the performance of each BO approach, averaged over 4 runs.

choice of this UCB parameter. Input noise was set to $\sigma_{\mathbf{x}} = 0.1$, and output noise was absent. With β_t fixed across iterations, IGP-UCB served as a BO baseline not accounting for input noise. As the results show, uGP-UCB is able to outperform both baselines. In addition, uGP-UCB shows consistently better performance than IGP-UCB across varying settings for β_t. These results demonstrate that the uGP-UCB algorithm is able to perform well in situations where its modelling assumptions are not exactly met, such as in scenarios involving physical systems.

6.3 Robotic exploration problem

This section presents results obtained in a simulated robotic exploration problem. In this experiment, a robot is set to explore an environmental process. The underlying process is based on the Broom's Barn dataset[2], consisting of the log-concentration of potassium in the soil of an experimental agricultural area. The robot is allowed to perform up to 30 measurements on different locations. Each BO method sequentially selects the locations where the robot should make a measurement in the usual online decision making process, based on the observations it gets. To simulate the robot, an ATRV platform, we used the OpenRobots' Morse simulator[3]. In this scenario, execution noise is not following a stationary distribution due to the dynamic constraints of the robot, imperfections in motion control, etc. We applied Gaussian noise to the pose information given by the simulator and used pure-pursuit path-following control to guide the robot to the target locations. Location estimates were provided by an extended Kalman filter (Thrun et al., 2006). Hyper-parameters for each GP were learnt online via log-marginal likelihood maximisation. The query noise model for uGP-UCB was set with $\hat{\sigma}^2_{\mathbf{x}} = 2$. We set β_t at a fixed value, again with $\beta_t = 3$. Figure 4b presents the performance of each algorithm in terms of regret. The plots show that uGP-UCB is able to outperform UEI, while performing still better than IGP-UCB in the long run, and with less variance in the outcomes. This result confirms that it is possible to obtain better performance in practical BO problems by taking advantage of distribution estimates and by directly considering execution uncertainty.

7 Conclusion

In this paper we proposed a novel method to optimise functions where both the sampling of the function as well as the location at which the function is sampled are stochastic. We provided theoretical guarantees for BO algorithms in noisy-inputs settings. In experiments we demonstrated that the proposed uGP-UCB shows competitive performance when compared to other BO approaches to input noise. Our method can be applied to problems where input variates or an agent's state is only partially observable, such as robotics, policy search, stochastic simulations, and others. For future work, it is worth investigating online-learning techniques for the approximate querying distribution $\hat{P}_{\mathbf{x}}$ and other bounds for the uncertain-inputs GP information gain.

Acknowledgements

We would like to thank the reviewers and Dr. Vitor Guizilini for helpful discussions and the funding agencies CAPES, Brazil, and Data61/CSIRO, Australia.

[2]Available at `http://www.kriging.com/datasets/`

[3]Morse: `https://www.openrobots.org/morse`

References

H. Bauer. *Probability theory and elements of measure theory*. Probability and mathematical statistics. Academic Press, 1981.

Justin J. Beland and Prasanth B. Nair. Bayesian Optimization Under Uncertainty. In *31st Conference on Neural Information Processing Systems (NIPS 2017) Workshop on Bayesian optimization (BayesOpt 2017)*, Long Beach, CA, 2017.

Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.

Adam D. Bull. Convergence Rates of Efficient Global Optimization Algorithms. *Journal of Machine Learning Research (JMLR)*, 12:2879–2904, 2011.

Sayak Ray Chowdhury and Aditya Gopalan. On Kernelized Multi-armed Bandits. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, Sydney, Australia, 2017.

Patrick Dallaire, Camille Besse, and Brahim Chaib-Draa. An approximate inference with Gaussian process to latent functions from uncertain data. *Neurocomputing*, 74:1945–1955, 2011.

Agathe Girard. *Approximate methods for propagation of uncertainty with Gaussian process models*. Ph. d, University of Glasgow, 2004.

Henry Lam. Advanced tutorial: input uncertainty and robust analysis in stochastic simulation. In *Proceedings of the 2016 Winter Simulation Conference*, pages 178–192, 2016.

Chun Kai Ling, Kian Hsiang Low, and Patrick Jaillet. Gaussian Process Planning with Lipschitz Continuous Reward Functions: Towards Unifying Bayesian Optimization, Active Learning, and Beyond. In *AAAI Conference on Artificial Intelligence*, 2016.

Roman Marchant and Fabio Ramos. Bayesian Optimisation for Intelligent Environmental Monitoring. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, October 2012.

Roman Marchant and Fabio Ramos. Bayesian Optimisation for Informative Continuous Path Planning. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 6136–6143, 2014.

Andrew Mchutchon and Carl E. Rasmussen. Gaussian Process Training with Input Noise. In *Advances in Neural Information Processing Systems*, pages 1341–1349, 2011.

Krikamol Muandet, Kenji Fukumizu, Francesco Dinuzzo, and Bernhard Schölkopf. Learning from Distributions via Support Measure Machines. In *Proceeding of the 26th Annual Conference on Neural Information Processing Systems (NIPS 2012)*, 2012.

Krikamol Muandet, Kenji Fukumizu, Bharath Sriperumbudur, and Bernhard Schölkopf. Kernel Mean Embedding of Distributions: A Review and Beyond. *arXiv*, 2016.

José Nogueira, Ruben Martinez-Cantin, Alexandre Bernardino, and Lorenzo Jamone. Unscented Bayesian Optimization for Safe Robot Grasping. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Daejeon, Korea, 2016.

Rafael Oliveira, Lionel Ott, Vitor Guizlini, and Fabio Ramos. Bayesian Optimisation for Safe Navigation under Localisation Uncertainty. In *International Symposium on Robotics Research (ISRR) (to appear)*, Puerto Varas, Chile, 2017.

Michael Pearce and Juergen Branke. Bayesian simulation optimization with input uncertainty. In W. K. V. Chan, A. D'Ambrogio, G. Zacharewicz, N. Mustafee, G. Wainer, and E. Page, editors, *Proceedings of the 2017 Winter Simulation Conference*, pages 2268–2278. IEEE, 2017.

Carl E. Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, Cambridge, MA, 2006.

Bernhard Schölkopf and Alexander J. Smola. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT Press, Cambridge, Mass, 2002.

Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P. Adams, and Nando De Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016.

Jasper Snoek, Hugo Larochelle, and Ryan P. Adams. Practical bayesian optimization of machine learning algorithms. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 2951–2959. Curran Associates, Inc., 2012.

Niranjan Srinivas, Andreas Krause, Sham M. Kakade, and Matthias Seeger. Gaussian Process Optimization in the Bandit Setting: No Regret and Experimental Design. In *Proceedings of the 27th International Conference on Machine Learning (ICML 2010)*, pages 1015–1022, 2010.

Bharath K. Sriperumbudur, Kenji Fukumizu, and Gert R. G. Lanckriet. Universality, Characteristic Kernels and RKHS Embedding of Measures. *Journal of Machine Learning Research (JMLR)*, 12:2389–2410, 2011.

Ingo Steinwart and Andreas Christmann. *Support Vector Machines*, chapter 4, pages 110–163. Springer New York, New York, NY, 2008.

Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic Robotics*. The MIT Press, Cambridge, MA, 2006.

Saul Toscano-Palmerin and Peter I. Frazier. Bayesian Optimization with Expensive Integrands. *arXiv*, 2018.

Charlie Vanaret, Jean-Baptiste Gotteland, Nicolas Durand, and Jean-Marc Alliot. Certified Global Minima for a Benchmark of Difficult Optimization Problems. Technical report, hal-00996713, 2014. Preprint.

Eric A. Wan and Rudolph van der Merwe. The Unscented Kalman Filter for Nonlinear Estimation. In *Adaptive Systems for Signal Processing, Communications, and Control Symposium (AS-SPCC)*, pages 153–158, 2000. ISBN 0780384822.

Aaron Wilson, Alan Fern, and Prasad Tadepalli. Using trajectory data to improve bayesian optimization for reinforcement learning. *Journal of Machine Learning Research*, 15:253–282, 2014.