

Bayesian Optimisation for Safe Navigation under Localisation Uncertainty

Rafael Oliveira, Lionel Ott, Vitor Guizilini and Fabio Ramos

Abstract In outdoor environments, mobile robots are required to navigate through terrain with varying characteristics, some of which might significantly affect the integrity of the platform. Ideally, the robot should be able to identify areas that are safe for navigation based on its own percepts about the environment while avoiding damage to itself. Bayesian optimisation (BO) has been successfully applied to the task of learning a model of terrain traversability while guiding the robot through more traversable areas. An issue, however, is that localisation uncertainty can end up guiding the robot to unsafe areas and distort the model being learnt. In this paper, we address this problem and present a novel method that allows BO to consider localisation uncertainty by applying a Gaussian process model for uncertain inputs as a prior. We evaluate the proposed method in simulation and in experiments with a real robot navigating over rough terrain and compare it against standard BO methods.

1 Introduction

Mobile robots have been successfully applied to many field applications, such as mining [9], planetary exploration [8], agriculture [25], and environmental monitoring [10], to name a few. In all these applications, robots face environments with physical characteristics that are *a priori* unknown and can heavily affect performance. In the case of ground robots, terrain roughness can affect the ability of a robot to navigate and even cause damage to its on-board hardware due to excessive vibration [23]. To aid in these problems, methods to enable the robot to automatically learn terrain properties from its sensory data have been presented in the literature [8, 18, 23]. However, such methods usually assume that localisation is accurate enough, without dealing with its inherent uncertainty.

Uncertainty in localisation can mislead learning algorithms with noise-corrupted estimates of the location where measurements are taken. In addition, localisation accuracy also affects navigation, since local path execution heavily depends on knowing where the robot is with respect to a given reference path. As a result, localisation inaccuracy can lead a robot into areas that are unsafe for navigation.

In this paper, we propose a method that allows a robot to actively search for, and learn a model of, traversable areas while keeping itself safe by considering both model and localisation uncertainty. In particular, we propose:

- a framework to account for input uncertainty in Bayesian optimisation [2] for mission planning; and
- an adaptation of the DUCB [10] acquisition function to the context of exploration under uncertain localisation.

In our method, we apply an extension of the conventional Gaussian process (GP) [20] regression models to contexts involving input uncertainty [5] as priors for the BO framework. Such GP models allow BO to take into account noise in both the execution of a query and in the observation’s location estimate. We apply the proposed method to the task of learning a model of terrain roughness from experienced vibration data, as in [23].

The remainder of this paper is organised as follows. In the next section, we review relevant prior work in the areas of terrain modelling and Bayesian optimisation. Section 3 revises the general GP-based BO framework, which does not consider uncertain inputs. In Section 4 we present our method for Bayesian optimisation under localisation uncertainty. Then, in Section 5, we present experimental results in simulation and on a physical robot to evaluate our approach. Finally, in Section 6, we conclude and propose some directions for future work.

2 Related Work

Traversability metrics estimate how hard it is for a vehicle to traverse a certain terrain. In robotics, terrain traversability is usually estimated from either LIDAR range measurements [16] or stereo vision information [6]. Among the different kinds of metrics, terrain roughness can play an important role in planning how a robot should navigate on the terrain [16]. Methods to learn terrain roughness using image data usually rely on learning image classification models [7], which generally do not provide a measure of how uncertain they are about their estimates and are computationally intensive. In [12], however, the authors propose using vehicle experience data to learn a Gaussian process (GP) regression model [20] to predict terrain traversability with uncertainty estimates. Though restrictive when compared to exteroceptive sensing information [6], proprioceptive sensing can still be a very viable option for traversability estimation with robots that are not equipped with stereo vision or 3D LIDAR sensors, or do not possess the computational resources to process this kind of information on-board online.

Souza et al. [23] presented a method to learn a GP model for terrain roughness from vehicle experience. The authors applied Bayesian optimisation (BO) [2] in an active perception approach to reduce experienced vibration during navigation while learning the model from IMU measurements online. In the BO framework, the terrain roughness model is learnt online as the algorithm drives the robot around selecting locations to visit balancing a trade-off between exploration and exploitation [2, 23]. Nevertheless, BO usually considers deterministic query locations within its search space, as BO is typically used in problems where a fixed number of parameters have to be optimised [22, 27].

Recent work [15] presented a method to apply BO to problems where the execution of a query is uncertain, such as robotic grasping. The authors propose querying BO's surrogate model using a Gaussian distribution by applying the unscented transform [26]. In this, even with uncertainty in the execution of the query, the algorithm chooses to sample the objective function at locations where interesting values are more likely to be observed, instead of trying to reach a narrow peak. Despite that, [15] still applies a deterministic-inputs GP model as a prior for BO. In navigation problems, the robot is usually able to obtain a probability distribution estimating its location from a localisation system. Making use of a GP model that takes into account such distributions as inputs should then allow BO to learn a better model of the true underlying objective function.

To deal with uncertain location estimates in observations, besides the usual noisy outputs, BO needs a statistical model that considers partial observations of the inputs when sampling the function being modelled, usually called the *objective* function. In BO the most common approach to modelling the objective function is using Gaussian process (GP) regression [20]. In the case of uncertain input observations, the work in [5] proposed methods to propagate input uncertainty through a GP model, developing analytical approximations to compute covariance function values between inputs represented as Gaussian distributions. Another approach is considered in [13], where the input location estimates are assumed to be corrupted by i.i.d. Gaussian noise, similar to the output observations in standard GP regression. The authors present a method to propagate the input uncertainty to the output of the model using a first-order Taylor approximation. Finally, a more general framework is presented in [4], where the true unobserved inputs of the GP model are considered as latent variables and a variational inference framework is applied to compute the posterior of this GP model under a set of assumptions about the input distributions. None of these methods, however, have been applied to the BO context, where the GP model is built in an online fashion as the algorithm proceeds.

3 Preliminaries

In this section, we review the concepts that form the basis of the method we propose. We start with a general introduction to Gaussian process (GP) regression, which is

used to model terrain properties. Then we follow with a short review of Bayesian optimisation (BO) using GP priors, which will be our planning technique.

3.1 Gaussian process regression

Gaussian process regression [20] is a Bayesian non-parametric framework that places a Gaussian distribution as a prior over the space of functions f , mapping inputs $\mathbf{x} \in \mathbb{R}^d$ to outputs $y \in \mathbb{R}$. In a similar way to a conventional Gaussian distribution, a GP model needs to specify the mean and the covariance for any given pair of values in its vector space, which for a GP is a reproducing kernel Hilbert space (RKHS) [1, 20]. Considering a mean function $\mu_0 : \mathbb{R}^d \rightarrow \mathbb{R}$ and a positive-definite covariance function $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$, a GP prior models the distribution of the function values $\mathbf{y} = [y_1, \dots, y_n]^T$ at a given set of input locations $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^n$ as:

$$\mathbf{y}|\mathbf{X} \sim \mathcal{N}(\mu_0(\mathbf{X}), k(\mathbf{X}, \mathbf{X})), \quad (1)$$

where $\mu_0(\mathbf{X}) = [\mu_0(\mathbf{x}_1), \dots, \mu_0(\mathbf{x}_n)]^T$ and $[k(\mathbf{X}, \mathbf{X})]_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$. In many practical applications the prior mean μ_0 is set to zero or a constant that can be learnt from a dataset. For the covariance function there are a few popular choices, one of them is the squared exponential kernel:

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{x}')^T \mathbf{W}^{-1}(\mathbf{x} - \mathbf{x}')\right), \quad (2)$$

where σ_f^2 is a signal variance parameter and \mathbf{W} is a $d \times d$ diagonal matrix with $[\mathbf{W}]_{ii} = \lambda_i^2$, and each λ_i is a length-scale parameter indicating how much the function can vary along the i -th dimension.

Given a set of observations $\mathcal{D} = \{\mathbf{x}_i, z_i\}$, containing both noise-corrupted outputs $z_i = f(\mathbf{x}_i) + \varepsilon_i$ with $\varepsilon_i \sim \mathcal{N}(0, \sigma_n^2)$ and deterministic input locations \mathbf{x}_i , the values for $y_* = f(\mathbf{x}_*)$ at a given query location \mathbf{x}_* are then distributed according to a Gaussian posterior:

$$y_*|\mathbf{z}, \mathbf{X}, \mathbf{x} \sim \mathcal{N}(\mu(\mathbf{x}_*), \sigma^2(\mathbf{x}_*)), \quad (3)$$

where:

$$\mu(\mathbf{x}_*) = \mu_0(\mathbf{x}_*) + k(\mathbf{x}_*, \mathbf{X})K_{\mathbf{X}}^{-1}(\mathbf{z} - \mu_0(\mathbf{X})) \quad (4)$$

$$\sigma^2(\mathbf{x}_*) = k(\mathbf{x}_*, \mathbf{x}_*) - k(\mathbf{x}_*, \mathbf{X})K_{\mathbf{X}}^{-1}k(\mathbf{X}, \mathbf{x}_*), \quad (5)$$

using $K_{\mathbf{X}} = k(\mathbf{X}, \mathbf{X}) + \sigma_n^2 I$. Therefore, GP models can be used as priors to model functions that are not directly observable during an optimisation process. Partial observations z_i can be collected from the function being optimised to incrementally update the GP model. If the GP prior is appropriately specified, as more observations are collected the variance in the posterior decreases, leading the mean to converge on top of the true f .

3.2 Bayesian optimisation

Consider the problem of searching for the global optimum of a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ in a certain compact region $\mathcal{S} \subset \mathbb{R}^d$, i.e.:

$$\mathbf{x}^* = \operatorname{argmin}_{\mathbf{x} \in \mathcal{S}} f(\mathbf{x}) . \quad (6)$$

Assume that f is unknown to us and we have only access to noisy observations of its output $z = f(\mathbf{x}) + \varepsilon$ with $\varepsilon \sim \mathcal{N}(0, \sigma_n^2)$ and that we can only sample the function N times.

Bayesian optimisation [2] assumes that f is a random variable itself and applies a probability distribution over it using a statistical model. Using a Gaussian process (GP) model, BO makes some prior assumptions about f encoded by the mean $\mu_0(\mathbf{x})$ and covariance function k .

Rather than directly searching over f , BO uses an *acquisition function* $h(\mathbf{x})$ as a guide to sequentially select input locations at which to observe f . The acquisition function uses the information provided by the GP prior and the observations of f to estimate an utility value for sampling f at a given \mathbf{x} . So at each iteration, BO queries the objective f at the location of highest utility according to the acquisition function, \mathbf{x}_t , such that:

$$\mathbf{x}_t = \operatorname{argmax}_{\mathbf{x} \in \mathcal{S}} h(\mathbf{x}) . \quad (7)$$

After observing f at \mathbf{x}_t , BO updates the GP's observations dataset with the new observation $\{\mathbf{x}_t, z_t\}$, which improves its belief about f . After the GP update the algorithm proceeds to the next iteration, choosing an \mathbf{x}_{t+1} . The BO loop runs until a stopping criteria is satisfied, which is usually defined by a maximum budget of iterations. In the end of this process, BO obtains a model that approximates the objective function f and an estimate of its optimum's location \mathbf{x}^* .

One example of acquisition function that can be applied to problems involving robotics navigation is the distance-based upper confidence bound (DUCB) [10]:

$$h_{\text{DUCB}}(\mathbf{x} | \mathcal{D}_{t-1}) = \mu_{t-1}(\mathbf{x}) + \kappa \sigma_{t-1}(\mathbf{x}) - \gamma d(\mathbf{x}_{t-1}, \mathbf{x}) , \quad (8)$$

where $d(\mathbf{x}_{t-1}, \mathbf{x})$ corresponds to the distance between the last sampled location and the candidate \mathbf{x} , μ_{t-1} and σ_{t-1} are given by the GP posterior with the observations in \mathcal{D}_{t-1} , and $\kappa > 0$, $\gamma > 0$ are parameters to be set. In this sense, κ controls the exploration-exploitation trade off, with higher values favouring areas of high uncertainty, while γ penalises large jumps, allowing shorter paths between observations. In addition, notice that for minimisation objectives, as in Equation 6, DUCB can be applied by simply flipping the sign of the GP posterior mean to negative instead.

When the hyper-parameters of the GP model are unknown a priori, a few methods can also be applied to adapt them online as BO runs. One of the simplest, but usually effective, of them is to optimise the log-marginal likelihood of the GP [20] after each observation or after collecting a batch of observations. This optimisation can be done




Fig. 1 Propagation of a Gaussian-distributed input x through a function f sampled from a GP prior. The resulting distribution of output values $f(x)$ is completely different from a Gaussian, since f is highly non-linear. The resulting stochastic process is no longer a Gaussian process.

by performing a few steps of gradient descent on the hyper-parameters after the GP is updated.

4 Bayesian optimisation under localisation uncertainty

In this section we present our framework for Bayesian optimisation under uncertain inputs. By inputs, we refer to points in the objective function f (Equation 6) input space, i.e. any $\mathbf{x} \in \mathbb{R}^d$. Input noise can affect both the execution of a query to observe the objective function and the location estimates of where these observations are taken, as we further explain in Section 4.2. In this sense, we first need to extend BO's surrogate model to incorporate uncertainty in the inputs. We propose to use a Gaussian process regression model with uncertain inputs [5] as a prior, which we then incorporate into BO to obtain a framework for optimisation under uncertain inputs. We start by explaining this GP model before going into details about the proposed BO approach.

4.1 Gaussian process priors with uncertain inputs

Consider that, due to measurement and execution noise, when evaluating a function f at a desired input $\mathbf{x}_* \in \mathbb{R}^d$, the actual input location $\tilde{\mathbf{x}}_* | \mathbf{x}_* \sim \mathbb{P}_*$ where the observation is collected at is not directly observable, but we have access to some estimate of

its probability distribution, here indicated by \mathbb{P}_* . Assuming that $f \sim GP(\mu_0, k)$, the distribution over $\tilde{y}_* = f(\tilde{\mathbf{x}}_*)$ under input noise is no longer Gaussian due to the non-linear relationship between the distribution of $y|\mathbf{x}$, given by Equation 1, and \mathbf{x} . Figure 1 presents an example of what can happen to the output distribution when f is highly non-linear. Therefore, the resulting stochastic process that represents f under random inputs is no longer Gaussian and lacking an analytic formulation [4, 5]. However, as demonstrated in [3], based on the work of [5], we can still recover a Gaussian process approximation for f by using the mean and the covariance of the resulting stochastic process under the influence of input noise. In particular, in the case of a constant deterministic mean function $\mu_0(\mathbf{x}) = m_0$, the mean and the covariance of this noisy process are:

$$\mathbb{E}_{\tilde{\mathbf{x}}_* \sim \mathbb{P}_*} [f(\tilde{\mathbf{x}}_*)] = m_0 \quad (9)$$

$$\text{Cov}_{\mathbb{P}_i, \mathbb{P}_j} [\tilde{y}_i, \tilde{y}_j] = \mathbb{E}_{\mathbb{P}_i, \mathbb{P}_j} [k(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j)] , \quad (10)$$

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, the *expected* covariance function is given by:

$$\mathbb{E}_{\mathbb{P}_i, \mathbb{P}_j} [k(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j)] = \int_{\tilde{\mathbf{x}}_i \in \mathbb{R}^d} \int_{\tilde{\mathbf{x}}_j \in \mathbb{R}^d} k(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j) d\mathbb{P}_i(\tilde{\mathbf{x}}_i) d\mathbb{P}_j(\tilde{\mathbf{x}}_j) = k_p(\mathbb{P}_i, \mathbb{P}_j) . \quad (11)$$

Depending on the type of input distributions and the original kernel for deterministic inputs $k(\mathbf{x}, \mathbf{x}')$, approximate and analytical solutions for Equation 11 may exist [3, 5, 17]. An example is the squared exponential kernel in Equation 2. In this case, if the input distributions are Gaussian, the resulting covariance function, according to [3], is given by:

$$k_p(\mathcal{N}(\mathbf{u}_i, \Sigma_i), \mathcal{N}(\mathbf{u}_j, \Sigma_j)) = \frac{\sigma_f^2 \exp\left(-\frac{1}{2}(\mathbf{u}_i - \mathbf{u}_j)^T (\mathbf{W} + \Sigma_i + \Sigma_j)^{-1} (\mathbf{u}_i - \mathbf{u}_j)\right)}{|I + \mathbf{W}^{-1}(\Sigma_i + \Sigma_j)(1 - \delta_{ij})|^{1/2}} , \quad (12)$$

where σ_f^2 and \mathbf{W} are the same hyper-parameters as described for the standard squared exponential kernel in Equation 2.

The posterior of the stochastic process representing $f \sim GP(\mu_0, k)$ under input noise is not Gaussian due to the complicated forms of Equation 4 and Equation 5 with respect to \mathbf{x} and \mathbf{X} . Yet we can still obtain a suitable approximation [3, 5] for the original f in the noisy input setting by doing inference over a GP with mean m_0 and covariance function k_p , as defined in Equation 11. This approximation is obtained as:

$$\tilde{y}_* | \mathcal{D}, \mathbb{P}_* \sim \mathcal{N}(m(\mathbb{P}_*), v(\mathbb{P}_*)) \quad (13)$$

with

$$m(\mathbb{P}_*) = m_0 + \mathbf{k}_*^T K_{\mathcal{D}}^{-1} (\mathbf{z} - \mathbf{m}_0) \quad (14)$$

$$v(\mathbb{P}_*) = k_p(\mathbb{P}_*, \mathbb{P}_*) - \mathbf{k}_*^T K_{\mathcal{D}}^{-1} \mathbf{k}_* , \quad (15)$$

where \mathbf{m}_0 is an n -dimensional mean vector with the constant m_0 as elements, $\mathbf{k}_* = [k_p(\mathbb{P}_*, \mathbb{P}_1), \dots, k_p(\mathbb{P}_*, \mathbb{P}_n)]^T$, $[K_{\mathcal{D}}]_{ij} = k_p(\mathbb{P}_i, \mathbb{P}_j) + \delta_{ij}\sigma_n^2$, and δ_{ij} denotes the Kronecker delta.

4.2 Bayesian optimisation under localisation uncertainty

To extend BO to the context of uncertain inputs in robotics, we first have to consider how input noise affects the BO process. In robotics problems involving location estimation, we can split input noise into two categories: *localisation* noise and *execution* noise. The first refers to noise affecting the location estimate provided by the robot’s localisation system, usually due to imperfections in motion sensing and in other kinds of sensors, such as GPS devices. The second type, execution noise, is the combined effect of everything affecting the execution of the robot’s path to a given target location, such as localisation noise, uncertain motion dynamics, etc.

In the BO context, execution noise determines the actual location at which an observation will be taken, while localisation noise affects the estimation of that location. Using an uncertain-inputs GP model, as defined in Section 4.1, allows us to take into account both execution and localisation noise in the BO algorithm. For instance, assuming additive zero-mean noise, given a target \mathbf{x} , the actual location where the observation will be taken is:

$$\tilde{\mathbf{x}}^* = \mathbf{x} + \varepsilon_e, \quad (16)$$

where $\varepsilon_e \sim \mathbb{P}_e$ represents execution noise. The query location can then be modelled as a random variable $\tilde{\mathbf{x}}^* \sim \mathbb{P}_{\mathbf{x}}$. In practice, we usually don’t know the execution noise distribution \mathbb{P}_e exactly, and consequently don’t know $\mathbb{P}_{\mathbf{x}}$. However, we can use an approximate *querying* distribution $\hat{\mathbb{P}}_{\mathbf{x}} \approx \mathbb{P}_{\mathbf{x}}$ to account for this uncertainty in the query process. Considering that, at each iteration t , the resulting BO loop should select as target location:

$$\mathbf{x}_t = \operatorname{argmax}_{\mathbf{x} \in \mathcal{S}} h(\hat{\mathbb{P}}_{\mathbf{x}}), \quad (17)$$

which is computed over the uncertain-inputs GP posterior. Under this formulation, $\hat{\mathbb{P}}_{\mathbf{x}}$ provides BO with a measure of how the true location distribution spreads around a given target, due to execution noise. Therefore, any $\hat{\mathbb{P}}_{\mathbf{x}}$ that can upper-bound the variance of $\mathbb{P}_{\mathbf{x}}$ should be sufficient for the algorithm to work.

After observing the objective function, we assume that BO is provided with another distribution estimating the robot’s new location $\tilde{\mathbf{x}}_t \sim \mathbb{P}_l$, given by localisation, such that its mean estimator is:

$$\hat{\mathbf{x}}_t = \mathbb{E}_{\mathbb{P}_l}[\tilde{\mathbf{x}}_t] = \tilde{\mathbf{x}}_t^* + \varepsilon_{l,t}, \quad (18)$$

where $\varepsilon_{l,t} \sim \mathbb{P}_l$ represents localisation noise. Standard BO would utilise $\hat{\mathbf{x}}_t$ as the input to be added with the corresponding outcome z_t into the observations dataset. However, if the localisation noise level is significant or if \mathbb{P}_l is multi-modal, $\hat{\mathbf{x}}_t$ might

not be a good estimator. Fortunately, the GP model in Section 4.1 allows us to use \mathbb{P}_t directly in the observations dataset, mitigating these effects.

As Equation 17 requires, the acquisition function needs to be able to handle probability distributions as inputs by using the uncertain-inputs GP model. For some acquisition functions, such as UCB [24], it can be as straight forward as replacing \mathbf{x} by $\mathbb{P}_{\mathbf{x}}$ and using the corresponding GP posterior mean $m(\mathbb{P}_{\mathbf{x}})$ (Equation 14) and variance $v(\mathbb{P}_{\mathbf{x}})$ (Equation 15). For others, some modifications need to be made. In the case of DUCB, a simple but effective modification to the distance-penalty term is to use the distance between the target location and the mean of the distribution of the last sampled location ($\mathbb{E}_{\mathbb{P}_{t-1}}[\tilde{\mathbf{x}}_{t-1}] = \hat{\mathbf{x}}_{t-1}$), i.e.:

$$h_{\text{DUCB}}(\hat{\mathbb{P}}_{\mathbf{x}}) = m_{t-1}(\hat{\mathbb{P}}_{\mathbf{x}}) + \kappa v_{t-1}^{1/2}(\hat{\mathbb{P}}_{\mathbf{x}}) - \gamma d(\hat{\mathbf{x}}_{t-1}, \mathbf{x}) . \quad (19)$$

5 Experiments

In this section, we present experimental results obtained with our method for uncertain-inputs Bayesian optimisation (UIBO) in the presence of localisation uncertainty. In all cases, we used the DUCB acquisition function [10] to guide the exploration process. We compare our method against two other approaches. The first is standard BO, which does not consider any uncertainty in the location estimates. The second is unscented Bayesian optimisation (UBO) [15], which considers execution noise by means of the unscented transform [26], but assumes that the location estimate of the observation is accurate.

5.1 Simulations

We performed simulations using randomly-generated 2D functions as a model for terrain roughness to be learnt by BO. The task is to find areas of low terrain roughness while staying away from areas of high vibration, which can cause damage to the robot. The learnt GP model should then be more accurate over areas of lower terrain roughness, which are more interesting in practice.

For each test trial, a function is drawn from a Hilbert space [21] with the reproducing kernel defined as the input-noise-free covariance function and combined with a constant mean to keep vibration values positive. In our case, we used the squared exponential (see Equation 2) and its uncertain-inputs equivalent (Equation 12) as covariance functions. To simulate input noise, we sample execution noise from a Gaussian $\mathcal{N}(0, \sigma_e^2 I)$ and localisation noise from another Gaussian $\mathcal{N}(0, \sigma_l^2 I)$. The querying distribution applied by UIBO is set to also be Gaussian, $\hat{\mathbb{P}}_{\mathbf{x}} = \mathcal{N}(\mathbf{x}, \sigma_x^2 I)$, and is the same distribution as applied by unscented BO (UBO). We set $\sigma_e = \sigma_l = 0.07$ and $\sigma_x = 0.1$. Besides input noise, observation noise was set to $\sigma_n = 0.1$. For DUCB's parameters, we set $\kappa = 10$ and $\gamma = 1$ for the three versions of BO, which were man-

ually tuned. The hyper-parameters for the GP models were fixed and identical for each BO method.

In the simulations, we also compared each BO approach using maximum entropy search (ES) as heuristics. These heuristics seek only to reduce the entropy of the corresponding GP posterior, choosing to visit areas of high uncertainty to gain information. ES methods should provide a baseline for comparisons, as they do not consider any estimate of the expected vibration in their search. Following ES guidance, the robot should experience high amounts of vibration and long paths. For fair comparisons, the methods based on deterministic-inputs GP models used the equivalent noise-free version of the uncertain-inputs GP covariance function and the same DUCB parameters, since these do not affect the GP model.


Figure 2 presents results obtained in simulation. As seen in the test case in Figure 2a, standard BO and UBO using DUCB end up having a much more exploratory behaviour due to noise, which leads the robot to execute long paths over the terrain. UIBO, on the other hand, is able to focus its exploration on areas of lower vibration intensity. Overall, Figure 2b shows that UIBO is able to monotonically improve over iterations, finding areas of low vibration for the robot to navigate through. Both the standard and the unscented versions of BO ended up having an average performance much similar to pure entropy search (ES). As entropy-based methods are purely exploratory, there is no improvement in terms of mean experienced vibration, since the algorithm is not optimising for that. In the case of BO-DUCB and UBO-DUCB, the uncertainty in localisation corrupts the location estimates in the observations dataset, and this effect is not taken into account by the GP model. As a result, both methods are misled into areas close to locations that they previously observed, ending up into this more exploratory behaviour, instead of exploiting previous information.

Table 1 presents a summary of the final performance in terms of different metrics for each method. The root mean square error (RMSE) and weighted RMSE (WRMSE) [10] values are computed from the corresponding GP posterior mean (μ) and the ground-truth values of f deterministically queried over a uniform grid covering the entire search space. In our case, the WRMSE emphasises the error in areas of lower vibration, being computed as:


$$WRMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n \left((\mu(\mathbf{x}_i) - f(\mathbf{x}_i)) \frac{\max_j f(\mathbf{x}_j) - f(\mathbf{x}_i)}{\max_j f(\mathbf{x}_j) - \min_j f(\mathbf{x}_j)} \right)^2} \quad (20)$$

This metric is a better performance indicator for this experiment’s task than the plain RMSE, since we are more interested in finding areas of lower vibration.

As seen in Table 1, UIBO is able to outperform the other methods when using DUCB, obtaining a good model, in terms of WRMSE, of the terrain roughness while travelling the shortest distance and experiencing the least amount of vibration. Entropy search using the uncertain-inputs GP was able to obtain the smallest RMSE and WRMSE values, but required a much longer distance to be travelled under the cost of high vibration. In general, in terms of relative experienced vibration, BO and UBO methods with DUCB performed worse due to the high amounts of localisation noise, which led them to behave similar to BO-ES. These results indicate that UIBO



(a) Paths and underlying vibration map



(b) Mean experienced vibration

Fig. 2 Simulation results for standard BO, Unscented BO (UBO) and Uncertain-Inputs BO (UIBO) methods for random functions modelling terrain-induced vibration. (a) presents the true paths taken by each BO method using *DUCB* over the true vibration map in one of the test trials. The markers along the paths indicate the locations where observations were taken. The big "X" mark indicates the starting location, which is the same for the three methods at each trial. The plots in (b) present performance results in terms of mean intensity of the experienced vibration. The results were averaged over 30 trials, each with different maps, and the shaded areas correspond to one standard deviation.

is able to effectively take into account the uncertainty in localisation, keeping the robot safe while exploring an unknown environment.

Method	RMSE	WRMSE	Distance Travelled	Relative Vibration
BO-DUCB	0.59 ± 0.09	0.31 ± 0.05	31.25 ± 1.98	0.95 ± 0.04
UBO-DUCB	0.64 ± 0.12	0.34 ± 0.08	35.49 ± 1.23	0.95 ± 0.04
UIBO-DUCB	0.59 ± 0.15	0.26 ± 0.08	16.12 ± 2.04	0.61 ± 0.17
BO-ES	0.58 ± 0.10	0.32 ± 0.07	57.07 ± 2.73	1.00
UBO-ES	0.64 ± 0.12	0.34 ± 0.08	67.15 ± 2.56	1.00 ± 0.03
UIBO-ES	0.38 ± 0.05	0.22 ± 0.05	62.85 ± 2.06	0.99 ± 0.03

Table 1 Simulation results: average performance comparisons for each BO approach under different metrics with the corresponding standard deviation.

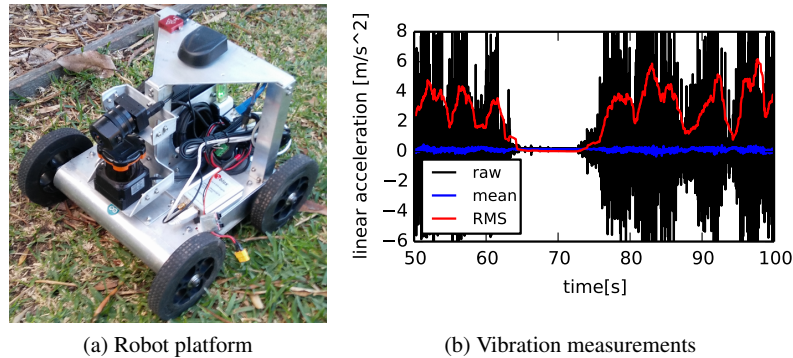


Fig. 3 Robot experiment details. (a) presents our experimental platform. (b) compares different methods to measure vibration using an IMU: *raw* corresponds to the raw vertical acceleration readings (discounting gravity), *mean* corresponds to the mean value of the raw measurements in a moving window of size 200, and *RMS* corresponds to the root-mean-square value of the raw measurements in the same window.

5.2 Experiments with a real robot

We performed experiments with a physical robot outdoors to test the performance of the proposed uncertain-inputs BO approach against standard BO with DUCB as in [23]. The purpose of this experiment is to highlight the differences in performance and in behaviour between both BO approaches when faced with localisation uncertainty in a real-world scenario, where the robot is tasked with learning terrain roughness.

5.2.1 Experimental setup

Our test platform was a small four-wheeled skid-steer robot, depicted in Figure 3a. The robot is equipped with an on-board computer running ROS ¹. The tests were performed in an area with terrain covered by grass and with some portions of harder ground exposed. The goal of the robot is to find areas of low terrain roughness, while avoiding areas that induce excessive vibration to the platform. In this scenario, high amounts of vibration can cause damage to the robot or affect its ability to perceive the environment.

Our robot was equipped with an IMU sensor, placed on top of its chassis. With the robot moving, vibration shows up mostly as linear acceleration along the robot's vertical axis, bouncing the robot up and down to the ground. We considered different methods to measure vibration from a fixed moving window of acceleration measurements. In particular, we considered the mean and the root mean square values. As we want the robot to avoid areas of high vibration, the RMS value demonstrated itself to

¹ The Robot Operating System: www.ros.org

be a better estimator, since it grows with the amplitude of the vibration, as pictured in Figure 3b. The mean, on the other hand, should be always around zero, as each upwards acceleration is immediately compensated by a downwards drop of the robot.

As observations for the BO algorithms, the RMS vibration estimates were computed over a sliding window of 100 IMU measurements and updated at a rate of 2 Hz as the robot is driving. To reduce the effect of different driving speeds on the readings, the vibration estimates were posted when the robot was driving at speeds between 0.4 and 0.7 m/s, where 0.7 m/s was the maximum speed allowed for the path following control. These observations were combined with location estimates from an extended Kalman filter (EKF) [14], which was configured to fuse wheel odometry, IMU and GPS estimates.

We ran both plain BO and UIBO for a fixed budget of 30 iterations, i.e. each algorithm was allowed to choose 30 target locations to take the robot to. At each iteration, the robot drove autonomously attempting to follow straight paths from its previous goal to the next goal, as given by the BO planner. The iteration was signalled as finished whenever the robot arrived within a given radius from the goal. After that, the GP model is updated, the hyper-parameters are re-learned by maximising the GP log-marginal likelihood using COBYLA [19], and the planner selects a new target. As the sample rate of observations is relatively low, both BO methods were able to repeatedly execute this process online in close to real time. The DUCB acquisition function was configured with an uncertainty factor $\kappa = 10$ and a distance factor $\gamma = 0.5$. UIBO applied as query distribution an isotropic Gaussian with variance $\sigma_x^2 = 4$ for each coordinate.

5.2.2 Performance

Figure 4 presents the GP model obtained by each BO method and their paths, according to noisy EKF location estimates. For both methods, the robot was placed at an initial location at the bottom left corner of the search space. As Figure 4 shows, UIBO is able to concentrate its search over areas of low vibration. On the other hand, standard BO ends up performing too much exploration passing more often over areas that cause excessive vibration to the robot.

For performance comparisons, we collected a set of vibration measurements during the experiments using a real time kinematic (RTK) GPS ² device, which was employed to obtain high-precision location estimates. These measurements were only collected for the purposes of validation, and not passed on to the BO planners, which relied solely on a conventional GPS device fused into the EKF estimates. The uncertainty of the EKF location estimates varied around 3 metres, while the RTK GPS device is able to provide centimetre accuracy. The estimates of a GP model built directly from the validation data is shown in Figure 5. This model is presented for visualisation purposes and was not used to compute performance metrics, which relied only on the raw validation data. The model's hyper-parameters were determined

² <https://emlid.com/reach/>

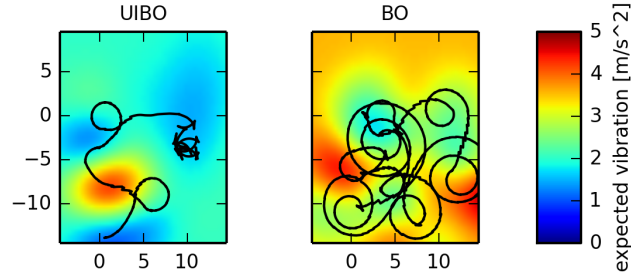


Fig. 4 Posterior mean of the GP model built by each BO method overlaid with their respective paths according to locations estimated by an EKF fusing conventional GPS, IMU and odometry.

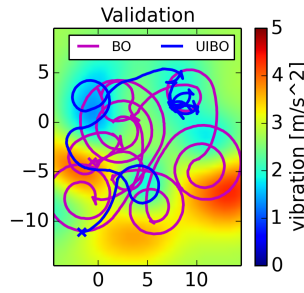


Fig. 5 For visualisation purposes, posterior mean of the GP model built from the validation data overlaid with the true paths taken by each method, according to RTK GPS location estimates.

Method	WRMSE [m/s^2]	RMSE [m/s^2]	Distance Travelled [m]	Mean Vibration [m/s^2]
BO	1.04	1.31	247	3.0
UIBO	0.83	1.33	92	2.1

Table 2 Field results: performance comparisons for each BO approach under different metrics. The WRMSE was computed between the posterior mean of the final GP model at the locations in the validation dataset and the corresponding vibration measurements at those locations.

by maximising its log-marginal likelihood as $\{m_0, l, \sigma_f, \sigma_n\} = \{2.7, 2.8, 0.8, 1.0\}$, where l is the kernel length-scale.

Table 2 presents a performance summary. The WRMSE value for this experiment was computed between the posterior mean from each method’s final GP model, built from noisy EKF location estimates, and the raw vibration measurements present in the validation dataset, collected using precise RTK GPS location estimates. We can see that, for this experiment, UIBO is able to outperform standard BO by producing a better model over areas of low roughness while experiencing less vibration and following a much shorter path, confirming facts previously observed in simulations.

6 Conclusion

In this paper, we presented a new method to Bayesian optimisation for active learning problems in robotics in the cases where uncertainty in the location estimates is significant. The proposed method provides a principled way to consider this uncertainty in the inputs of a Gaussian process model, which is applied as a prior by the BO algorithm. Execution noise is also considered by means of a query distribution when optimising BO's acquisition function. The method was proved to outperform other BO approaches in simulation and in an experiment with a physical robot. Therefore, the proposed method can be applied to problems where we are interested in mapping the traversability of a terrain, but need to keep the robot safe in the midst of localisation uncertainty.

Some topics were not addressed by this paper but are worthy of future research work. One of them is the estimation of execution noise, which could be done in an online way by other statistical methods, for example, using maximum likelihood estimates for the parameters of a given distribution. Another topic is the extension of the DUCB-based exploration to informative continuous path planning [11] or to trajectory optimisation in goal-directed navigation [18]. For that, however, a method to propagate execution noise through a candidate path would have to be developed taking into account hard-to-model factors from both the robot and the environment, such as stochastic motion dynamics and imperfections in sensor noise.

Acknowledgements This work was partly funded by CAPES, Brazil, scholarship BEX 13224/13-1, and by Data61/CSIRO, Australia.

References

1. Aronszajn, N.: Theory of Reproducing Kernels. *Transactions of the American Mathematical Society* 68(3), 337–404 (1950)
2. Brochu, E., Cora, V.M., de Freitas, N.: A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical Reinforcement Learning. Tech. rep., University of British Columbia (2010)
3. Dallaire, P., Besse, C., Chaib-Draa, B.: An approximate inference with Gaussian process to latent functions from uncertain data. *Neurocomputing* 74, 1945–1955 (2011)
4. Damianou, A.C., Titsias, M.K., Lawrence, N.D.: Variational Inference for Latent Variables and Uncertain Inputs in Gaussian Processes. *Journal of Machine Learning Research* 17(1), 1–62 (2016)
5. Girard, A.: Approximate methods for propagation of uncertainty with Gaussian process models. Ph. d, University of Glasgow (2004)
6. Ho, K., Peynot, T., Sukkarieh, S.: Traversability estimation for a planetary rover via experimental kernel learning in a Gaussian process framework. In: *IEEE International Conference on Robotics and Automation (ICRA)*. pp. 3475–3482. IEEE, Karlsruhe, Germany (2013)
7. Komma, P., Weiss, C., Zell, A.: Adaptive bayesian filtering for vibration-based terrain classification. In: *Proceedings - IEEE International Conference on Robotics and Automation*. pp. 3307–3313. IEEE, Kobe, Japan (2009)

8. Lang, T., Plagemann, C., Burgard, W.: Adaptive non-stationary kernel regression for terrain modelling. In: Burgard, W., Brock, O., Stachniss, C. (eds.) Proc. of the Robotics: Science and Systems Conference (RSS). Atlanta, GA (2007)
9. Maekawa, T., Noda, T., Tamura, S., Ozaki, T., Machida, K.i.: Curvature continuous path generation for autonomous vehicle using B-spline curves. *Computer-Aided Design* 42(4), 350–359 (2010)
10. Marchant, R., Ramos, F.: Bayesian Optimisation for Intelligent Environmental Monitoring. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE (Oct 2012)
11. Marchant, R., Ramos, F.: Bayesian Optimisation for Informative Continuous Path Planning. In: IEEE International Conference on Robotics and Automation (ICRA). pp. 6136–6143 (2014)
12. Martin, S., Murphy, L., Corke, P.: Building Large Scale Traversability Maps Using Vehicle Experience. In: The 13th International Symposium on Experimental Robotics (ISER). vol. 88, pp. 891–905. Springer (2013)
13. Mchutchon, A., Rasmussen, C.E.: Gaussian Process Training with Input Noise. In: Advances in Neural Information Processing Systems. pp. 1341–1349 (2011)
14. Moore, T., Stouch, D.: A generalized extended kalman filter implementation for the robot operating system. In: Proceedings of the 13th International Conference on Intelligent Autonomous Systems (IAS-13). Springer (July 2014)
15. Nogueira, J., Martinez-Cantin, R., Bernardino, A., Jamone, L.: Unscented Bayesian Optimization for Safe Robot Grasping. In: IEEE International Conference on Robotics and Automation (ICRA). pp. 1967–1972. Daejeon, Korea (2016)
16. Nordin, P.: Mobile Robot Traversability Mapping. Licentiate thesis, Linköping University (2012)
17. O’Callaghan, S.T., Ramos, F.T.: Gaussian process occupancy maps. *The International Journal of Robotics Research (IJRR)* 31(1), 42–62 (2012)
18. Oliveira, R., Ott, L., Ramos, F.: Active Perception for Modelling Energy Consumption in Off-Road Navigation. In: Australasian Conference on Robotics and Automation (ACRA). Brisbane, QLD, Australia (2016)
19. Powell, M.: A view of algorithms for optimization without derivatives. Tech. rep., Cambridge University DAMTP, Cambridge, United Kingdom (2007)
20. Rasmussen, C.E., Williams, C.K.I.: Gaussian Processes for Machine Learning. The MIT Press, Cambridge, MA (2006)
21. Schölkopf, B., Smola, A.J.: Learning with kernels: support vector machines, regularization, optimization, and beyond. MIT press (2002)
22. Snoek, J., Larochelle, H., Adams, R.P.: Practical bayesian optimization of machine learning algorithms. In: Pereira, F., Burges, C.J.C., Bottou, L., Weinberger, K.Q. (eds.) Advances in Neural Information Processing Systems 25, pp. 2951–2959. Curran Associates, Inc. (2012)
23. Souza, J.R., Marchant, R., Ott, L., Wolf, D.F., Ramos, F.: Bayesian Optimisation for Active Perception and Smooth Navigation. In: IEEE International Conference on Robotics and Automation (ICRA) (2014)
24. Srinivas, N., Krause, A., Kakade, S.M., Seeger, M.W.: Information-Theoretic Regret Bounds for Gaussian Process Optimization in the Bandit Setting. *IEEE Transactions on Information Theory* 58(5), 1–16 (2012)
25. Underwood, J., Wendel, A., Schofield, B., McMurray, L., Kimber, R.: Efficient in-field plant phenomics for row-crops with an autonomous ground vehicle. *Journal of Field Robotics* (2017)
26. Wan, E.A., van der Merwe, R.: The Unscented Kalman Filter for Nonlinear Estimation. In: Adaptive Systems for Signal Processing, Communications, and Control Symposium (AS-SPCC). pp. 153–158 (2000)
27. Wilson, A., Fern, A., Tadepalli, P.: Using Trajectory Data to Improve Bayesian Optimization for Reinforcement Learning. *Journal of Machine Learning Research* 15, 253–282 (2014)