

Leveraging Localisation Information into Bayesian Optimisation for Planning in Robotics

Rafael Oliveira*, Lionel Ott, Vitor Guizilini and Fabio Ramos

Abstract—Bayesian optimisation (BO) algorithms have been successfully applied to a wide range of problems where the objective function is expensive to evaluate and the observed values are corrupted by noise. In areas such as robotics, however, it is common to find problems where the querying of the objective function is noisy as well, for example due to localisation uncertainty. In these cases, the actual query point is not fully observable and its uncertainty should be considered in the model. In this paper we present a principled BO approach that takes input uncertainty into account. In particular, we propose a general upper-confidence bound (UCB) BO algorithm that is able to deal with problems where input locations are uncertain, both in terms of query execution and location estimation. For that, we employ a Gaussian process model that uses probability distributions as inputs. We provide theoretical results for this algorithm and for an adaptation to input noise for a recently proposed UCB method, analysing regret bounds. Finally, we evaluate their performance experimentally in simulations, comparing them to another BO approach to problems involving uncertain inputs.

I. INTRODUCTION

Bayesian optimisation (BO) [1] 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 [2], policy search [3], environmental monitoring [4], robotic grasping [5], etc. Although taking into account that we might have a noisy observation of the function’s output value, conventional BO approaches assume that the function has been sampled precisely at the specified query location within the given search space. While this is true for many applications of BO, there are certain problems, especially in areas of robotics and 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 certain 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* makes the robot unable to reach 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 some representation of a probability distribution $\tilde{\mathbf{x}}_t \sim \mathbb{P}_t^L$, which takes into account *localisation*




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

noise, as depicted in Figure 1. As a result of each query, we get a noisy observation of the environmental process $z_t = f(\tilde{\mathbf{x}}_t) + \nu_t$, where ν_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.

In this paper, we investigate optimisation problems where input noise affects both the query execution process and the estimation of the actual query location in the input space. In particular, we analyse the standard BO approach when running the improved Gaussian process upper-confidence bound (IGP-UCB) [6] algorithm under input noise, and we propose the *uncertain-inputs Gaussian process upper confidence bound (uGP-UCB)* algorithm, an extension of the original IGP-UCB that is equipped with a GP model for uncertain-inputs, in a similar framework to Oliveira et al. [7]. We present theoretical guarantees for uGP-UCB and provide empirical performance evaluations for the proposed algorithm, comparing it to the standard BO approach and the recently proposed unscented BO [5] in simulated problems involving input noise.

II. RELATED WORK

Bayesian optimisation [1] usually considers deterministic query locations within the objective function’s input space, as BO is typically applied to problems where a fixed number of parameters have to be optimised [2, 3]. Recently, however, BO frameworks to deal with input noise have been proposed in the literature. Nogueira et al. [5] presented a method to apply BO to problems where the execution of a query is uncertain,

The authors are with the School of Information Technologies, The University of Sydney, Australia. (*) Corresponding author: rdos6788@uni.sydney.edu.au. This work was supported by Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES), Brazil, and by Data61/CSIRO, Australia.

such as robotic grasping. The authors propose applying the unscented transform [8] to BO’s acquisition function to cope with execution noise, represented by a Gaussian distribution. This way, the algorithm is able to find optima that are robust to execution noise. Another approach has been proposed in Oliveira et al. [7], where the authors employ a Gaussian process (GP) model for uncertain inputs [9] as a prior, instead of conventional GP models for deterministic inputs. In that framework, besides being able to consider noise in the query execution, the BO algorithm is also able to take full advantage of input location estimates represented by probability distributions, as is often the case in robotics [8, 10]. This feature allows the algorithm to obtain better estimates of the underlying function being optimised. However, they consider a distance-penalised UCB as search heuristics, having a different objective, which includes reducing travelled distance. In this paper, we adopt a BO framework similar to Oliveira et al. [7], with a slight modification to the formulation of the uncertain-inputs GP model, which allows us to derive the first theoretical guarantees for BO algorithms in uncertain-inputs optimisation problems.

Different methods to incorporate input uncertainty into a GP model have been proposed in the machine learning literature. Girard et al. [11] presented a method to propagate input uncertainty through GP models in multi-step predictions, which has been later extended to incorporate observations with uncertain inputs [12, Section 3.6.1]. This work forms the basis for Dallaire et al. [9], which formulates an expected Gaussian process model, based on the joint inputs distribution. Besides these ones, GP models propagating input noise through a local linearisation of the latent function [13] and methods involving variational inference to infer the true latent inputs [14] have also been proposed. However, the local linearisation method [13] assumes that the input locations in the GP observations dataset are corrupted by noise coming from the same distribution, in contrast to Girard [12], which allows distributions to be arbitrary. On the other hand, methods based on variational inference [14] might not be suitable for BO due to the consequent computational overhead in optimising the variational parameters at each GP update, which would be done online in BO. In addition, none of the present methods for GP models with uncertain inputs provides theoretical guarantees on the approximation errors with respect to the true underlying function. In this paper, we shed new light over the GP models in Girard [12] and Dallaire et al. [9], connecting them to some of the theory of support measure machines [15] and kernel mean embeddings [16], which allows us to upper bound the approximation error for BO in the uncertain-inputs setting when running an extension of the GP-UCB algorithm [17, 6].

Problems like the one illustrated in Figure 1 can be present whenever the objective is a function of an agent’s unknown state, which is controllable through a continuous set of actions. Therefore, it is possible to relate the work in this paper to partially-observable Markov decision processes (POMDPs) as well, which have been previously connected to BO [18, 19]. POMDP problems are usually concerned with finding an

optimal policy to maximise an expected sum of rewards over some time horizon. In this paper, however, we are more interested in a general optimisation setup, analogous to an agent only concerned with maximising its immediate expected reward under the state’s distribution.

III. PRELIMINARIES

Bayesian optimisation (BO) is a strategy for global optimisation that incorporates some prior knowledge about the objective function by means of a statistical model, which is a GP in our case. Therefore, we begin by reviewing some concepts about GP regression followed by an introduction to the general BO strategy.

A. Gaussian process regression

Gaussian process (GP) regression [20] provides a non-parametric framework for Bayesian inference over function spaces. A zero-mean GP model with a positive-definite covariance function $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ places a distribution over the space of functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$. Under the GP assumption, a finite collection of T function values $\mathbf{f}_T = [f(\mathbf{x}_1), \dots, f(\mathbf{x}_T)]^\top$ at a given set of inputs $\mathcal{Q}_T = \{\mathbf{x}_i\}_{i=1}^T \subset \mathbb{R}^d$ follows a multivariate Gaussian distribution $\mathbf{f}_T | \mathcal{Q}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{K}_T)$, where $[\mathbf{K}_T]_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$. In this way, a GP model allows us to do inference over the value of f at an unobserved input \mathbf{x} , given noisy observations \mathbf{z}_T , where $z_t = f(\mathbf{x}_t) + \eta_t$ with $\eta_t \sim \mathcal{N}(0, \sigma_\eta^2)$, at the given inputs \mathcal{Q}_T . In this sense, the posterior distribution over f is just another Gaussian process with mean μ_T and covariance k_T :

$$\mu_T(\mathbf{x}) = \mathbf{k}_T(\mathbf{x})^\top (\mathbf{K}_T + \sigma_\eta^2 \mathbf{I})^{-1} \mathbf{z}_T \quad (1)$$

$$k_T(\mathbf{x}, \mathbf{x}') = k(\mathbf{x}, \mathbf{x}') - \mathbf{k}_T(\mathbf{x})^\top (\mathbf{K}_T + \sigma_\eta^2 \mathbf{I})^{-1} \mathbf{k}_T(\mathbf{x}') \quad (2)$$

$$\sigma_T^2(\mathbf{x}) = k_T(\mathbf{x}, \mathbf{x}) \quad , \quad (3)$$

where $\mathbf{k}_T(\mathbf{x}_*) = [k(\mathbf{x}_1, \mathbf{x}_*), \dots, k(\mathbf{x}_T, \mathbf{x}_*)]^\top$. In this paper, we assume that $k(\mathbf{x}, \mathbf{x}) = 1$ for any $\mathbf{x} \in \mathbb{R}^d$, as it’s usual in the GP-UCB literature [17].

B. Bayesian optimisation

Consider the problem of searching for the global optimum of a possibly non-convex function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ over a given compact search space $\mathcal{S} \subset \mathbb{R}^d$, such as determining:

$$\mathbf{x}^* = \operatorname{argmax}_{\mathbf{x} \in \mathcal{S}} f(\mathbf{x}). \quad (4)$$

Assume that f is only partially observable via noisy estimates $z = f(\mathbf{x}) + \eta$ with $\eta \sim \mathcal{N}(0, \sigma_\eta^2)$ and that we can only sample the function T times. Bayesian optimisation [1] assumes that f is a random variable itself and applies a probability distribution over it using a statistical model, which is usually a GP, to represent its belief about f .

To select points at which to query f , BO uses an acquisition function $h(\mathbf{x})$ as a guide that makes use of the information provided by the GP prior and previous observations to estimate a utility value for sampling f at a given \mathbf{x} . Each query point \mathbf{x}_t is then selected by maximising over h , as:

$$\mathbf{x}_t = \operatorname{argmax}_{\mathbf{x} \in \mathcal{S}} h(\mathbf{x}). \quad (5)$$

After collecting an observation z_t , BO updates the GP model with the pair $\{\mathbf{x}_t, z_t\}$ and starts the next iteration, with an improved belief about f . This BO loop runs until a stopping criteria is met, which is usually defined as reaching the maximum budget of queries to the objective function, T .

a) Upper-confidence bound criteria: A commonly used acquisition function in the BO literature is the upper-confidence bound (UCB):

$$h_{\text{UCB}}(\mathbf{x}) = \mu_{t-1}(\mathbf{x}) + \beta_t \sigma_{t-1}(\mathbf{x}), \quad (6)$$

where $\beta_t > 0$ is a parameter of the algorithm and μ_{t-1} and σ_{t-1} are given by the GP posterior with the observations collected up to iteration $t-1$. Under certain assumptions, it is possible to properly set β_t to allow $h_{\text{UCB}}(\mathbf{x})$ to behave as an upper bound for the value of $f(\mathbf{x})$, valid at every iteration and across the entire search space with high probability, which leads to sub-linear regret bounds [17, 6].

IV. BAYESIAN OPTIMISATION UNDER UNCERTAIN INPUTS

In this section we present the framework for Bayesian optimisation under uncertain inputs. To extend BO to the case where the input locations \mathbf{x} are uncertain, we first need to extend the surrogate model to incorporate uncertainty in the inputs. We propose to use Gaussian process regression with uncertain inputs as the surrogate model which we then incorporate into the BO framework to obtain a framework for optimisation with uncertain inputs. We start by explaining the GP with uncertain inputs model before going into details of the final BO algorithm.

A. Gaussian process priors with uncertain inputs

Consider that, due to query execution noise, when evaluating a function f at a desired input location $\mathbf{x} \in \mathbb{R}^d$, the actual input where the observation is collected at is $\tilde{\mathbf{x}}|\mathbf{x} \sim \mathbb{P}_{\mathbf{x}}^E$. In addition, consider that the value of $\tilde{\mathbf{x}}$ is not directly observable, however, after the query, we have access to a probability distribution estimating its location, given by $\tilde{\mathbf{x}} \sim \mathbb{P}_{\tilde{\mathbf{x}}}^L$. Assuming that $f \sim \mathcal{GP}(0, k)$, the distribution over $f(\tilde{\mathbf{x}})$, when we only know \mathbf{x} , is in general not Gaussian any longer due to the non-linear relationship between f and its inputs. As shown by Girard [12], we can still recover the mean and the covariance function of the resulting stochastic process as $\mathbb{E}_{\tilde{\mathbf{x}}_* \sim \mathbb{P}_*} [f(\tilde{\mathbf{x}}_*)] = 0$ and $\text{Cov}_{\mathbb{P}_i, \mathbb{P}_j} [f(\tilde{\mathbf{x}}_i), f(\tilde{\mathbf{x}}_j)] = \mathbb{E}_{\mathbb{P}_i, \mathbb{P}_j} [k(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j)]$, where it is assumed that $\tilde{\mathbf{x}}_i$ and $\tilde{\mathbf{x}}_j$ are independent. Here \mathbb{E} denotes *expected value* and *Cov* stands for *covariance*. Therefore, we can define an *expected* covariance function, given by:

$$\begin{aligned} \mathbb{E}_{\mathbb{P}_i, \mathbb{P}_j} [k(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j)] &= \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} k(\mathbf{x}, \mathbf{x}') d\mathbb{P}_i(\mathbf{x}) d\mathbb{P}_j(\mathbf{x}') \\ &= \hat{k}(\mathbb{P}_i, \mathbb{P}_j). \end{aligned} \quad (7)$$

Analytical solutions for Equation 7 are available for several commonly-applied covariance functions [12, 9], and approximations are also feasible [12].

Although the resulting stochastic process describing $f \sim \mathcal{GP}(0, k)$ under observations with uncertain inputs is no longer Gaussian [12], the noisy-inputs covariance \hat{k} also

represents a positive-definite kernel function from a subclass of kernels defined over kernel-mean embeddings of probability measures [15]. This fact allows us to connect this GP model to a few well-established theoretical results available for kernel embeddings methods [16] and to derive the first theoretical guarantees for a BO algorithm involving it.

Let \mathcal{P} denote the set containing all probability measures on a measurable space $(\mathcal{X}, \mathcal{A})$, where \mathcal{A} is a σ -algebra of subsets of $\mathcal{X} \subseteq \mathbb{R}^d$. Then a GP model with zero mean and a positive-definite covariance function $\hat{k} : \mathcal{P} \times \mathcal{P} \rightarrow \mathbb{R}$ places a prior over the space of functions $g : \mathcal{P} \rightarrow \mathbb{R}$. Given a set of observations $\mathcal{D} = \{\mathbb{P}_i, \bar{z}_i\}_{i=1}^T$, where $\bar{z}_i = g(\mathbb{P}_i) + \eta_i$ and $\mathbb{P}_i \in \mathcal{P}$, under the assumptions that $g \sim \mathcal{GP}(0, \hat{k})$ and $\eta_i \sim \mathcal{N}(0, \sigma_\eta^2)$, we can do inference about g in the usual GP way, i.e., for a given $\mathbb{P}_* \in \mathcal{P}$:

$$g|\mathcal{D}, \mathbb{P}_* \sim \mathcal{N}(m_T(\mathbb{P}_*), v_T(\mathbb{P}_*)),$$

where:

$$m_T(\mathbb{P}_*) = \hat{\mathbf{K}}_T(\mathbb{P}_*)^\top (\hat{\mathbf{K}}_T + \sigma_\eta^2 \mathbf{I})^{-1} \bar{\mathbf{z}}_T \quad (8)$$

$$\hat{k}_T(\mathbb{P}, \mathbb{P}') = \hat{k}(\mathbb{P}, \mathbb{P}') - \hat{\mathbf{K}}_T(\mathbb{P})^\top (\hat{\mathbf{K}}_T + \sigma_\eta^2 \mathbf{I})^{-1} \hat{\mathbf{K}}_T(\mathbb{P}') \quad (9)$$

$$v_T(\mathbb{P}_*) = \hat{k}_T(\mathbb{P}_*, \mathbb{P}_*) \quad (10)$$

with $\hat{\mathbf{K}}_T(\mathbb{P}_*) = [\hat{k}(\mathbb{P}_*, \mathbb{P}_1), \dots, \hat{k}(\mathbb{P}_*, \mathbb{P}_T)]^\top$ and $[\hat{\mathbf{K}}_T]_{ij} = \hat{k}(\mathbb{P}_i, \mathbb{P}_j)$. Here we use m and v for mean and variance, respectively, to distinguish them from the usual notation describing GP models with deterministic inputs, following the convention in Girard [12]. Despite the description in terms of probability measures, in practice, each distribution \mathbb{P}_i can be represented by its parameters or approximated by a collection of samples at each entry in the GP dataset.

B. uGP-UCB algorithm for Bayesian optimisation under uncertain inputs

In this section, we describe the *uncertain-inputs Gaussian process upper confidence bound* (uGP-UCB) method. We start by formally specifying the problem of Bayesian optimisation under uncertain-inputs. Then we present how we applied the uncertain-inputs GP model from the previous section to this BO framework. Lastly we present the corresponding uGP-UCB acquisition function.

We consider an optimisation problem where we are allowed to select target locations \mathbf{x}_t in a compact search space $\mathcal{S} \subset \mathbb{R}^d$ at which to query a function f to find its global optimum. However, the query execution itself is a stochastic process, leading the query to actually be made at some $\tilde{\mathbf{x}}|\mathbf{x}_t \sim \mathbb{P}_{\mathbf{x}_t}^E$. Hence, the goal is defined as finding:

$$\mathbf{x}^* = \underset{\mathbf{x} \in \mathcal{S}}{\text{argmax}} \mathbb{E}_{\tilde{\mathbf{x}} \sim \mathbb{P}_{\mathbf{x}}^E} [f(\tilde{\mathbf{x}})], \quad (11)$$

which defines the target location that maximises the expected value of the function f under the querying process noise. This execution noise could be the outcome of a sequence of actions taken by the agent to reach the target location \mathbf{x}_t , which would then result in dependence on the starting point \mathbf{x}_{t-1} and the actions themselves. However, here we assume that

\mathbb{P}_x^E is a stationary distribution marginalising over all possible starting points and actions taken to reach a given target location \mathbf{x} , i.e. an average case scenario. We also assume that, *after* each query, we are informed with a distribution representing an independent estimate for the actual query location $\tilde{\mathbf{x}}_t \sim \mathbb{P}_t^L$ and observations $z_t = f(\tilde{\mathbf{x}}_t) + \nu_t$, where ν_t is σ_ν -sub-Gaussian noise [6]. Notice that this assumption on the noise is more general than usual BO assumptions, since it is not only met by Gaussians, but also by bounded random variables and others [21].

1) *Querying with a distribution*:: We use the uncertain-inputs GP model, described in Section IV-A, as the model for the expected value of the objective function f under arbitrary inputs distributions. Besides taking into account uncertain-inputs in the GP dataset $\{\mathbb{P}_t^L, z_t\}_t$, which are affected by localisation noise, we can query the GP prior during the acquisition function optimisation with a distribution $\hat{\mathbb{P}}_x$, instead of a deterministic point \mathbf{x} , to take execution noise into account. In general, we do not know the true distribution of the noise affecting queries. Then $\hat{\mathbb{P}}_x$ encodes our assumptions about the distribution of the actual query location $\tilde{\mathbf{x}}$ given a target \mathbf{x}_t , i.e. $\hat{\mathbb{P}}_x \approx \mathbb{P}_x^E$. For instance, we could assume that the execution noise follows a zero-mean Gaussian measure, so that $\tilde{\mathbf{x}}|\mathbf{x} \sim \mathcal{N}(\tilde{\mathbf{x}}; \mathbf{x}, \Sigma_x)$ for a given covariance matrix Σ_x . Notice that $\hat{\mathbb{P}}_x$ is only applied to query the acquisition function. Determining its true distribution can be a challenging task, beyond the scope of this paper. However, as our theoretical results present, we can still guarantee convergence even under an imperfect approximation for \mathbb{P}_x^E .

2) *Upper-confidence bound*:: With $\hat{\mathbb{P}}_x$ specified, we can then estimate the mean and the variance of $\hat{f}(\mathbb{P}_x^E) = \mathbb{E}_{\tilde{\mathbf{x}} \sim \mathbb{P}_x^E}[f(\tilde{\mathbf{x}})] \approx \mathbb{E}_{\tilde{\mathbf{x}} \sim \hat{\mathbb{P}}_x}[f(\tilde{\mathbf{x}})]$ under the GP assumptions using equations 8 and 10. We can then define an upper-confidence bound acquisition function over the expected outcome of observing f at a location $\tilde{\mathbf{x}} \sim \hat{\mathbb{P}}_x$ as follows:

$$h(\hat{\mathbb{P}}_x|\mathcal{D}_{t-1}) = m_{t-1}(\hat{\mathbb{P}}_x) + \beta_t v_{t-1}^{1/2}(\hat{\mathbb{P}}_x), \quad (12)$$

where β_t is a parameter that regulates the size of the confidence bound. At each iteration, BO selects the target location $\tilde{\mathbf{x}}_t$ that maximises $h(\hat{\mathbb{P}}_{\tilde{\mathbf{x}}_t}|\mathcal{D}_{t-1})$. By doing so, we select points that maximise an upper bound on the expected value of f . After each query, uGP-UCB updates the GP model with a location estimate \mathbb{P}_t^L , which is independently provided, and an observation z_t . Since the upper-confidence bound in Equation 12 is defined over a GP model for uncertain inputs, we name this Bayesian optimisation algorithm as *uncertain Gaussian process upper confidence bound* (uGP-UCB).

V. THEORETICAL ANALYSIS

When designing an optimisation algorithm we are interested in finding the location of the global optimum \mathbf{x}_* of a function f as efficiently as possible. The measure of how close the algorithm got to the global optimum at round t is usually quantified by the instant regret $r_t := f(\mathbf{x}_*) - f(\mathbf{x}_t)$. In our case we need to take the query uncertainty into account, and as such we need to consider that $f(\mathbf{x}^*)$ might be a very unlikely outcome and possibly unsafe in certain situations [5].

Therefore, we are interested in the *expected regret*, which we define as:

$$\bar{r}_t := \max_{\mathbf{x} \in \mathcal{S}} \mathbb{E}_{\mathbb{P}_x^E}[f] - \mathbb{E}_{\mathbb{P}_{\tilde{\mathbf{x}}_t}^E}[f]. \quad (13)$$

A desirable property of an optimisation algorithm is to have no regret, i.e. $\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \bar{r}_t = 0$. Consequently, if we bound the *cumulative expected regret* of the algorithm, $\bar{R}_T := \sum_{t=1}^T \bar{r}_t$, we can quantify how close it is able to get to the global optimum after T rounds, since $\bar{r}_t \leq \bar{R}_T/T$.

Srinivas et al. [17] provided the theoretical foundations to derive bounds for the cumulative regret of GP-UCB algorithms and several works followed after that. Recently, Chowdhury and Gopalan [6] presented a version of the algorithm that is able to converge under more general noise assumptions, deriving regret bounds for settings involving additive sub-Gaussian noise on the observed outcomes. In this section, we extend these results to the case of uGP-UCB and standard GP-UCB in the noisy-inputs setting. Some of the main issues in this case is that assumptions involving deterministic inputs or an Euclidean input space for the GP model do not hold. In addition, the true query distribution \mathbb{P}_x^E is unknown.

A. Uncertain-Inputs GP Analysis

We base our analysis for the uncertain-inputs GP model on the theory of support measure machines [15] and kernel mean embeddings [16]. Assume that $f \in \mathcal{H}_k$, i.e. the *reproducing kernel Hilbert space* (RKHS) [22] associated with the positive-definite kernel k , which implies that f is *not* a sample from $\mathcal{GP}(0, k)$ [see 23, Chapter 1]. In this case, we can define the mean map:

$$\mu : \mathcal{P} \rightarrow \mathcal{H}_k, \mathbb{P} \rightarrow \int_{\tilde{\mathbf{x}} \in \mathcal{X}} k(\cdot, \tilde{\mathbf{x}}) d\mathbb{P}(\tilde{\mathbf{x}}). \quad (14)$$

Assuming that $k(\cdot, \mathbf{x})$ is bounded for any $\mathbf{x} \in \mathcal{X}$,

$$\mathbb{E}_{\tilde{\mathbf{x}} \sim \mathbb{P}}[f(\tilde{\mathbf{x}})] = \langle \mu_{\mathbb{P}}, f \rangle_k, \quad \forall f \in \mathcal{H}_k, \quad (15)$$

where $\langle \cdot, \cdot \rangle_k$ denotes the usual inner product in the RKHS \mathcal{H}_k , and we use a notation shortcut $\mu_{\mathbb{P}} = \mu(\mathbb{P})$. In this sense, we have that:

$$\hat{k}(\mathbb{P}_i, \mathbb{P}_j) = \langle \mu_{\mathbb{P}_i}, \mu_{\mathbb{P}_j} \rangle_k \quad (16)$$

is a valid positive definite kernel [15], and we can also define $\hat{k}(\mathbf{x}, \mathbb{P}) = \langle k(\cdot, \mathbf{x}), \mu_{\mathbb{P}} \rangle_k$ for a given (deterministic) $\mathbf{x} \in \mathcal{X}$. We base our analysis on an uncertain-inputs GP model defined as in Equation 16.

Notice that this model that we choose is different from the *expected GP* model in Dallaire et al. [9]. The entries on the diagonal of the covariance matrix in the latter are $\mathbb{E}_{\tilde{\mathbf{x}} \sim \mathbb{P}_i}[k(\tilde{\mathbf{x}}, \tilde{\mathbf{x}})] \neq \langle \mu_{\mathbb{P}_i}, \mu_{\mathbb{P}_i} \rangle_k$. Since Dallaire et al. view their GP covariance function as the *expected covariance* under the joint distribution of the two locations, the diagonal entries of the expected covariance matrix are derived from the same random input location, becoming a single expectation.

B. Regret Bounds

We start by establishing sufficient conditions to make our optimisation problem to be analogous to the setup in Chowdhury and Gopalan [6]. As we have only noisy (uncertain) point estimates of $f(\tilde{\mathbf{x}}_t) + \nu_t$ available, we need the equivalent observation noise, i.e. $\eta_t = z_t - \mathbb{E}_{\mathbb{P}_t^L}[f]$, to be sub-Gaussian so that we can apply an equivalent form of Theorem 1 in Chowdhury and Gopalan [6] to our case. The following lemma states a few general cases where that condition can be satisfied. For proofs, please, refer to the appendix.

Lemma 1. *Let $f \in \mathcal{H}_k$ and $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ be an at least twice-differentiable positive-definite kernel with finite $L_k^2 \geq \sup_{\mathbf{x} \in \mathbb{R}^d} \sup_{i \in [d]} \frac{\partial^2 k(\mathbf{x}, \mathbf{x}')}{\partial x_i \partial x'_i} \Big|_{\mathbf{x}=\mathbf{x}'}$. Then, for $\mathbb{P} \in \mathcal{P}$ and $\tilde{\mathbf{x}} \sim \mathbb{P}$, we have that $\Delta f_{\mathbb{P}} := f(\tilde{\mathbf{x}}) - \mathbb{E}_{\mathbb{P}}[f(\tilde{\mathbf{x}})]$ is a σ_F -sub-Gaussian random variable in the following cases:*

- 1) \mathbb{P} is Gaussian, which implies $\sigma_F = \|f\|_k L_k \text{tr}(\Sigma)^{1/2}$, where Σ is \mathbb{P} 's covariance matrix;
- 2) \mathbb{P} has compact support, with $|\tilde{x}_i - \hat{x}_i| \leq \frac{1}{2}\sigma_i$ for each coordinate i , where $\hat{\mathbf{x}} = \mathbb{E}_{\mathbb{P}}[\tilde{\mathbf{x}}]$, which implies $\sigma_F = \frac{1}{2}\|f\|_k L_k \sqrt{\sum_{i=1}^d \sigma_i^2}$.

Notice that, in the case of conventional deterministic-inputs BO, the assumptions in the lemma above can be met if the execution-noise is uniformly bounded or Gaussian. Therefore, under certain assumptions, it is possible to guarantee asymptotic convergence for BO methods that can tolerate sub-Gaussian observation noise, such as IGP-UCB [6], in the noisy-inputs setting. In the theorem below we formalise that for the IGP-UCB case.

Theorem 1. *Let $f \in \mathcal{H}_k$, where k is a twice-differentiable and translation-invariant kernel. Assume that, for all $\mathbf{x} \in \mathcal{S}$, $\mathbb{P}_{\mathbf{x}}^E$ satisfies the conditions for $\Delta f_{\mathbb{P}_{\mathbf{x}}^E}$ to be σ_F -sub-Gaussian, for a given $\sigma_F > 0$. In addition, assume that, $\forall t \geq 1$, $\epsilon_t := \tilde{\mathbf{x}}_t - \mathbf{x}_t$ is independent and identically distributed when $\tilde{\mathbf{x}}_t \sim \mathbb{P}_{\mathbf{x}_t}^E$. Let $\delta \in (0, 1)$, and $B \geq \|f\|_k$, and ν_t be σ_ν -sub-Gaussian noise. Then, running IGP-UCB [6] with $\beta_t = B + \sqrt{2(\sigma_F^2 + \sigma_\nu^2)(\hat{\gamma}_{t-1} + 1 + \ln(1/\delta))}$, we obtain cumulative expected regret bounded by $\mathcal{O}(\sqrt{T}(B\sqrt{\gamma_T} + \gamma_T))$ with probability at least $1 - \delta$.*

The result above basically states that as long as the GP output noise model has a large enough variance to accommodate for the additional noise in the outputs observations due to input noise, IGP-UCB should be able to converge. However, it is possible that the required output noise variance is impractically large to compensate for all the execution noise distributions for every target $\mathbf{x} \in \mathcal{S}$.

Now consider cases where the query execution error is not guaranteed to be Gaussian, uniformly-bounded, target-location-independent, or specially, cases where the error might even satisfy these conditions, but it is large enough to significantly interfere with the optimisation problem. In these cases, we can still take advantage of probability estimates for the query locations given by independent systems, such as

localisation systems in the case of mobile-robotics applications. In addition, we can have a reasonably accurate adaptive model $\hat{\mathbb{P}}_{\mathbf{x}}^t$ for the actual querying distribution $\mathbb{P}_{\mathbf{x}}^E$, which is learnt online from the executed trajectories. Considering that, given the assumptions in Lemma 1, we extend Theorem 3 in Chowdhury and Gopalan [6] to bound the cumulative expected regret for uGP-UCB.

Theorem 2. *Let $\delta \in (0, 1)$, $f \in \mathcal{H}_k$, and $B \geq \|f\|_k$. Consider ν_t as σ_ν -sub-Gaussian noise. Assume that both k and location estimates \mathbb{P}_t^L satisfy the conditions for $\Delta f_{\mathbb{P}_t^L}$ to be σ_F -sub-Gaussian, for a given $\sigma_F > 0$, according to Lemma 1, for all $t \geq 1$. Then, running uGP-UCB with $\beta_t = B + \sqrt{2(\sigma_F^2 + \sigma_\nu^2)(\hat{\gamma}_{t-1} + 1 + \ln(1/\delta))}$ and a model $\hat{\mathbb{P}}_{\mathbf{x}}^t$, such that:*

$$\forall t \geq 1, \exists \rho_t \geq 0 : \max_{\mathbf{x} \in \mathcal{S}} \left| \mathbb{E}_{\mathbb{P}_{\mathbf{x}}^E}[f] - \mathbb{E}_{\hat{\mathbb{P}}_{\mathbf{x}}^t}[f] \right| \leq \rho_t,$$

we obtain cumulative expected regret bounded by $\mathcal{O}(\sqrt{T}(B\sqrt{\gamma_T} + \hat{\gamma}_T) + \sum_{t=1}^T \rho_t)$ with probability at least $1 - \delta$.

For this theorem we basically reapplied the theoretical results for IGP-UCB in Chowdhury and Gopalan [6], whose derivation is mostly independent from the kind of input space, while taking into account the effects of distribution mismatch between $\hat{\mathbb{P}}_{\mathbf{x}}^t$ and $\mathbb{P}_{\mathbf{x}}^E$. Notice that although β_t has a similar formulation for both IGP-UCB and uGP-UCB, σ_F is derived from \mathbb{P}_t^L for the latter, while depending on $\mathbb{P}_{\mathbf{x}}^E$ for the former. Hence, lower-noise estimates from \mathbb{P}_t^L can benefit uGP-UCB.

Another factor in Theorem 2 is the approximation error between the model $\hat{\mathbb{P}}_{\mathbf{x}}^t$ and the actual $\mathbb{P}_{\mathbf{x}}^E$, quantified in terms of difference in expectations. Among several possibilities for learning an adaptive model for the execution noise, this error term is implementation dependent and can be bounded in terms of divergence measures via Pinsker's inequality [21]. As a remark, if $\hat{\mathbb{P}}_{\mathbf{x}}^t$ can asymptotically converge to the true $\mathbb{P}_{\mathbf{x}}^E$, with $\rho_t \rightarrow 0$ as $t \rightarrow \infty$, it is possible to obtain sub-linear regret for uGP-UCB. However, in the following, we provide a constant upper bound for ρ_t , which acts as a lower bound on \bar{r}_t , for a trivial case, when $\mathbb{P}_{\mathbf{x}}^E$ is Gaussian with a constant, target-independent covariance matrix.

Proposition 1. *Let $f \in \mathcal{H}_k$ and $B \geq \|f\|_k$. We have that, for any given $\mathbf{x} \in \mathcal{S}$:*

$$\forall t \geq 1, \left| \mathbb{E}_{\mathbb{P}_{\mathbf{x}}^E}[f] - \mathbb{E}_{\hat{\mathbb{P}}_{\mathbf{x}}^t}[f] \right| \leq B \sqrt{\frac{1}{2} D_{\text{KL}}(\mathbb{P}_{\mathbf{x}}^E \parallel \hat{\mathbb{P}}_{\mathbf{x}}^t)},$$

where $D_{\text{KL}}(\mathbb{P}_{\mathbf{x}}^E \parallel \hat{\mathbb{P}}_{\mathbf{x}}^t)$ represents the Kullback-Leibler divergence between the two distributions. In particular, if $\tilde{\mathbf{x}} \sim \mathbb{P}_{\mathbf{x}}^E$ is distributed according to a Gaussian with mean \mathbf{x} and covariance Σ^E for all $\mathbf{x} \in \mathcal{S}$ and every $t \geq 1$, using a Gaussian model for $\mathbb{P}_{\mathbf{x}}^t$ with same mean and a given constant covariance matrix $\hat{\Sigma}$, we have that $\forall t \geq 1, \forall \mathbf{x} \in \mathcal{S}$:

$$\left| \mathbb{E}_{\mathbb{P}_{\mathbf{x}}^E}[f] - \mathbb{E}_{\hat{\mathbb{P}}_{\mathbf{x}}^t}[f] \right| \leq \frac{B}{2} \left(\text{tr}(\hat{\Sigma}^{-1} \Sigma^E) - d + \ln \frac{|\hat{\Sigma}|}{|\Sigma^E|} \right)^{\frac{1}{2}}.$$

The remaining important term that appears in our regret bounds is the maximum information gain $\hat{\gamma}_T$ of the uncertain-inputs GP model. Considering a compact set $\mathcal{P}_S \subset \mathcal{P}$, we define:

$$\hat{\gamma}_T := \max_{\mathcal{V} \subset \mathcal{P}_S: |\mathcal{V}| \leq T} I(\hat{\mathbf{z}}_{\mathcal{V}}; \hat{\mathbf{f}}_{\mathcal{V}}), \quad (17)$$

where $I(\hat{\mathbf{z}}_{\mathcal{V}}; \hat{\mathbf{f}}_{\mathcal{V}})$ is the mutual information between the function $\hat{f} : \mathcal{P} \rightarrow \mathbb{R}$ modelled by the GP and a vector of observations $\hat{\mathbf{z}}_{\mathcal{V}} = \hat{\mathbf{f}}_{\mathcal{V}} + \boldsymbol{\eta}_{\mathcal{V}}$, with $\hat{\mathbf{f}}_{\mathcal{V}} = [\hat{f}(\mathbb{P})]_{\mathbb{P} \in \mathcal{V}}$ and $\boldsymbol{\eta}_{\mathcal{V}} \sim \mathcal{N}(\mathbf{0}, \sigma_{\eta}^2 \mathbf{I})$. The mutual information can be computed in the same way as in the deterministic-inputs case [17], i.e. $I(\hat{\mathbf{z}}_{\mathcal{V}}; \hat{\mathbf{f}}_{\mathcal{V}}) = \frac{1}{2} \log |\mathbf{I} + \sigma_{\eta}^{-2} \hat{\mathbf{K}}_{\mathcal{V}}|$, since under the GP assumption \hat{f} is Gaussian. Although more general bounds for $\hat{\gamma}_T$ are object of future work, in the case of the squared exponential kernel [20], we can still upper-bound it in terms of the maximum information gain for the corresponding deterministic-inputs GP model, as presented below.

Proposition 2. Consider $\mathcal{P}_S = \{\mathbb{P} \in \mathcal{P} \mid \mathbb{P}[S] = 1\}$, the set of probability measures with support fully contained in \mathcal{S} , which is assumed to be a compact and connected subspace of \mathcal{X} . Let k be a continuous positive-definite kernel over \mathcal{S} and \hat{k} its uncertain-inputs counterpart, according to Equation 16. Then we have that:

$$\hat{\gamma}_T \leq \gamma_T,$$

where $\hat{\gamma}_T$ is the maximum information gain after collecting $T \geq 1$ observations over \mathcal{P}_S for a GP with covariance \hat{k} , and γ_T is the equivalent for a GP with covariance k over \mathcal{S} .

The above proposition suggests that we can apply upper bounds available for the deterministic-inputs γ_T [17, Theorem 5] to bound $\hat{\gamma}_T$. This is mainly due to the fact that the values of the entries of the kernel matrix $\hat{\mathbf{K}}_T$ can be upper-bounded by those of a deterministic-inputs version \mathbf{K}_T for some set $\{\mathbf{x}_i\}_{i=1}^T \subset \mathcal{S}$, which leads to a smaller information gain at each round, as perceived by the uncertain-inputs GP model. Moreover, since the regret bounds for IGP-UCB and uGP-UCB are a function of their respective maximum information gain, which is also true for other theoretically-proven GP-UCB strategies [17], Theorem 2 also implies that uGP-UCB should be able to obtain lower expected regret in optimisation problems with input noise than other GP-UCB methods.

VI. 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 1), and the unscented Expected improvement (UEI) heuristics [5], which applies the unscented transform to the EI 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 function and the location at which the sample is taken are significantly stochastic.

In all experiments, the inputs were corrupted by Gaussian noise. As kernel for all methods we used the squared exponential [20] with its corresponding version for uncertain




Fig. 2: Mean expected regret for IGP-UCB, UEI, and uGP-UCB in the optimisation of functions in the same RKHS: in (a) the UCB confidence-bound parameter β_t was set according to the theoretical results, while in (b), it was fixed at $\beta_t = 3$ for all t . The results were averaged over 20 trials, and the shaded areas correspond to one standard deviation.

inputs, whose analytical solution to Equation 7 is known [12, Eq. 3.53].

In every experiment, the optimisation search space \mathcal{S} is a continuous, compact subset of \mathbb{R}^d . As main performance metrics, we used the mean expected regret, i.e. $\frac{1}{T} \sum_{t=1}^T \bar{r}_t$ (Equation 13), except for the robotic exploration problem, where we used the conventional average regret, $\frac{1}{T} \sum_{t=1}^T y^* - f(\mathbf{x}_t)$, where $y^* = \max_{\mathbf{x} \in \mathcal{S}} f(\mathbf{x})$ is the deterministic-inputs maximiser.

A. Objective functions in the same RKHS

In this experiment, at each trial a different objective function was generated by uniformly sampling a set of 30 weights α_i and support points \mathbf{x}_i for the kernels, with $f(\cdot) = \sum_{i=1}^n \alpha_i k(\cdot, \mathbf{x}_i) \in \mathcal{H}_k$. The search space was set to the unit box $\mathcal{S} = [0, 1]^2 \subset \mathbb{R}^2$. As parameters to verify the theoretical results for the UCB algorithms, we set $\delta = 0.4$, and computed B directly from the f . The querying execution noise on $\mathbb{P}_{\mathbf{x}}^E$ was generated from i.i.d from $\mathcal{N}(\mathbf{0}, \sigma_{\mathbf{x}}^2 \mathbf{I})$ with $\sigma_{\mathbf{x}} = 0.1$. The output noise parameters for the GP model were computed according to Lemma 1, but each method was configured assuming an underestimated execution noise distribution $\mathcal{N}(\mathbf{0}, \hat{\sigma}_{\mathbf{x}}^2 \mathbf{I})$, with $\hat{\sigma}_{\mathbf{x}} = \frac{1}{2} \sigma_{\mathbf{x}}$. Noise on the localisation estimates \mathbb{P}_t^L was also set at half the standard deviation of the execution noise. Kernel length-scales were set isotropic to 0.1. The maximum information gain was set according to the results in Srinivas et al. [17].

Figure 2a presents performance results, in terms of mean expected regret, obtained by uGP-UCB under the assumptions of Theorem 2, IGP-UCB with parameters adapted according to Theorem 1's settings, and UEI. As the results show, in this setup, uGP-UCB is able to outperform IGP-UCB, if following the setup in Theorem 2. However, UEI is able to significantly outperform both UCB methods. One of the main reasons for that is that cumulative regret bounds are valid up to a constant factor, due to the use of the \mathcal{O} notation, and their focus is on guaranteeing asymptotic convergence, i.e. $\lim_{T \rightarrow \infty} \bar{R}_T/T = 0$, as most theoretical results in the UCB literature [17, 6]. Thus,

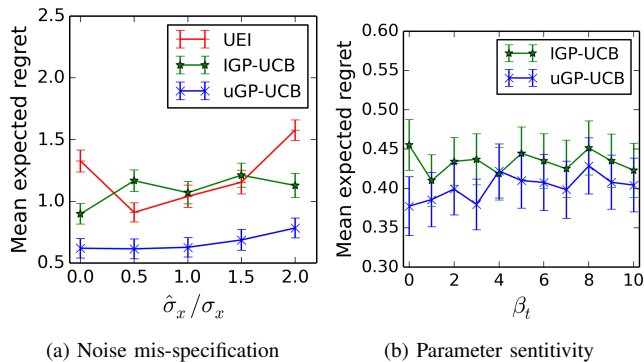


Fig. 3: Robustness: (a) presents how each algorithm’s performance is affected under mis-specification of the execution noise variance for the experiment with functions in the same RKHS, while (b) presents how the UCB methods performance varies under different settings for β_t in the case of the objective function in a different RKHS. In both cases, the mean expected regret shown is computed at the end of 200 iterations and averaged over 10 trials.

the method might need many more iterations to converge, at prohibitive computational costs due to the GP $\mathcal{O}(n^3)$ update complexity [20].

As Figure 2b presents, for a fixed β_t , UCB-based methods can quickly outperform UEI. For the latter, the expected improvement over the best objective function value observed so far is computed over the GP model by applying the unscented transform to the standard EI estimate. This transformed estimate corresponds to the expected value of the EI under the input-noise distribution, which in this experiment is under-estimated with a quarter of the true variance. As Figure 3a presents, UCB methods are more robust to mis-specification of the execution noise having smaller variation in performance, being actually favoured by assuming less execution noise than the true amount. This behaviour can be explained by considering that assuming less execution noise, GP-UCB’s estimated upper bound on the objective function value is higher than needed, but still a valid upper bound, maintaining the algorithm’s assumptions.

B. Objective function in different RKHS

To verify uGP-UCB’s performance under wrong kernel assumptions, we performed tests with an objective function in a different RKHS. This function, shown in Figure 4a also appeared in experiments in Nogueira et al. [5]. Figure 4b presents performance results for fixed $\beta_t = 3$, while in Figure 3b we evaluated each algorithm’s sensitivity to the choice of β_t , so we could assess the robustness of the methods when theoretical assumptions are not met. Input noise was set with $\sigma_x = 0.01$. As seen, both the proposed uGP-UCB and IGP-UCB can outperform the unscented BO approach UEI. In addition, we can see that uGP-UCB’s shows consistently better performance than that of IGP-UCB across varying settings for β_t .

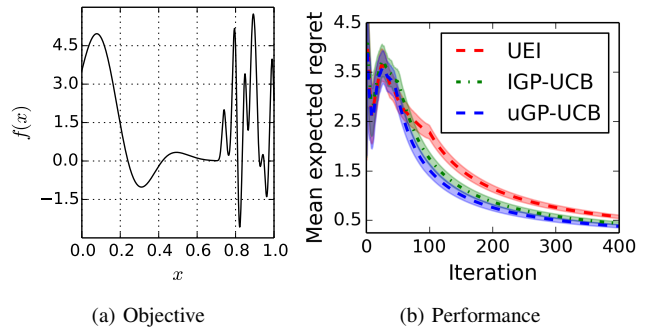


Fig. 4: Performance results for objective function in different RKHS: (a) presents the objective function and (b) the mean expected regret for each algorithm. Results were averaged over 10 trials.

C. Robotic exploration problem


In this section, we present results obtained in a simulated robotic exploration problem. In this experiment, a robot is set to explore an environmental process by measuring the concentration of potassium in the soil of a farm. The underlying process is based on the Broom’s Barn dataset ¹, consisting of the log-concentration of potassium in the soil of an experimental agricultural area. This data was spread over a 36 by 62 metres search space as shown in Figure 5a. 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 we used the OpenRobots’ Morse simulator ². In this scenario, execution noise is not following a stationary distribution due to the dynamic constraints of the robot, the imperfections in motion control etc. Figure 5b presents a typical execution noise distribution, when marginalised over different starting and goal locations from a given run of the experiment. As seen in the figure, the noise distribution is not Gaussian, but it’s concentrated on average 3 metres away from the target location, since this was the threshold given to the robot’s path-following control as the minimum distance to a target location. The noise around the ring formed by the distribution is a consequence of localisation uncertainty and imperfections on the robot’s motion control. The hyper-parameters for each GP were learnt online via log-marginal likelihood maximisation, after each sample. We set β_t at a fixed value for this experiment, again with $\beta_t = 3$. Figure 5c presents the performance of each algorithm in terms of regret. The plots show that uGP-UCB is able to significantly outperform UEI, while performing still better than IGP-UCB in the long run.

VII. CONCLUSION


In this paper we proposed a novel way to optimise functions where both the sampling of the function as well as the location

¹Dataset available at <http://www.kriging.com/datasets/>

²<https://www.openrobots.org/morse>



(a) Potassium (K) concentration (b) Execution error



(c) Robotics problem

Fig. 5: Robotics exploration experiment: (a) presents the Brom’s barn data as distributed over the search space; (b) presents a typical execution error distribution as observed from one of the runs in the experiment; and (c) shows the performance of each BO approach, averaged over 4 runs.

at which a function is sampled are stochastic. We provided theoretical guarantees on the optimality of the proposed uGP-UCB algorithm and the standard IGP-UCB method, when adapted to the noisy-inputs settings. In experiments we were able to show competitive performance when compared to state-of-the-art methods [5], while at the same time providing the first theoretical guarantees for this kind of setup. For future work, it would be worth investigating learning techniques for the approximate querying distribution $\hat{\mathbb{P}}_{\mathbf{x}}$ and better bounds for the uncertain-inputs GP model information gain, probably by means of information geometry [24] techniques.

APPENDIX

In this section we derive the proofs for the theoretical results in the paper. We start by providing a summary of the notation used throughout the paper. Following that, we present some theoretical background with definitions and basic lemmas applied in the proofs. And then we follow with the proof for each result, starting with Lemma 1. For reference, a notation summary is presented in Table I.

A. Background

Definition 1. A real-valued random variable η is said to be σ_η -sub-Gaussian if:

$$\forall \lambda \in \mathbb{R}, \mathbb{E}[e^{\lambda\eta}] \leq e^{\lambda^2\sigma_\eta^2/2}. \quad (18)$$

TABLE I: Notation

\mathbb{R}	the field of real numbers, or the real line
\mathbb{R}^d	the Euclidean space of d -dimensional vectors
\mathcal{S}	BO’s search space, a subset of \mathbb{R}^d
\mathbf{x}	a location vector, considered as an element of \mathbb{R}^d
$\tilde{\mathbf{x}}$	random location, an \mathbb{R}^d -valued random variable
f	objective function
\mathcal{X}	input domain of the objective function, often \mathbb{R}^d
\mathcal{P}	set of all probability measures on \mathcal{X}
\mathbb{P}	a probability measure or distribution
\mathbb{P}_t^L	location distribution estimated by localisation
\mathbb{P}_x^E	query location distribution given a target \mathbf{x}
$\hat{\mathbb{P}}_{\mathbf{x}}$	model for \mathbb{P}_x^E used by uGP-UCB
$\mu_{\mathbb{P}}$	mean map of \mathbb{P}
k	a positive-definite kernel
\mathcal{H}_k	reproducing kernel Hilbert space of k

Definition 2. Let $\mathbf{x} = [x_1, \dots, x_d]^\top$ and $\mathbf{x}'_i = [x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_d]^\top$, where $x_i, x'_i \in \mathcal{Q}_i \subset \mathbb{R}$ and $\mathcal{Q} = (\mathcal{Q}_1 \times \dots \times \mathcal{Q}_d)$. A function $f : \mathcal{Q} \rightarrow \mathbb{R}$ has the bounded differences property if:

$$|f(\mathbf{x}) - f(\mathbf{x}'_i)| \leq c_i, \quad \forall i \in \{1, \dots, d\}, \quad (19)$$

where c_i are non-negative constants.

Lemma 2 (Corollary 4.36 in Christmann and Steinwart [25]). Let $f \in \mathcal{H}_k$, where $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is a twice-differentiable kernel on $\mathcal{X} \subseteq \mathbb{R}^d$. Then f has bounded first-order partial derivatives, such that for any $\mathbf{x} \in \mathcal{X}$:

$$\left| \frac{\partial f(\mathbf{x})}{\partial x_i} \right| \leq \|f\|_k \sqrt{\left. \frac{\partial^2 k(\mathbf{x}, \mathbf{x}')}{\partial x_i \partial x'_i} \right|_{\mathbf{x}=\mathbf{x}'}}. \quad (20)$$

Lemma 3 (Theorem 5.5 in Boucheron et al. [21]). Let $\tilde{\mathbf{x}} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ be an \mathbb{R}^d -valued standard Gaussian random vector. Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ denote a L -Lipschitz function, i.e. $\exists L > 0$, such that:

$$|f(\mathbf{x}) - f(\mathbf{x}')| \leq L\|\mathbf{x} - \mathbf{x}'\|_2, \quad \forall \mathbf{x}, \mathbf{x}' \in \mathbb{R}^d. \quad (21)$$

Then, for all $\lambda \in \mathbb{R}$:

$$\mathbb{E}[e^{\lambda(f(\tilde{\mathbf{x}}) - \mathbb{E}[f(\tilde{\mathbf{x}})])}] \leq e^{\frac{1}{2}\lambda^2 L^2}. \quad (22)$$

B. Proof of Lemma 1

In this section, we derive proofs for Lemma 1, which is split in two parts. We start with the case where the inputs follow a Gaussian distribution and then the case for arbitrary probability distributions with bounded support.

1) *Gaussian inputs:* In the case of Gaussian inputs, Lemma 1 is a direct consequence of Lemma 3 when applied to functions $f \in \mathcal{H}_k$. Notice that, by the definition of \mathcal{H}_k , any f in it is continuously differentiable, according to Lemma 2, and therefore Lipschitz continuous. All we have to do is to generalise the inequality in Equation 22 for the case of general Gaussian random vectors $\tilde{\mathbf{x}} \sim \mathcal{N}(\hat{\mathbf{x}}, \Sigma)$.

If $\tilde{\mathbf{x}}^s$ is a standard Gaussian random vector, $\tilde{\mathbf{x}} = \hat{\mathbf{x}} + \mathbf{A}\tilde{\mathbf{x}}^s$, where $\Sigma = \mathbf{A}\mathbf{A}^\top$, due to the translational and rotational invariance of Gaussian random vectors. We can define a function g , such that:

$$g(\tilde{\mathbf{x}}^s) = f(\hat{\mathbf{x}} + \mathbf{A}\tilde{\mathbf{x}}^s) = f(\tilde{\mathbf{x}}). \quad (23)$$

Since f is Lipschitz continuous, g also is, for some Lipschitz constant L_g . Then we can apply Lemma 3 to g , which yields:

$$\mathbb{E}[e^{\lambda(g(\tilde{\mathbf{x}}^s) - \mathbb{E}[g(\tilde{\mathbf{x}}^s)])}] \leq e^{\frac{1}{2}\lambda^2 L_g^2}. \quad (24)$$

In addition, by definition (Equation 23), $g(\tilde{\mathbf{x}}^s)$ and $f(\tilde{\mathbf{x}})$ follow the same distribution, so that $\mathbb{E}[f(\tilde{\mathbf{x}})] = \mathbb{E}[g(\tilde{\mathbf{x}}^s)]$. As a result, $\Delta f_{\mathbb{P}} = f(\tilde{\mathbf{x}}) - \mathbb{E}[f(\tilde{\mathbf{x}})]$ is L_g -sub-Gaussian, according to Definition 1.

Now L_g is any constant uniformly upper-bounding the Euclidean norm of g 's gradient, and $\|\nabla g\|_2 = \|\mathbf{A}^\top \nabla f\|_2$. Without loss of generality, let's assume that Σ is a matrix of diagonal entries $\sigma_i^2, 1 \leq i \leq d$. Then we have that:

$$\begin{aligned} \|\nabla g\|_2^2 &= \|\mathbf{A}^\top \nabla f\|_2^2 \\ &= \nabla f^\top \mathbf{A} \mathbf{A}^\top \nabla f \\ &= \nabla f^\top \Sigma \nabla f \\ &= \sum_{i=1}^d \sigma_i^2 \left| \frac{\partial f}{\partial x_i} \right|^2 \leq L_f^2 \text{tr}(\Sigma), \end{aligned} \quad (25)$$

where $L_f = \|f\|_k L_k$. Therefore, the inequality in Equation 24 holds with $L_g = L_f \sqrt{\text{tr}(\Sigma)}$.

For a non-diagonal Σ , by spectral decomposition, we have that $\Sigma = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^\top$, where $\mathbf{\Lambda}$ is a matrix whose diagonal is composed of Σ 's eigenvalues and $\mathbf{V} \mathbf{V}^\top = \mathbf{I}$. Observe that the result in Equation 25 would also hold for a zero-mean Gaussian random vector $\tilde{\mathbf{x}}^v$ with covariance matrix $\mathbf{\Lambda}$. Then we could define $h(\tilde{\mathbf{x}}^v) = f(\tilde{\mathbf{x}} + \mathbf{V} \tilde{\mathbf{x}}^v)$ and follow similar steps to the ones we took for g . However, f and h , as defined, have the same Lipschitz constant, since:

$$\|\mathbf{V} \mathbf{x} - \mathbf{V} \mathbf{x}'\|_2^2 = (\mathbf{x} - \mathbf{x}')^\top \mathbf{V}^\top \mathbf{V} (\mathbf{x} - \mathbf{x}') = \|\mathbf{x} - \mathbf{x}'\|_2^2, \quad (26)$$

where we applied $\mathbf{V}^\top \mathbf{V} = \mathbf{V} \mathbf{V}^\top = \mathbf{I}$. In addition, as Σ is positive definite, $\text{tr}(\Sigma) = \text{tr}(\mathbf{\Lambda})$. Therefore, the same result in Equation 25 holds for general Σ and $\tilde{\mathbf{x}}$, which can also be seen as a consequence of the translational and rotational invariance of Gaussian random vectors. Making $\sigma_F = L_g = \|f\|_k L_k \text{tr}(\Sigma)^{1/2}$ concludes the proof.

2) *Bounded inputs:* By Lemma 2, we can observe that $f \in \mathcal{H}_k(\mathcal{X})$ is Lipschitz continuous with respect to the 1-norm on \mathbb{R}^d , in particular:

$$|f(\mathbf{x}) - f(\mathbf{x}')| \leq \|f\|_k L_k \|\mathbf{x} - \mathbf{x}'\|_1, \quad \forall \mathbf{x}, \mathbf{x}' \in \mathbb{R}^d, \quad (27)$$

where $L_k \geq 0$ is any constant, such that $L_k^2 \geq \sup_{\mathbf{x} \in \mathcal{X}} \sup_{\mathbf{x}' \in \mathcal{X}} \frac{\partial^2 k(\mathbf{x}, \mathbf{x}')}{\partial x_i \partial x'_i} \Big|_{\mathbf{x}=\mathbf{x}'}$. Therefore, according to Definition 2, f satisfies the bounded differences property for any \mathbf{x} in the support of \mathbb{P} with $c_i = \|f\|_k L_k \sigma_i$. Applying McDiarmid's inequality [26], we have that:

$$\Pr \{|f(\tilde{\mathbf{x}}) - \mathbb{E}_{\mathbb{P}}(f)| \geq t\} \leq 2 \exp \left(- \frac{2t^2}{\|f\|_k^2 L_k^2 \sum_{i=1}^d \sigma_i^2} \right). \quad (28)$$

As a result, $\Delta f_{\mathbb{P}}$ is σ_F -sub-Gaussian with $\sigma_F = \frac{1}{2} \|f\|_k L_k \sqrt{\sum_{i=1}^d \sigma_i^2}$, according to Definition 1 and Lemma 2.2 in Boucheron et al. [21].

C. Proof of Theorem 1

Theorem 1 is a direct application of the results for IGP-UCB in Chowdhury and Gopalan [6, Theorem 3], which are conditioned on a few important assumptions. Firstly, the objective function needs to be a member of \mathcal{H}_k . For that, we define a function $g : \mathcal{S} \rightarrow \mathbb{R}$ mapping a point $\mathbf{x} \in \mathcal{S}$ to the expected value of $f \in \mathcal{H}_k$ under the querying probability distribution, i.e. $g(\mathbf{x}) = \mathbb{E}_{\mathbb{P}_{\tilde{\mathbf{x}}}}[f]$. For a general mapping $\mathbf{x} \rightarrow \mathbb{P}_{\tilde{\mathbf{x}}}^E$, g is not necessarily a member of \mathcal{H}_k , since the location-noise ϵ_t distribution might change with \mathbf{x} . However, under the assumptions of Theorem 1, Lemma 4 shows that there's a $g \in \mathcal{H}_k$ that can be used as a proxy for the expected value of f given a target \mathbf{x} . A second requirement is that the observation noise as perceived by the algorithm, i.e. $\eta_t = z_t - \mathbb{E}_{\mathbb{P}_{\tilde{\mathbf{x}}_t}^E}[f]$, must be conditionally sub-Gaussian for the regret bounds to hold [6]. Due to the assumptions on k and $\mathbb{P}_{\tilde{\mathbf{x}}}^E$, this condition is more easily satisfied, which we will show in the second part of this proof.

Lemma 4. *Let \mathcal{H}_k be the reproducing kernel Hilbert space (RKHS) associated with a positive-definite kernel $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$. Assume that k is translation-invariant, so that, for $\epsilon \sim \mathbb{P}_q$ and any given $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$:*

$$\int k(\mathbf{x} + \epsilon, \mathbf{x}') d\mathbb{P}_q(\epsilon) = \int k(\mathbf{x}, \mathbf{x}' - \epsilon) d\mathbb{P}_q(\epsilon). \quad (29)$$

As a result, for any given $f \in \mathcal{H}_k$, we have that:

$$\exists g \in \mathcal{H}_k, \text{ s.t. } g(\mathbf{x}) = \mathbb{E}_{\mathbb{P}_q}[f(\mathbf{x} + \epsilon)], \quad \forall \mathbf{x} \in \mathcal{X}. \quad (30)$$

and $\|g\|_k \leq \|f\|_k$.

Proof. By definition, any function in an RKHS \mathcal{H}_k can be described as:

$$f(\mathbf{x}) = \sum_{i=1}^n \alpha_i k(\mathbf{x}, \mathbf{x}_i) \quad (31)$$

for $\mathbf{x} \in \mathcal{X}$, $n \in \mathbb{N}$, $\alpha_i \in \mathbb{R}$ and $\mathbf{x}_i \in \mathcal{X}$, $\forall i \in \{1, \dots, n\}$. Consequently, by linearity in the expectation, the expected value of $f(\mathbf{x} + \epsilon)$, with $\epsilon \sim \mathbb{P}_q$ is given by:

$$\mathbb{E}_{\epsilon \sim \mathbb{P}_q}[f(\mathbf{x} + \epsilon)] = \sum_{i=1}^n \alpha_i \mathbb{E}_{\mathbb{P}_q}[k(\mathbf{x} + \epsilon, \mathbf{x}_i)]. \quad (32)$$

With k satisfying the assumptions in the lemma, we have that:

$$\begin{aligned} \forall i, \mathbb{E}_{\epsilon \sim \mathbb{P}_q}[k(\mathbf{x} + \epsilon, \mathbf{x}_i)] &= \mathbb{E}_{\epsilon \sim \mathbb{P}_q}[k(\mathbf{x}, \mathbf{x}_i - \epsilon)] \\ &= \mathbb{E}_{\tilde{\mathbf{x}}_i \sim \mathbb{P}_i}[k(\mathbf{x}, \tilde{\mathbf{x}}_i)] \\ &= \langle k(\cdot, \mathbf{x}), \mu_{\mathbb{P}_i} \rangle_k, \end{aligned} \quad (33)$$

where, in the second line, we defined $\tilde{\mathbf{x}}_i := \mathbf{x}_i - \epsilon_q$ and \mathbb{P}_i is $\tilde{\mathbf{x}}_i$'s equivalent distribution. In the last line, we applied the mean map $\mu_{\mathbb{P}}$ definition. As a result, we obtain:

$$\begin{aligned} \mathbb{E}_{\epsilon_q \sim \mathbb{P}_q}[f(\mathbf{x} + \epsilon_q)] &= \sum_{i=1}^n \alpha_i \langle \mu_{\mathbb{P}_i}, k(\cdot, \mathbf{x}) \rangle_k \\ &= \left\langle \sum_{i=1}^n \alpha_i \mu_{\mathbb{P}_i}, k(\cdot, \mathbf{x}) \right\rangle_k = g(\mathbf{x}), \end{aligned} \quad (34)$$

where $g = \sum_{i=1}^n \alpha_i \mu_{\mathbb{P}_i} \in \mathcal{H}_k$. In addition, since $\sup_{\mathbf{x} \in \mathcal{X}} g(\mathbf{x}) \leq \sup_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$, we have that $\|g\|_k \leq \|f\|_k$. \square

Basically Lemma 4 tells us that, case the execution noise $\epsilon_t := \tilde{\mathbf{x}}_t - \mathbf{x}_t$, with $\tilde{\mathbf{x}}_t \sim \mathbb{P}_{\mathbf{x}_t}^E$, has a resulting distribution \mathbb{P}_q , regardless of \mathbf{x}_t , we can find a proxy $g \in \mathcal{H}_k$ that can behave as an objective function for IGP-UCB. Such condition on \mathbb{P}_q is met under the assumptions of Theorem 1. Now we proceed to show that the observation noise is conditionally sub-Gaussian.

In our case, the observation noise is given by:

$$\eta_t = z_t - \mathbb{E}_{\mathbb{P}_{\mathbf{x}_t}^E}[f] = \nu_t + f(\tilde{\mathbf{x}}_t) - \mathbb{E}_{\mathbb{P}_{\mathbf{x}_t}^E}[f], \quad (35)$$

where $\tilde{\mathbf{x}}_t$ is distributed according to $\mathbb{P}_{\mathbf{x}_t}^E$ and ν_t is sub-Gaussian noise by definition. We also have that, for a given \mathbf{x}_t , the difference $f(\tilde{\mathbf{x}}_t) - \mathbb{E}_{\mathbb{P}_{\mathbf{x}_t}^E}[f]$ is a σ_F -sub-Gaussian random variable, according to Lemma 1 and the assumptions in Theorem 1. By its elementary properties, the sum of two independent sub-Gaussian random variables is also sub-Gaussian [27]. Therefore, the observation noise term η_t is σ_η -sub-Gaussian, with $\sigma_\eta \leq \sqrt{\sigma_\nu^2 + \sigma_F^2}$. Plugging this term back into the definitions in Theorem 1 concludes the proof.

D. Proof of Theorem 2

In a similar fashion to Theorem 1, Theorem 2 is also derived from the work in Chowdhury and Gopalan [6]. Though most of their results are stated in terms of an input space that is a subset of \mathbb{R}^d , none of the results that lead up to the regret bounds on IGP-UCB were derived based on particular characteristics, such as metrics, of Euclidean input spaces. Therefore, in this section we re-state the main results from that work, fitting them according to our settings, and then we analyse under which conditions on f the bounds would hold and what are the effects of an imperfect $\hat{\mathbb{P}}_{\mathbf{x}}$.

Generalising our problem formulation, consider a filtration $\{\mathcal{F}_t\}_{t=0}^\infty$, where \mathcal{F}_{t-1} is the σ -algebra [28] generated by the random variables \mathbb{P}_t and $\{\mathbb{P}_i, \eta_i\}_{i=1}^{t-1}$. The observation noise sequence $\{\eta_t\}_{t=1}^\infty$ is then a stochastic process that is adapted to that filtration. We also have that the sequence $\{\mathbb{P}_i\}_{i=1}^\infty$ can be modelled as a stochastic process that is predictable with respect to the same filtration $\{\mathcal{F}_t\}_{t=0}^\infty$. In the mobile robotics case, for example, this formulation abstracts what a localisation system would be doing when the robot obtains a location estimate \mathbb{P}_t^L at every round t . This can be considered as a stochastic process, since the distribution \mathbb{P}_t^L is actually a function of parameters estimated by the localisation system, which are subject to noise due to sensor imperfections, atmospheric interferences, etc. Hence, we can apply the following lemma to our settings.

Lemma 5 (Chowdhury and Gopalan [6], Theorem 1). *Let $\{\mathbb{P}_t\}_{t=1}^\infty$ be a \mathcal{P} -valued discrete-time stochastic process predictable with respect to the filtration $\{\mathcal{F}_t\}_{t=0}^\infty$, i.e. \mathbb{P}_t is \mathcal{F}_{t-1} -measurable $\forall t \geq 1$. Let $\{\eta_t\}_{t=1}^\infty$ be a real-valued stochastic process such that, for some $R \geq 0$ and for all $t \geq 1$, η_t is (a) \mathcal{F}_t -measurable, and (b) R -sub-Gaussian conditionally on \mathcal{F}_{t-1} . Let $\hat{k} : \mathcal{P} \times \mathcal{P} \rightarrow \mathbb{R}$ be a symmetric,*

positive-semidefinite kernel, and let $0 < \delta \leq 1$. For a given $\zeta > 0$, with probability at least $1 - \delta$, the following holds simultaneously for all $t \geq 0$:

$$\|\boldsymbol{\eta}_T\|_{((\hat{\mathbf{K}}_T + \zeta \mathbf{I})^{-1} + \mathbf{I})^{-1}}^2 \leq 2R^2 \log \frac{\sqrt{|(1 + \zeta)\mathbf{I} + \hat{\mathbf{K}}_T|}}{\delta} \quad (36)$$

Proof. The lemma above is simply a re-statement of Theorem 1 in Chowdhury and Gopalan [6]. Since the proof for that Theorem does not depend on any feature of the kernel input space itself, but only on the conditional sub-Gaussianity of the noise and the definition, it can be extended to a GP model that has as input domain a set of probability measures, such as \mathcal{P} . \square

Lemma 6. *Under the same hypotheses as those of Lemma 5, let $g : \mathcal{P} \rightarrow \mathbb{R}$ be a member of the RKHS of real-valued functions on \mathcal{P} with kernel \hat{k} , with RKHS norm bounded by B . Then, with probability at least $1 - \delta$, the following holds for all $\mathbb{P} \in \mathcal{P}$ and $t \geq 1$:*

$$|m_{t-1}(\mathbb{P}) - g(\mathbb{P})| \leq (B + R\sqrt{2(\hat{\gamma}_{t-1} + 1 + \log(1/\delta))}v_{t-1}^{1/2}(\mathbb{P})), \quad (37)$$

where $\hat{\gamma}_{t-1}$ is the maximum information gain after $t - 1$ rounds, and m_{t-1} and v_{t-1} are the posterior mean and variance of a Gaussian process model with covariance function \hat{k} with σ_η^2 set to $1 + \zeta$, and $\zeta = 2/T$.

Proof. This result follows from Theorem 2 in Chowdhury and Gopalan [6], being only re-stated for an input space \mathcal{P} of probability measures. The proof for this result in Chowdhury and Gopalan [6] is entirely based on Lemma 5 and the associated feature map of \hat{k} , which exists for any reproducing kernel [25]. \square

Now consider the set $\mathcal{P}_S^E = \{\mathbb{P}_{\mathbf{x}}^E \in \mathcal{P} : \mathbf{x} \in \mathcal{S}\}$, which is the image of the mapping $\mathbf{x} \rightarrow \mathbb{P}_{\mathbf{x}}^E$ given by the query execution process. Let's assume that a second mapping $\mathbf{x} \rightarrow \hat{\mathbb{P}}_{\mathbf{x}}^E$ satisfies the following condition:

$$\forall t \geq 1, \exists \rho_t \geq 0 : \max_{\mathbf{x} \in \mathcal{S}} \left| \mathbb{E}_{\mathbb{P}_{\mathbf{x}}^E}[f] - \mathbb{E}_{\hat{\mathbb{P}}_{\mathbf{x}}^E}[f] \right| \leq \rho_t. \quad (38)$$

Then, by Lemma 6, we would have that:

$$\forall \mathbf{x} \in \mathcal{S}, |m_{t-1}(\hat{\mathbb{P}}_{\mathbf{x}}^E) - \mathbb{E}_{\mathbb{P}_{\mathbf{x}}^E}[f]| \leq \beta_t v_{t-1}^{1/2}(\hat{\mathbb{P}}_{\mathbf{x}}^E) + \rho_t, \quad (39)$$

and:

$$\max_{\mathbf{x} \in \mathcal{S}} \mathbb{E}_{\mathbb{P}_{\mathbf{x}}^E}[f] \leq \rho_t + \max_{\mathbf{x} \in \mathcal{S}} m_{t-1}(\hat{\mathbb{P}}_{\mathbf{x}}^E) + \beta_t v_{t-1}^{1/2}(\hat{\mathbb{P}}_{\mathbf{x}}^E), \quad (40)$$

both with probability at least $1 - \delta$ for all $t \geq 1$. Consequently:

$$\bar{r}_t = \max_{\mathbf{x} \in \mathcal{S}} \mathbb{E}_{\mathbb{P}_{\mathbf{x}}^E}[f] - \mathbb{E}_{\mathbb{P}_{\mathbf{x}_t}^E}[f] \leq 2\beta_t v_{t-1}^{1/2}(\hat{\mathbb{P}}_{\mathbf{x}_t}^E) + 2\rho_t, \quad (41)$$

where \mathbf{x}_t is the maximiser of the right-hand side in Equation 40. Hence, $\sum_{t=1}^T \bar{r}_t \leq 2\beta_T \sum_{t=1}^T (v_{t-1}^{1/2}(\hat{\mathbb{P}}_{\mathbf{x}_t}^E) + \rho_t)$. In addition, $\sum_{t=1}^T v_{t-1}^{1/2}(\hat{\mathbb{P}}_{\mathbf{x}_t}^E) \leq \mathcal{O}(\sqrt{T\hat{\gamma}_T^E})$ [6, Lemma 4], where for us $\hat{\gamma}_T^E$ is the maximum information gain after T rounds for a GP over the set \mathcal{P}_S^E . Combining that to the maximum information gain $\hat{\gamma}_T^L$ after T rounds for a GP over a set \mathcal{P}^L

composed of all possible location estimates \mathbb{P}_t^L , we have the following result:

$$\sum_{t=1}^T \bar{r}_t = \mathcal{O} \left(B\sqrt{T\hat{\gamma}_T^E} + \sqrt{T\hat{\gamma}_T^E(\hat{\gamma}_T^L + \log(1/\delta))} + \sum_{t=1}^T \rho_t \right) \quad (42)$$

with probability at least $1 - \delta$. Since the maximum information gain $\hat{\gamma}_T$ over a larger set \mathcal{P} , where $\mathcal{P}_S^E \cup \mathcal{P}^L \subset \mathcal{P}$ upper bounds both $\hat{\gamma}_T^E$ and $\hat{\gamma}_T^L$ [17], we can simplify the equation above to:

$$\bar{R}_T = \mathcal{O} \left(\sqrt{T}(B\sqrt{\hat{\gamma}_T} + \hat{\gamma}_T) + \sum_{t=1}^T \rho_t \right), \quad (43)$$

which holds with probability at least $1 - \delta$.

E. Proof of Proposition 1

Proposition 1 refers to the approximation error between the model $\hat{\mathbb{P}}_{\mathbf{x}}^t$ and the actual distribution $\mathbb{P}_{\mathbf{x}}^E$ in terms of difference in the expected value of a function $f \in \mathcal{H}_k$. For the first part of the result, we simply applied Pinsker's inequality [21, Sec. 4.11].

We can define densities $\hat{p}_{\mathbf{x}}^t = \frac{d\hat{\mathbb{P}}_{\mathbf{x}}^t}{dv}$ and $p_{\mathbf{x}} = \frac{d\mathbb{P}_{\mathbf{x}}^E}{dv}$, where v denotes a common dominating measure for $\hat{\mathbb{P}}_{\mathbf{x}}^t$ and $\mathbb{P}_{\mathbf{x}}^E$, e.g. the Lebesgue measure on \mathcal{X} . Then we can write the difference in expectations in terms of $\mathcal{L}^2(v)$ inner products and apply Hölder's inequality:

$$\mathbb{E}_{\mathbb{P}_{\mathbf{x}}^E}[f] - \mathbb{E}_{\hat{\mathbb{P}}_{\mathbf{x}}^t}[f] = \langle f, \hat{p}_{\mathbf{x}}^t - p_{\mathbf{x}} \rangle_{\mathcal{L}^2(v)} \leq \|f\|_{\infty} \|\hat{p}_{\mathbf{x}}^t - p_{\mathbf{x}}\|_1, \quad (44)$$

where $\|\cdot\|_{\infty}$ and $\|\cdot\|_1$ denote the $\mathcal{L}^{\infty}(v)$ and the $\mathcal{L}^1(v)$ norms, respectively. Now note that, by Cauchy-Schwartz, for any $\mathbf{x} \in \mathcal{X}$, $f(\mathbf{x}) = \langle f, k(\cdot, \mathbf{x}) \rangle_k \leq \|f\|_k \|k(\cdot, \mathbf{x})\|_k = \|f\|_k$, since $k(\mathbf{x}, \mathbf{x}) = 1$ under our general assumptions. Then we have that:

$$\|f\|_{\infty} \leq \|f\|_k \leq B. \quad (45)$$

The second term in the Equation 44 is equivalent to half the total variation distance between the two probability measures. Therefore, we apply Pinsker's inequality and have that:

$$\|\hat{p}_{\mathbf{x}}^t - p_{\mathbf{x}}\|_1 \leq \sqrt{\frac{1}{2} D_{\text{KL}}(\mathbb{P}_{\mathbf{x}}^E \|\hat{\mathbb{P}}_{\mathbf{x}}^t)}. \quad (46)$$

Plugging this and (45) back into Equation 44 yields the first result in Proposition 1.

For the second part of Proposition 1, the Kullback-Leibler divergence between two Gaussians $\mathbb{P}_{\mathbf{x}}^E$ and $\hat{\mathbb{P}}_{\mathbf{x}}^t$ with same mean and covariance matrices as stated is given by:

$$D_{\text{KL}}(\mathbb{P}_{\mathbf{x}}^E \|\hat{\mathbb{P}}_{\mathbf{x}}^t) = \frac{1}{2} \left(\text{tr}(\hat{\Sigma}^{-1} \Sigma^E) - d + \ln \frac{|\hat{\Sigma}|}{|\Sigma^E|} \right), \quad (47)$$

which comes from a known result [20, Eq. A.23] and the fact that $\text{tr}(\hat{\Sigma}^{-1}(\Sigma^E - \hat{\Sigma})) = \text{tr}(\hat{\Sigma}^{-1} \Sigma^E) - \text{tr}(\hat{\Sigma}^{-1} \hat{\Sigma})$. Replacing this term in the first equation in the proposition yields its second result.

F. Proof of Proposition 2

Given a compact set \mathcal{S} where a BO algorithm chooses its target locations from, the maximum information gain after T iterations is defined as:

$$\gamma_T := \max_{\mathcal{Q} \subset \mathcal{S}; |\mathcal{Q}| \leq T} I(\mathbf{z}_{\mathcal{Q}}; \mathbf{f}_{\mathcal{Q}}), \quad (48)$$

where $I(\mathbf{z}_{\mathcal{Q}}; f)$ is the mutual information between f and a vector of observations $\mathbf{z}_{\mathcal{Q}} = \mathbf{f}_{\mathcal{Q}} + \boldsymbol{\eta}_{\mathcal{Q}}$, with $\mathbf{f}_{\mathcal{Q}} = [f(\mathbf{x})]_{\mathbf{x} \in \mathcal{Q}}$ and $\boldsymbol{\eta}_{\mathcal{Q}} \sim \mathcal{N}(\mathbf{0}, \sigma_{\eta}^2 \mathbf{I})$. In the Gaussian process case, mutual information is given by [17]:

$$I(\mathbf{z}_{\mathcal{Q}}; \mathbf{f}_{\mathcal{Q}}) = \frac{1}{2} \log |\mathbf{I} + \sigma_{\eta}^{-2} \mathbf{K}_{\mathcal{Q}}| = \frac{1}{2} \sum_{i=1}^{|\mathcal{Q}|} \log(1 + \sigma_{\eta}^{-2} \lambda_i), \quad (49)$$

where λ_i denote the eigenvalues of $\mathbf{K}_{\mathcal{Q}}$.

For the uncertain-inputs GP model, by definition, equivalently we have:

$$\hat{\gamma}_T := \max_{\mathcal{V} \subset \mathcal{P}_S; |\mathcal{V}|=T} I(\hat{\mathbf{z}}_{\mathcal{V}}; \hat{\mathbf{f}}_{\mathcal{V}}). \quad (50)$$

As defined in Proposition 2, \mathcal{P}_S can be viewed as a compact subspace of \mathcal{P} under some metric, such as a q -norm, for $q \geq 1$, so that, by the maximum value theorem [29, Ch. 3, Thr. 6.4], the maximum in the above equation exists.

For the remainder we will make use of the following lemma.

Lemma 7. *Let $k(\mathcal{S} \times \mathcal{S})$ denote the image set of $\mathcal{S} \times \mathcal{S}$ under k . Under the assumptions of Proposition 2, we have that $\hat{k}(\mathcal{P}_S \times \mathcal{P}_S) \subset k(\mathcal{S} \times \mathcal{S})$.*

Proof. The proof follows by applying the generalised intermediate value theorem [29, p. 154]. For any $\mathbb{P}, \mathbb{P}' \in \mathcal{P}_S$, we have that:

$$\hat{k}(\mathbb{P}, \mathbb{P}') \leq \max_{\mathbf{x}, \mathbf{x}' \in \mathcal{S}} k(\mathbf{x}, \mathbf{x}') \quad (51)$$

$$\hat{k}(\mathbb{P}, \mathbb{P}') \geq \min_{\mathbf{x}, \mathbf{x}' \in \mathcal{S}} k(\mathbf{x}, \mathbf{x}'). \quad (52)$$

Since \mathcal{S} is connected and k is continuous, we have that $k(\mathcal{S} \times \mathcal{S}) \subset \mathbb{R}$ is connected [29, Ch. 3, Thr. 1.5]. As a result, (51) and (52) lead to $k(\mathcal{P}_S \times \mathcal{P}_S) \subset k(\mathcal{S} \times \mathcal{S})$. \square

Given any $\mathcal{V} \subset \mathcal{P}_S$, the above lemma implies that all entries of $\hat{\mathbf{K}}_{\mathcal{V}}$ have values in the image set of the deterministic-inputs k . Consequently, although the diagonal entries $[\hat{\mathbf{K}}_{\mathcal{V}}]_{ii} = \hat{k}(\mathbb{P}_i, \mathbb{P}_i)$ would be lower-valued, we can find a set of points $\mathcal{Q} \subset \mathcal{S}$ such that $\mathbf{K}_{\mathcal{Q}}$ matches $\hat{\mathbf{K}}_{\mathcal{V}}$ in the off-diagonal entries, making $\mathbf{K}_{\mathcal{Q}} - \hat{\mathbf{K}}_{\mathcal{V}}$ positive semi-definite. Now considering any $\mathcal{V}^* \in \arg\max_{\mathcal{V} \subset \mathcal{P}_S} I(\hat{\mathbf{z}}_{\mathcal{V}}; \hat{\mathbf{f}}_{\mathcal{V}})$, the latter implies that [30, 7.7.4]:

$$\lambda_i \geq \hat{\lambda}_i^*, \quad \forall i \in \{1, \dots, |\mathcal{V}^*|\}, \quad (53)$$

for some $\mathcal{Q} \subset \mathcal{S}$, $|\mathcal{Q}| = |\mathcal{V}^*|$, where $\hat{\lambda}_i^*$ and λ_i denote the eigenvalues of $\hat{\mathbf{K}}_{\mathcal{V}^*}$ and $\mathbf{K}_{\mathcal{Q}}$, respectively. As a result, we have:

$$\hat{\gamma}_T = \frac{1}{2} \sum_{i=1}^T \log(1 + \sigma_{\eta}^{-2} \hat{\lambda}_i^*) \leq \frac{1}{2} \sum_{i=1}^T \log(1 + \sigma_{\eta}^{-2} \lambda_i) \leq \gamma_T, \quad (54)$$

which concludes the proof.

ACKNOWLEDGEMENT

The authors would like to thank Kelvin Hsu and Prof. Lamiae Azizi for their helpful comments on the initial derivations of the theoretical results presented in this paper.

REFERENCES

- [1] E. Brochu, V. M. Cora, and N. de Freitas, “A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical Reinforcement Learning,” University of British Columbia, Tech. Rep., 2010.
- [2] J. Snoek, H. Larochelle, and R. P. Adams, “Practical bayesian optimization of machine learning algorithms,” in *Advances in Neural Information Processing Systems*, 2012.
- [3] A. Wilson, A. Fern, and P. Tadepalli, “Using Trajectory Data to Improve Bayesian Optimization for Reinforcement Learning,” *Journal of Machine Learning Research*, vol. 15, pp. 253–282, 2014.
- [4] R. Marchant and F. Ramos, “Bayesian Optimisation for Intelligent Environmental Monitoring,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2012.
- [5] J. Nogueira, R. Martinez-Cantin, A. Bernardino, and L. Jamone, “Unscented Bayesian Optimization for Safe Robot Grasping,” in *IEEE International Conference on Robotics and Automation (ICRA)*, Daejeon, Korea, 2016, pp. 1967–1972.
- [6] S. R. Chowdhury and A. Gopalan, “On Kernelized Multi-armed Bandits,” in *Proceedings of the 34th International Conference on Machine Learning (ICML)*, Sydney, Australia, 2017.
- [7] R. Oliveira, L. Ott, V. Guizilini, and F. Ramos, “Bayesian Optimisation for Safe Navigation under Localisation Uncertainty,” *arXiv:1709.02169*, 2017.
- [8] E. A. Wan and R. van der Merwe, “The Unscented Kalman Filter for Nonlinear Estimation,” in *Adaptive Systems for Signal Processing, Communications, and Control Symposium (AS-SPCC)*, 2000, pp. 153–158.
- [9] P. Dallaire, C. Besse, and B. Chaib-Draa, “An approximate inference with Gaussian process to latent functions from uncertain data,” *Neurocomputing*, vol. 74, pp. 1945–1955, 2011.
- [10] G. Pillonetto, G. Erinc, and S. Carpin, “Online Estimation of Covariance Parameters using Extended Kalman Filtering and Application to Robot Localization,” *Advanced Robotics*, vol. 26, no. 18, pp. 2169–2188, 2012.
- [11] A. Girard, C. E. Rasmussen, J. Q. Candela, and R. Murray-Smith, “Gaussian process priors with uncertain inputs-application to multiple-step ahead time series forecasting,” in *Advances in neural information processing systems*, 2003, pp. 545–552.
- [12] A. Girard, “Approximate methods for propagation of uncertainty with Gaussian process models,” Ph. D, University of Glasgow, 2004.
- [13] A. Mchutchon and C. E. Rasmussen, “Gaussian Process Training with Input Noise,” in *Advances in Neural Information Processing Systems*, 2011, pp. 1341–1349.
- [14] A. C. Damianou, M. K. Titsias, and N. D. Lawrence, “Variational Inference for Latent Variables and Uncertain Inputs in Gaussian Processes,” *Journal of Machine Learning Research*, vol. 17, no. 1, pp. 1–62, 2016.
- [15] K. Muandet, K. Fukumizu, F. Dinuzzo, and B. Schölkopf, “Learning from Distributions via Support Measure Machines,” in *Advances in Neural Information Processing Systems*, 2012.
- [16] K. Muandet, K. Fukumizu, B. Sriperumbudur, and B. Schölkopf, “Kernel Mean Embedding of Distributions: A Review and Beyond,” *arXiv*, 2016.
- [17] N. Srinivas, A. Krause, S. M. Kakade, and M. W. Seeger, “Information-Theoretic Regret Bounds for Gaussian Process Optimization in the Bandit Setting,” *IEEE Transactions on Information Theory*, vol. 58, no. 5, 2012.
- [18] R. Marchant, F. Ramos, and S. Sanner, “Sequential Bayesian Optimisation for Spatial-Temporal Monitoring,” in *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2014.
- [19] C. K. Ling, K. H. Low, and P. Jaillet, “Gaussian Process Planning with Lipschitz Continuous Reward Functions: Towards Unifying Bayesian Optimization, Active Learning, and Beyond,” in *AAAI*, 2016.
- [20] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA: The MIT Press, 2006.
- [21] S. Boucheron, G. Lugosi, and P. Massart, *Concentration inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.
- [22] B. Schölkopf and A. J. Smola, *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2002.
- [23] G. Wahba, *Spline Models for Observational Data*. Philadelphia, PA: SIAM, 1990.
- [24] N. Ay, J. Jost, H. V. Lê, and L. Schwachhöfer, *Information Geometry*. Springer, 2017.
- [25] A. Christmann and I. Steinwart, *Support Vector Machines*. New York, NY: Springer New York, 2008, ch. 4, pp. 110–163.
- [26] C. McDiarmid, *On the method of bounded differences*, ser. London Mathematical Society Lecture Note Series. Cambridge University Press, 1989, p. 148188.
- [27] G. Pisier, “Subgaussian sequences in probability and Fourier analysis,” Texas A&M University and UPMC-Paris VI, Tech. Rep., 2016.
- [28] H. Bauer, *Probability theory and elements of measure theory*, ser. Probability and mathematical statistics. Academic Press, 1981.
- [29] J. R. Munkres, *Topology: a first course*. Edgewood Cliffs, NJ: Prentice Hall, 1975.
- [30] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge University Press, 1985.