
Noisy-Input Entropy Search for Efficient Robust Bayesian Optimization

Lukas P. Fröhlich^{1,2}

Edgar D. Klenske¹

Julia Vinogradska¹

Christian Daniel¹

¹Bosch Center for Artificial Intelligence
Renningen, Germany

Melanie N. Zeilinger²

²ETH Zürich
Zürich, Switzerland

Abstract

We consider the problem of robust optimization within the well-established Bayesian optimization (BO) framework. While BO is intrinsically robust to noisy evaluations of the objective function, standard approaches do not consider the case of uncertainty about the input parameters. In this paper, we propose *Noisy-Input Entropy Search (NES)*, a novel information-theoretic acquisition function that is designed to find robust optima for problems with both input and measurement noise. NES is based on the key insight that the robust objective in many cases can be modeled as a Gaussian process, however, it cannot be observed directly. We evaluate NES on several benchmark problems from the optimization literature and from engineering. The results show that NES reliably finds robust optima, outperforming existing methods from the literature on all benchmarks.

1 Introduction

Bayesian optimization (BO) is a well-established technique for optimization of black-box functions with applications in a wide range of domains (Brochu et al., 2010; Shahriari et al., 2016). The two key benefits of BO are its sample-efficiency and its intrinsic robustness to noisy function evaluations, rendering it particularly powerful when function evaluations are either time consuming or costly, e.g., for drug design or (robot) controller tuning (Calandra et al., 2016; Cully et al.,

2015; Griffiths and Hernández-Lobato, 2017). The sample-efficiency of BO stems from two key ingredients: (i) a Bayesian surrogate model that approximates the objective function based on previous evaluations, e.g., Gaussian process (GP) regression, and (ii) an acquisition function that defines the next evaluation point based on the surrogate model. Several acquisition functions have been proposed that heuristically trade off between exploration and exploitation (Kushner, 1964; Moćkus, 1975; Cox and John, 1992). More recently, the family of entropy search acquisition functions has been introduced. These acquisition functions use an information-theoretic approach and choose the next evaluation point to maximize information about the global optimum. At the cost of computational complexity, entropy-based methods are generally more sample-efficient than other acquisition functions (Henning and Schuler, 2012; Hernández-Lobato et al., 2014; Wang and Jegelka, 2017).

In addition to sample efficiency, robustness with respect to model uncertainties or perturbations on the input is critical in many applications (see, e.g., (Beyer and Sendhoff, 2007) for a survey). Examples are numerous in fields such as control (Başar and Bernhard, 2008), design engineering (Chen et al., 1996) and operations research (Adida and Perakis, 2006). In its standard formulation, BO is intrinsically robust with respect to noisy function evaluations, however, it leads to sub-optimal solutions in the presence of perturbations on the input. While robust optimization has been considered in the context of BO before, previous work is based on heuristic acquisition functions. To the best of our knowledge, entropy-based acquisition functions for robust optimization problems that fully leverage the potential of BO have not been addressed to date.

bilistic formulation of robustness where the parameters found during optimization are randomly perturbed at the implementation stage, i.e., broad optima are preferable over narrow ones. Due to their sample-efficiency, we build on entropy-based acquisition functions and propose to choose the next evaluation point in order to maximize the information about the robust optimum. Our method is based on the key insight that the robust objective can be modeled with a GP just as in the standard BO setting. However, the robust objective is not directly observable, but needs to be constructed from (noisy) evaluations of the original function without perturbations on the input. We base our framework on the recently proposed max-value entropy search (MES) (Wang and Jegelka, 2017), due to the low computational demand. The resulting formulation requires knowledge of the GP’s predictive distribution conditioned on the robust maximum value, which is an analytically intractable distribution. We propose two methods to approximate this distribution (i) based on rejection sampling, which in the limit of infinitely many samples is exact but computationally expensive, and (ii) based on expectation propagation (EP) (Minka, 2001), which is computationally more efficient. We evaluate the proposed acquisition function on a wide range of benchmark problems and compare against related approaches from the literature. Moreover, we apply the proposed method to a simulated aerospace task to demonstrate the importance of robust black-box optimization in practice.

Related Work Closely related to our method are the approaches presented by Nogueira et al. (2016) and Beland and Nair (2017), both of which consider the same probabilistic robust objective as considered in this paper (see Eq. (1)). Nogueira et al. (2016) proposed to use the expectation of the expected improvement (EI) acquisition function with respect to the input noise. The expectation is approximated using the unscented transformation (Julier and Uhlmann, 2004), which is computationally efficient, but the approximation accuracy strongly depends on the choice of hyperparameters. In the paper by Beland and Nair (2017), the robust objective is also modeled as a GP, however, it is implicitly assumed that the robust objective can be observed directly. In contrast to the two aforementioned methods, our method uses an information-theoretic approach. We compare our method to both Nogueira et al. (2016) and Beland and Nair (2017).

Besides random perturbations on the optimization parameters, other robust optimization settings have been investigated in the context of BO. In recent work, Bogunovic et al. (2018) consider the worst-case perturbation within a given set (or minimax setting) instead of random perturbations. Moreover, the authors provide

rigorous convergence guarantees for their acquisition function, based on the results from Srinivas et al. (2010). Chen et al. (2017) consider a finite set of non-convex objective functions and seek the maximizer that is robust with respect to the choice of objective function from the given set. In the setting considered by Martinez-Cantin et al. (2018), some evaluations are corrupted such that their value is perturbed much stronger than the observation noise, thus biasing the surrogate model. However, this setting does not extend to the case of perturbations on the input. Groot et al. (2010); Tesch et al. (2011); Toscano-Palmerin and Frazier (2018) assume that the objective function depends on two types of input parameters: the control parameters to be optimized and environmental parameters against which the maximizer should be robust. This differs from our setting, in which we aim at finding an optimum that is robust with respect to the control parameters.

Similar to Nogueira et al. (2016); Beland and Nair (2017); Bogunovic et al. (2018), we assume exact knowledge of the control parameters during the optimization and require robustness when deploying the optimal parameters. In contrast, Oliveira et al. (2019) proposed a method that deals with uncertain inputs during the optimization process, however, their goal is to find the global optimum instead of the robust optimum.

2 Preliminaries

In this section, we briefly review Bayesian optimization (BO) and discuss how it relates to the robust optimization setting considered in this paper. As the robust objective will be approximated with GP regression, we furthermore summarize how perturbations on the input parameters can be included in the posterior predictive distribution.

Bayesian Optimization In BO we seek the maximizer of the unknown objective function $f(\mathbf{x}) : \mathcal{X} \rightarrow \mathbb{R}$ over a compact set $\mathcal{X} \subseteq \mathbb{R}^d$ despite only having access to noisy observations, $y_i = f(\mathbf{x}_i) + \epsilon$ with $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$. Furthermore, no gradient information is available and each evaluation of $f(\mathbf{x})$ takes a considerable amount of time or effort. Thus, the goal is to find the maximum in as few evaluations as possible. The core idea of BO is to model the unknown objective $f(\mathbf{x})$ with a Bayesian surrogate model based on past observations $\mathcal{D}_n = \{(\mathbf{x}_i, y_i)\}_{i=1:n}$. Common choices for the model are Bayesian neural networks (Snoek et al., 2015) or GPs (Rasmussen and Williams, 2006). In this paper, we consider the latter. Based on the surrogate model, the next query point is chosen by maximizing a so-called acquisition function $\alpha(\mathbf{x})$.

Acquisition functions quantify the exploration-

exploitation trade-off between regions with large predicted values (exploitation) and regions of high uncertainty (exploration). Entropy-based acquisition functions address this trade-off by minimizing the uncertainty of the belief about global optimum’s location $p(\mathbf{x}^*|\mathcal{D}_n)$ (Hennig and Schuler, 2012). The next evaluation point \mathbf{x}_{n+1} is chosen to maximize the mutual information between the global optimum \mathbf{x}^* and the next evaluation point, given by $I((\mathbf{x}, y); \mathbf{x}^*|\mathcal{D}_n)$. Recently, Wang and Jegelka (2017) introduced the MES acquisition function, which considers the optimum’s value y^* instead of its location, i.e., $\alpha_{\text{MES}}(\mathbf{x}) = I((\mathbf{x}, y); y^*|\mathcal{D}_n)$. This formulation significantly reduces the computational burden compared to its predecessors.

By design, BO is able to efficiently optimize non-convex black-box functions. However, it is generally not able to find optima that are robust with respect to perturbations of the input parameters.

Robust Bayesian Optimization In this paper, we consider a probabilistic formulation of robustness, i.e., we assume that the optimization parameters are randomly perturbed at implementation stage. In the presence of input noise, broad optima should be preferred over narrow ones. Thus, instead of optimizing $f(\mathbf{x})$ directly, we aim at maximizing the *robust objective*,

$$g(\mathbf{x}) = \mathbb{E}_{\boldsymbol{\xi} \sim p(\boldsymbol{\xi})} [f(\mathbf{x} + \boldsymbol{\xi})] = \int f(\mathbf{x} + \boldsymbol{\xi})p(\boldsymbol{\xi})d\boldsymbol{\xi}, \quad (1)$$

such that the robust optimizer is given by $\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathcal{X}} g(\mathbf{x})$. The random perturbations acting on the input parameters \mathbf{x} are characterized by the distribution $p(\boldsymbol{\xi})$. To this end, we assume $p(\boldsymbol{\xi}) \sim \mathcal{N}(0, \boldsymbol{\Sigma}_x)$ with $\boldsymbol{\Sigma}_x = \text{diag}[\sigma_{x,1}^2, \dots, \sigma_{x,d}^2]$ and $\sigma_{x,i}^2$ to be known for all i . Other choices are of course possible, e.g., $p(\boldsymbol{\xi})$ could be chosen as a uniform distribution. Note that for vanishing input noise, $\sigma_{x,i} \rightarrow 0$ for all i , the distribution $p(\boldsymbol{\xi})$ converges to the Dirac delta distribution and we obtain the standard, non-robust optimization setting.

Gaussian Process Regression Gaussian process (GP) regression is a non-parametric method to model an unknown function $f(\mathbf{x}) : \mathcal{X} \mapsto \mathbb{R}$ from data \mathcal{D}_n (see, e.g., (Rasmussen and Williams, 2006)). A GP defines a prior distribution over functions, such that any finite number of function values are normally distributed with mean $\mu_f(\mathbf{x})$ and covariance specified by the kernel function $k_f(\mathbf{x}, \mathbf{x}')$ for any $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$ (w.l.o.g. we assume $\mu_f(\mathbf{x}) \equiv 0$). Conditioning the prior distribution on observed data \mathcal{D}_n leads to the posterior predictive mean and variance,

$$\begin{aligned} m_f(\mathbf{x}|\mathcal{D}_n) &= \mathbf{k}_f(\mathbf{x})^\top \mathbf{K}^{-1} \mathbf{y}, \\ v_f(\mathbf{x}|\mathcal{D}_n) &= k_f(\mathbf{x}, \mathbf{x}) - \mathbf{k}_f(\mathbf{x})^\top \mathbf{K}^{-1} \mathbf{k}_f(\mathbf{x}), \end{aligned} \quad (2)$$

at any $\mathbf{x} \in \mathcal{X}$ with $[\mathbf{k}_f(\mathbf{x})]_i = k_f(\mathbf{x}, \mathbf{x}_i)$, $[\mathbf{K}]_{ij} = k_f(\mathbf{x}_i, \mathbf{x}_j) + \delta_{ij}\sigma_\epsilon^2$, $[\mathbf{y}]_i = y_i$ and δ_{ij} denotes the Kronecker delta.

In the context of Bayesian optimization, GP regression is commonly used as a surrogate model for the objective $f(\mathbf{x})$. Since the expectation is a linear operator and GPs are closed under linear operations (Rasmussen and Williams, 2006), the robust objective $g(\mathbf{x})$ can be modeled as a GP as well, based on noisy observations of $f(\mathbf{x})$. The predictive distribution for the robust objective then becomes

$$\begin{aligned} m_g(\mathbf{x}|\mathcal{D}_n) &= \mathbf{k}_{gf}(\mathbf{x})^\top \mathbf{K}^{-1} \mathbf{y}, \\ v_g(\mathbf{x}|\mathcal{D}_n) &= k_g(\mathbf{x}, \mathbf{x}) - \mathbf{k}_{gf}(\mathbf{x})^\top \mathbf{K}^{-1} \mathbf{k}_{fg}(\mathbf{x}), \end{aligned} \quad (3)$$

where the respective kernel functions are given by $k_g(\mathbf{x}, \mathbf{x}') = \iint k_f(\mathbf{x} + \boldsymbol{\xi}, \mathbf{x}' + \boldsymbol{\xi}')p(\boldsymbol{\xi})p(\boldsymbol{\xi}')d\boldsymbol{\xi}d\boldsymbol{\xi}'$ and $k_{gf}(\mathbf{x}, \mathbf{x}') = \int k_f(\mathbf{x} + \boldsymbol{\xi}, \mathbf{x}')p(\boldsymbol{\xi})d\boldsymbol{\xi}$. For the well-known squared exponential and Matérn kernel functions, k_{gf} and k_g can be computed in closed-form for normally and uniformly distributed input noise $\boldsymbol{\xi}$ (see, e.g., (Dallaire et al., 2009)).

3 Noisy-Input Entropy Search

In this section, we elaborate on the main contribution of this paper. We first present our robust acquisition function and give an overview of the challenges associated with the proposed approach. The main challenge is that the robust formulation requires the GP’s predictive distribution conditioned on the robust maximum value, which is analytically intractable. We propose two approximation schemes: The first is based on rejection sampling (RS), which gives the exact result in the limit of infinitely many samples, but is computationally challenging. The second approach is based on expectation propagation (EP) (Minka, 2001) and is computationally more efficient, albeit not unbiased.

As discussed in Sec. 2, entropy-based acquisition functions quantify the information gain about the global optimum of $f(\mathbf{x})$. Hence, the next evaluation point \mathbf{x}_{n+1} is selected to be maximally informative about \mathbf{x}^* (or y^* for MES). For robust optimization, we aim at finding the maximizer of the robust objective $g(\mathbf{x})$ instead. We build on the work of Wang and Jegelka (2017) and consider the mutual information between \mathbf{x} and the objective’s maximum value. Consequently, we maximize the information about the robust maximum value $g^* = \max_{\mathbf{x} \in \mathcal{X}} g(\mathbf{x})$ and propose the *Noisy-Input Entropy Search (NES)* acquisition function

$$\begin{aligned} \alpha_{\text{NES}}(\mathbf{x}) &= I((\mathbf{x}, y); g^*|\mathcal{D}_n) \\ &= H[p(y(\mathbf{x})|\mathcal{D}_n)] - \mathbb{E}_{g^*|\mathcal{D}_n} \left[H[p(y(\mathbf{x})|\mathcal{D}_n, g^*)] \right], \end{aligned} \quad (4)$$

where $I(\cdot; \cdot | \cdot)$ denotes the conditional mutual information and $H[\cdot]$ the differential entropy. Note how NES reasons about g^* while only (noisily) observing $f(\mathbf{x})$ as opposed to the naïve approach of applying maximum entropy search (MES) to the GP model of the robust objective, which assumes access to observations of $g(\mathbf{x})$. The corresponding mutual information would be $I((\mathbf{x}, z); g^* | \mathcal{D}_n)$, with the hypothetical observation model $z = g(\mathbf{x}) + \eta$ and $\eta \sim \mathcal{N}(0, \sigma_\eta^2)$. This, however, is not possible as $g(\mathbf{x})$ cannot be observed directly.

The first term in Eq. (4) corresponds to the entropy of the GP’s predictive posterior distribution. For Gaussian distributions, the entropy can be computed analytically such that $H[p(y(\mathbf{x}) | \mathcal{D}_n)] = 0.5 \log[2\pi e(v_f(\mathbf{x} | \mathcal{D}_n) + \sigma_\epsilon^2)]$. The second term in Eq. (4) has no analytic solution and requires approximations for the following reasons: (i) The expectation is with respect to the unknown distribution over g^* and (ii) it is not obvious how conditioning on the robust maximum value g^* influences the predictive distribution $p(y(\mathbf{x}) | \mathcal{D}_n)$. In what follows we will address these challenges.

3.1 Approximating the Expectation Over Robust Maximum Values

The belief over the robust maximum value $p(g^* | \mathcal{D}_n)$ in Eq. (4) cannot be computed in closed form. In the standard BO setting, the corresponding expectation has been approximated via Monte Carlo sampling (Hernández-Lobato et al., 2014; Wang and Jegelka, 2017) We follow this approach and approximate the expectation over $p(g^* | \mathcal{D}_n)$ as

$$\mathbb{E}_{g^* | \mathcal{D}_n} \left[H[p(y(\mathbf{x}) | \mathcal{D}_n, g^*)] \right] \approx \frac{1}{K} \sum_{g_k^* \in G^*} H[p(y(\mathbf{x}) | \mathcal{D}_n, g_k^*)], \quad (5)$$

where G^* is a set of K samples drawn from $p(g^* | \mathcal{D}_n)$. We generate samples $g_k^* \in G^*$ via a two-step process: (i) sample a function $\tilde{g}_k(\mathbf{x})$ from $p(g(\mathbf{x}) | \mathcal{D}_n)$ and (ii) maximize it such that $g_k^* = \max_{\mathbf{x} \in \mathcal{X}} \tilde{g}_k(\mathbf{x})$.

For efficient function sampling from $p(g(\mathbf{x}) | \mathcal{D}_n)$ and subsequent maximization, we employ the sparse spectrum Gaussian process (SSGP) approximation (Lázaro-Gredilla et al., 2010). The advantage of SSGPs is that the sampled functions can be efficiently optimized with a gradient-based optimizer. In this case, we can sample functions from $p(f(\mathbf{x}) | \mathcal{D}_n)$ that are of the form $\tilde{f}_k(\mathbf{x}) = \mathbf{a}^T \boldsymbol{\phi}_f(\mathbf{x})$, where $\boldsymbol{\phi}_f(\mathbf{x}) \in \mathbb{R}^M$ is a vector of random feature functions. The components of the feature vector $\boldsymbol{\phi}_f(\mathbf{x}) \in \mathbb{R}^M$ are given by $\phi_{f,i}(\mathbf{x}) = \cos(\mathbf{w}_i^T \mathbf{x} + b_i)$, with $b_i \sim \mathcal{U}(0, 2\pi)$ and $\mathbf{w}_i \sim p(\mathbf{w}) \propto s(\mathbf{w})$ where $s(\mathbf{w})$ is the Fourier dual of the kernel function k_f . The weight vector \mathbf{a} is distributed according to $\mathcal{N}(\mathbf{A}^{-1} \boldsymbol{\Phi}_f^T \mathbf{y}, \sigma_\epsilon^2 \mathbf{A}^{-1})$

with $\mathbf{A} = \boldsymbol{\Phi}_f^T \boldsymbol{\Phi}_f + \sigma_\epsilon^2 \mathbf{I}$, $\boldsymbol{\Phi}_f^T = [\phi_f(\mathbf{x}_1), \dots, \phi_f(\mathbf{x}_n)]$, $\mathcal{D}_n = \{(\mathbf{x}_i, y_i)\}_{i=1:n}$ and $\mathbf{y} = [y_1, \dots, y_n]$ (see, e.g., (Lázaro-Gredilla et al., 2010) or (Hernández-Lobato et al., 2014) for details).

We can now generate $\tilde{g}_k(\mathbf{x})$ from a function sample $\tilde{f}_k(\mathbf{x})$ by taking the expectation w.r.t. the input noise. As each $\tilde{f}_k(\mathbf{x})$ is a linear combination of M cosine functions, we can compute this expectation in closed form. For normally distributed input noise, $\boldsymbol{\xi} \sim \mathcal{N}(0, \boldsymbol{\Sigma}_x)$ with $\boldsymbol{\Sigma}_x = \text{diag}[\sigma_{x,1}^2, \dots, \sigma_{x,d}^2]$, this operation reduces to a scaling of the feature functions,

$$\begin{aligned} \phi_{g,i}(\mathbf{x}) &= \int \phi_{f,i}(\mathbf{x} + \boldsymbol{\xi}) p(\boldsymbol{\xi}) d\boldsymbol{\xi} \\ &= \phi_{f,i}(\mathbf{x}) \exp\left(-\frac{1}{2} \sum_{j=1}^d \mathbf{w}_{i,j}^2 \sigma_{x,j}^2\right). \end{aligned} \quad (6)$$

Thus, we can efficiently sample $g_k^* \sim p(g^* | \mathcal{D}_n)$ exploiting the fact that $g_k^* = \max_{\mathbf{x} \in \mathcal{X}} \tilde{g}_k(\mathbf{x})$ and $\tilde{g}_k(\mathbf{x}) = \mathbf{a}^T \boldsymbol{\phi}_g(\mathbf{x})$. We present a detailed derivation of Eq. (6) as well as a discussion on the number of samples needed for a sufficient approximation accuracy of Eq. (5) in the supplementary material.

3.2 Approximating the Conditional Predictive Distribution

In the previous section we discussed how to sample robust maximum values $g_k^* \sim p(g^* | \mathcal{D}_n)$. To evaluate the proposed acquisition function $\alpha_{\text{NES}}(\mathbf{x})$, we need to compute the entropy of the predictive distribution conditioned on a sampled robust maximum value, i.e., $H[p(y(\mathbf{x}) | \mathcal{D}_n, g_k^*)]$. Conditioning on g_k^* imposes $g(\mathbf{x}) \leq g_k^*$, which renders the computation of $p(y(\mathbf{x}) | \mathcal{D}_n, g_k^*)$ intractable. In this section, we propose two approximation schemes (i) based on rejection sampling (RS) which is exact in the limit of infinite samples and (ii) a computationally more efficient approach based on EP (Minka, 2001).

3.2.1 Using Rejection Sampling

While no closed-form for $p(y(\mathbf{x}) | \mathcal{D}_n, g_k^*)$ is known, it is straightforward to sample from this distribution via rejection sampling (RS). For the RS, a sampled function $\tilde{f}(\mathbf{x})$ from $p(y(\mathbf{x}) | \mathcal{D}_n)$ is generated and its robust counterpart $\tilde{g}(\mathbf{x})$ is computed. Given a robust maximum value sample $g_k^* \sim p(g^* | \mathcal{D}_n)$, a sample is accepted when the maximum of $\tilde{g}(\mathbf{x})$ is smaller than g_k^* . This process is repeated until L samples have been accepted. Pseudo-code for this procedure is shown in Alg. 1. Given the set \tilde{Y} of L accepted samples, we can approximate the entropy of the sample distribution as proposed by Ahmad and Lin (1976), $H[p(y(\mathbf{x}) | \mathcal{D}_n, g_k^*)] \approx -\frac{1}{L} \sum_{\tilde{y}_i \in \tilde{Y}} \ln[\hat{p}(\tilde{y}_i)]$, with $\hat{p}(\cdot)$ being the kernel density estimate of $p(y(\mathbf{x}) | \mathcal{D}_n, g_k^*)$ based

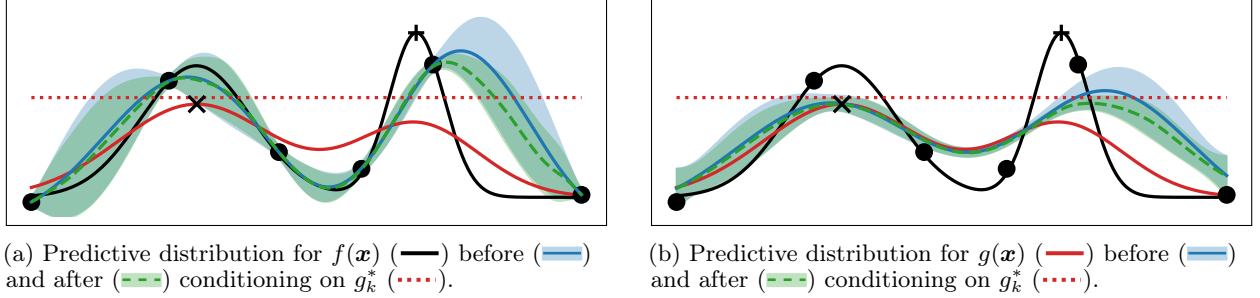


Figure 1: Comparison of predictive distributions for the objective $f(\mathbf{x})$ (left) and the robust objective $g(\mathbf{x})$ (right) before and after conditioning on the sampled robust maximum value g_k^* . The goal is to find the robust maximum (\times) instead of the global maximum ($+$); only $f(\mathbf{x})$ can be observed (\bullet).

Algorithm 1 Rejection sampling for $p(y(\mathbf{x})|\mathcal{D}_n, g_k^*)$

- 1: **Input:** GP posterior predictive distribution $p(y(\mathbf{x})|\mathcal{D}_n)$, robust maximum value sample g_k^*
 - 2: **Output:** Set \tilde{Y} of L accepted samples
 - 3: $\tilde{Y} \leftarrow \emptyset$
 - 4: **while** $|\tilde{Y}| \leq L$ **do**
 - 5: $\tilde{f}(\mathbf{x}) \sim p(y(\mathbf{x})|\mathcal{D}_n)$ // Generate sample
 - 6: $\tilde{g}(\mathbf{x}) \leftarrow \int \tilde{f}(\mathbf{x} + \boldsymbol{\xi}) p(\boldsymbol{\xi}) d\boldsymbol{\xi}$ // Robust sample
 - 7: **if** $\max_{\mathbf{x}} \tilde{g}(\mathbf{x}) \leq g_k^*$ **then**
 - 8: $\tilde{Y} \leftarrow \tilde{Y} \cup \{\tilde{f}(\mathbf{x}) + \epsilon\}$ // Store sample
 - 9: **end if**
 - 10: **end while**
 - 11: **return** \tilde{Y}
-

on \tilde{Y} (Rosenblatt, 1956). Note that $\hat{p}(\cdot)$ sums over all samples in \tilde{Y} , such that the entropy computation scales quadratically with L due to a nested summation over \tilde{Y} . In the experiments, we found that $L = 1000$ samples result in a sufficiently accurate approximation. Due to the Monte Carlo approximation in Eq. (5), the RS step is conducted for each $g_k^* \in G^*$, which renders the optimization of $\alpha_{\text{NES}}(\mathbf{x})$ costly. In the following, we develop a more efficient approximation scheme.

3.2.2 Using Expectation Propagation

For a computationally more efficient approximation of $H[p(y(\mathbf{x})|\mathcal{D}_n, g_k^*)]$, we exploit the fact that the entropy of a normal distribution is given analytically. As the observation noise is additive, we can approximate the predictive distribution $p(f(\mathbf{x})|\mathcal{D}_n, g_k^*)$ and then add the observation noise to compute the entropy. In the remainder of this section we discuss how to compute a Gaussian approximation to $p(f(\mathbf{x})|\mathcal{D}_n, g_k^*)$ with EP. More details are given in the supplementary material.

We rewrite the conditioned posterior predictive distribution

of $f(\mathbf{x})$ as

$$p(f(\mathbf{x})|\mathcal{D}_n, g_k^*) = \int p(f(\mathbf{x})|\mathcal{D}_n, g(\mathbf{x})) p(g(\mathbf{x})|\mathcal{D}_n, g_k^*) dg(\mathbf{x}). \quad (7)$$

The first distribution, $p(f(\mathbf{x})|\mathcal{D}_n, g(\mathbf{x}))$, can be computed from GP arithmetic. Note that the joint distribution $p(f(\mathbf{x}), \mathcal{D}_n, g(\mathbf{x}))$ is a multivariate normal distribution and conditioning on \mathcal{D}_n and $g(\mathbf{x})$ results in $p(f(\mathbf{x})|\mathcal{D}_n, g(\mathbf{x})) = \mathcal{N}(f(\mathbf{x})|\boldsymbol{\mu}_f, \boldsymbol{\Sigma}_f)$. The second distribution in the integral, $p(g(\mathbf{x})|\mathcal{D}_n, g_k^*)$, is the predictive distribution for $g(\mathbf{x})$ with the constraint $g(\mathbf{x}) \leq g_k^*$. This constraint can either be incorporated by a truncated normal distribution (Wang and Jegelka, 2017) or by a Gaussian approximation (Hoffman and Ghahramani, 2015). We follow the latter such that our approximation of the integral in Eq. (7) has an analytic solution:

1. The constraint $g(\mathbf{x}) \leq g_k^*$ implies in particular that $g(\mathbf{x}_i) \leq g_k^*$ for all $\mathbf{x}_i \in \mathcal{D}_n$, which results in a truncated normal distribution for $\mathbf{g} = [g(\mathbf{x}_1), \dots, g(\mathbf{x}_n)]^\top$. Aside from the univariate case, there exist no closed-form expressions for the mean and covariance of this distribution. We approximate the respective moments with EP (Herbrich, 2005), denoting the indicator function by $\mathbf{1}_{\{\cdot\}}$,

$$p(\mathbf{g}|\mathcal{D}_n, g_k^*) \propto p(\mathbf{g}|\mathcal{D}_n) \prod_{i=1}^n \mathbf{1}_{\{\mathbf{x}_i | g(\mathbf{x}_i) \leq g_k^*\}} \stackrel{\text{(EP)}}{\approx} \mathcal{N}(\mathbf{g}|\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1). \quad (8)$$

2. By marginalizing out the latent function values \mathbf{g} , we obtain a predictive distribution. Deriving $p(g(\mathbf{x})|\mathbf{g}, \mathcal{D}_n)$ from GP arithmetic and substituting Eq. (8) results in

$$p_0(g(\mathbf{x})|\mathcal{D}_n, g_k^*) = \int p(g(\mathbf{x})|\mathbf{g}, \mathcal{D}_n) p(\mathbf{g}|\mathcal{D}_n, g_k^*) d\mathbf{g} \approx \mathcal{N}(g(\mathbf{x})|m_0(\mathbf{x}), v_0(\mathbf{x})). \quad (9)$$

3. Next, we incorporate the constraint that $g(\mathbf{x}) \leq g_k^*$ for all $\mathbf{x} \in \mathcal{X}$ by moment matching,

$$p(g(\mathbf{x})|\mathcal{D}_n, g_k^*) \propto p_0(g(\mathbf{x})|\mathcal{D}_n, g_k^*) \mathbb{1}_{\{g(\mathbf{x}) \leq g_k^*\}} \\ \approx \mathcal{N}(g(\mathbf{x})|\hat{m}(\mathbf{x}), \hat{v}(\mathbf{x})).$$

With the shorthand notation $\beta = (g_k^* - m_0(\mathbf{x})) / \sqrt{v_0(\mathbf{x})}$ and $r = \varphi(\beta) / \Phi(\beta)$, mean and variance are given by $\hat{m}(\mathbf{x}) = m_0(\mathbf{x}) - \sqrt{v_0(\mathbf{x})}r$ and $\hat{v}(\mathbf{x}) = v_0(\mathbf{x}) - v_0(\mathbf{x})r(r + \beta)$ (see, e.g., (Jawitz, 2004)), where $\varphi(\cdot)$ and $\Phi(\cdot)$ denote the probability density function and cumulative density function of the standard normal distribution, respectively. The influence of conditioning $p(g(\mathbf{x})|\mathcal{D}_n)$ on the robust maximum value g_k^* is visualized in Fig. 1b.

4. Approximating $p(g(\mathbf{x})|\mathcal{D}_n, g_k^*)$ with a Gaussian offers the benefit that the integral in Eq. (7) can be solved analytically as it is the marginalization over a product of Gaussians. Thus, the approximation to the posterior predictive distribution for $f(\mathbf{x})$ conditioned on g_k^* is given by

$$p(f(\mathbf{x})|\mathcal{D}_n, g_k^*) \approx \mathcal{N}(f(\mathbf{x})|\tilde{m}_k(\mathbf{x}), \tilde{v}_k(\mathbf{x})).$$

An exemplary visualization of this approximation is displayed in Fig. 1a. Note that both, mean $\tilde{m}_k(\mathbf{x})$ and variance $\tilde{v}_k(\mathbf{x})$, are strongly influenced in regions of large predicted values.

The final form of the NES acquisition function based on EP is then given by

$$\alpha_{\text{NES-EP}}(\mathbf{x}) = \frac{1}{2} \left[\log \left(v_f(\mathbf{x}|\mathcal{D}_n) + \sigma_\epsilon^2 \right) - \frac{1}{K} \sum_{g_k^* \in G^*} \log \left(\tilde{v}_k(\mathbf{x}) + \sigma_\epsilon^2 \right) \right]. \quad (10)$$

For each evaluation of Eq. (10), the variance $\tilde{v}_k(\mathbf{x})$ is computed for every sample g_k^* separately. The EP step iterates over all data points. During experiments we found that it converges within 2–5 sweeps. Eq. (9) dominates the computational cost due to the inversion of a kernel matrix of size $2n$, with n being the number of data points. The overall complexity of one evaluation is then $\mathcal{O}(Kn^3)$. Please note that, unlike the RS-based approach that relies on a kernel density estimation, the entropy of the Gaussian approximation obtained with EP can be computed analytically. In the following, we evaluate the proposed NES acquisition function and compare the RS- and EP-based approximations.

4 Experiments

In this section, we evaluate the Noisy-Input Entropy Search (NES) acquisition function and compare it to

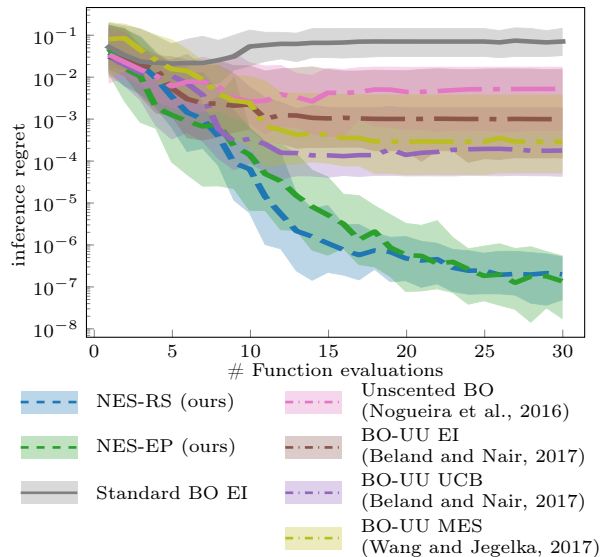


Figure 2: Within-model comparison in terms of the inference regret $r_n = |g(\mathbf{x}_n^*) - g^*|$. We present the median (lines) and 25/75th percentiles (shaded areas) across 50 different function samples from a GP prior.

other methods from the literature on a range of benchmark problems. Furthermore, we consider an application from aerospace engineering, for which robustness of the design parameters is crucial. For all experiments, we use a squared exponential (SE) kernel $k_f(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp(-0.5 \|\mathbf{x} - \mathbf{x}'\|_{\Lambda^{-1}}^2)$ with $\Lambda = \text{diag}[\ell_1^2, \dots, \ell_d^2]$. For this choice, $k_{gf}(\mathbf{x}, \mathbf{x}')$ and $k_g(\mathbf{x}, \mathbf{x}')$ (see Eq. (3)) can be computed analytically. As performance metric we choose the inference regret (IR) $r_n = |g(\mathbf{x}_n^*) - g^*|$, where \mathbf{x}_n^* is the estimate of the robust optimum at iteration n . For all experiments, we perform 100 independent runs, each with different sets of initial points, albeit the same set across all methods. The result figures show the median across all runs as line and the 25/75th percentiles as shaded area. The initial observations are uniformly sampled and the number of initial points is chosen depending on the dimensionality of the objective ($n_0 = 3, 5, 10$ for $d = 1, 2, 3$, respectively). We describe all evaluated methods below. All approaches were implemented based on GPy (GPy, since 2012) and the code to reproduce all results is publicly available at <https://github.com/boschresearch/NoisyInputEntropySearch>.

- **Noisy-Input Entropy Search (NES):** The proposed acquisition function using either rejection sampling (NES-RS) or expectation propagation (NES-EP). For both variants of NES, we use $M = 500$ random features for the SSGP and $K = 1$ samples for the Monte Carlo estimate of the expectation over $p(g^*|\mathcal{D}_n)$. The number of accepted samples for NES-RS is set to $L = 1000$.

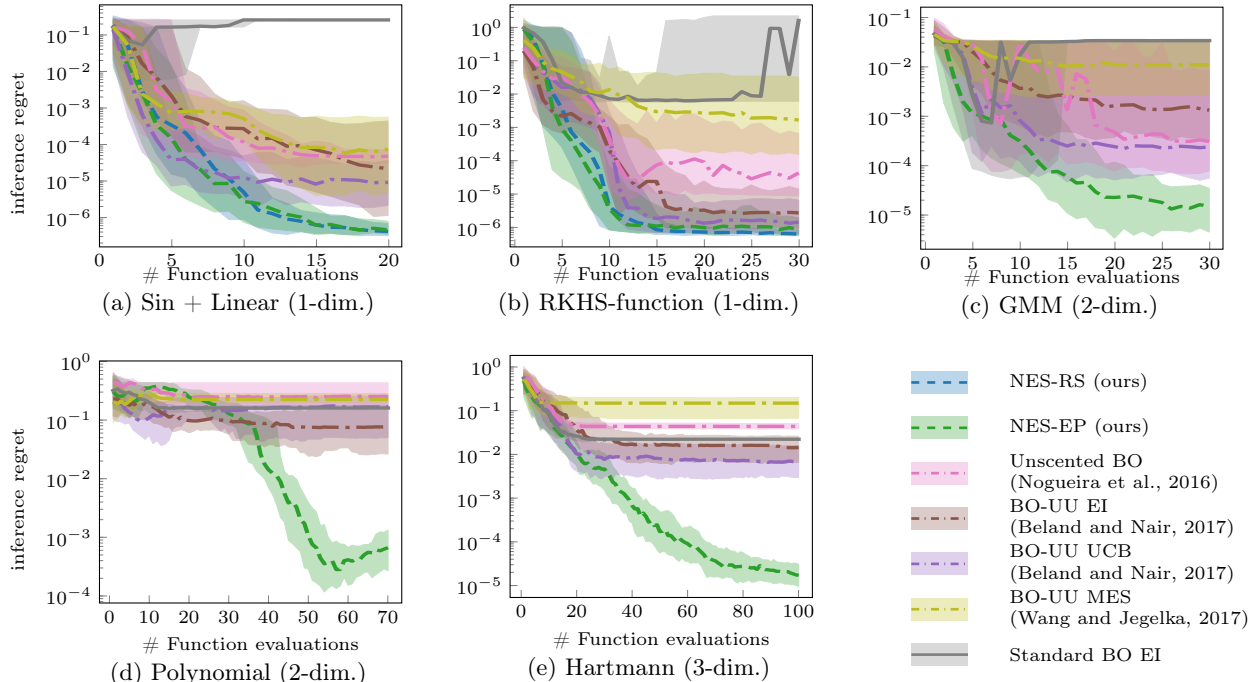


Figure 3: Inference regret $r_n = |g(\mathbf{x}_n^*) - g^*|$ on synthetic benchmark problems. We present the median (lines) and 25/75th percentiles (shaded areas) across 100 independent runs with randomly sampled initial points.

- **BO Under Uncertainty (BO-UU):** The method presented by Beland and Nair (2017) which models the robust objective $g(\mathbf{x})$ as a GP, but assumes that it can be observed directly. We evaluate BO-UU with expected improvement (EI), upper confidence bound (UCB) and MES (Wang and Jegelka, 2017).
- **Unscented BO:** The method presented by Nogueira et al. (2016) where the expectation over the input noise is approximated using an unscented transformation (Julier and Uhlmann, 2004).
- **Standard BO:** Furthermore, we compare against standard BO with EI as acquisition function, which in general gives rise to non-robust optima.

4.1 Within-Model Comparison

In a first step, we follow Hennig and Schuler (2012) and perform a within-model comparison. For this analysis, we draw 50 function samples from a 1-dim. GP prior (SE-kernel with $\sigma_f = 0.5$, $\ell = 0.05$) and for each sample we try to find the robust optimum assuming the input noise $\sigma_x = 0.05$. During optimization, the GP hyperparameters are fixed to their true values. The benefit of this analysis is to isolate the influence of the acquisition functions from other factors such as the inference of hyperparameters or a potential model mismatch between objective and GP model. For unknown objective functions, however, this procedure is not possible and the hyperparameters need to be inferred during optimization (see Sec. 4.2). The results of the within-model

comparison are presented in Fig. 2. Clearly, the two proposed acquisition functions NES-RS and NES-EP outperform all other approaches. We observed that the NES acquisition functions continue to explore the vicinity of the robust optimum even at later stages of the optimization. The other acquisition functions, however, stop exploring prematurely, which explains why the IR-curves level off early in Fig. 2. Furthermore, the performance of both NES variants is very similar in terms of IR, indicating that the EP-based approach is able to approximate the entropy terms in Eq. (4) similarly well compared to the RS-based approach, but at lower computational cost.

4.2 Synthetic Benchmark Functions

We evaluate the aforementioned methods on the following functions:

- $f(\mathbf{x}) = \sin(5\pi\mathbf{x}^2) + 0.5\mathbf{x}$, with $\mathbf{x} \in [0, 1]$ and input noise $\sigma_x = 0.05$,
- RKHS-function (1-dim.) (Assael et al., 2014) with $\sigma_x = 0.03$, also used by Nogueira et al. (2016),
- Gaussian mixture model (2-dim.) with $\Sigma_x = 0.1^2\mathbf{I}$, also used by Nogueira et al. (2016),
- Polynomial (2-dim.) (Bertsimas et al., 2010) with $\Sigma_x = 0.6^2\mathbf{I}$, also used by Bogunovic et al. (2018). Here, we scaled and shifted the objective $f(\mathbf{x})$ s.t. $\mathbb{E}[f(\mathbf{x})] = 0.0$ and $\mathbb{V}[f(\mathbf{x})] = 1.0$,
- Hartmann (3-dim.) with $\Sigma_x = 0.1^2\mathbf{I}$.

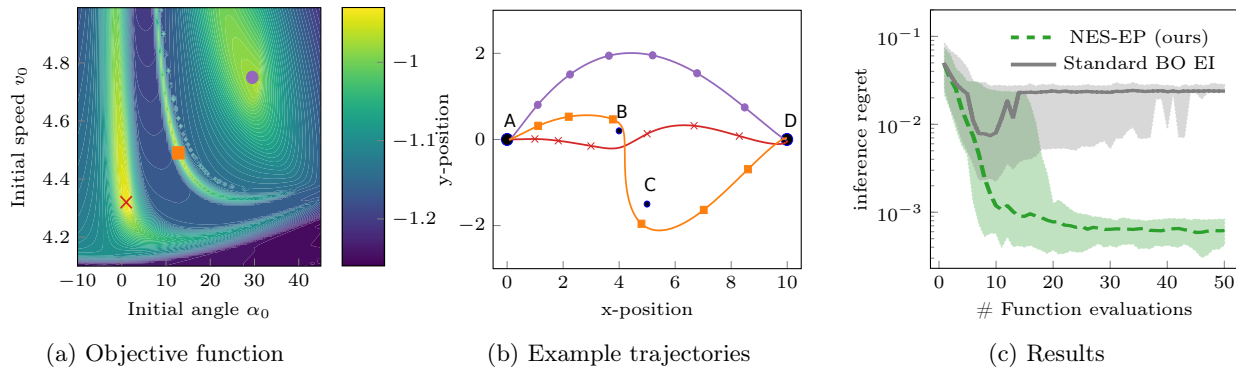


Figure 4: Gravity assist maneuver: The goal of the maneuver is to get from planet A to planet D (see Fig. 4b) by choosing an appropriate initial speed and starting angle. During the flight the engines are turned off such that all direction changes happen due to gravitational forces of planets A–D. The objective penalizes high values for the initial speed as well as the distance to the target planet, which results in the objective function as shown in Fig. 4a. Results are shown in Fig. 4c.

Visualizations of the 1- and 2-dimensional functions are shown in the supplementary material. The kernel hyperparameters as well as the observation noise are inferred via marginal likelihood maximization after each function evaluation. Additionally, we chose a log-normal hyperprior for the kernel lengthscales, in order to relate them with the magnitude of the input noise which led to significantly more stable convergence for all acquisition functions.

In general, NES shows better convergence and IR across all benchmark functions, and this performance benefit increases with the dimensionality of the problem. For all other methods that are designed to find the robust optimum, the performance is strongly task dependent. Moreover, standard BO always finds the global optimum, which is, however, sub-optimal in the robust setting. Note that the performance of standard BO appears to be competitive in terms of IR in Fig. 3d/3e, but the location of the estimated optimum is far off (see also the supplementary material for an evaluation of the distance to the optimum).

4.3 Application to Gravity Assist Maneuver

We evaluate the proposed acquisition function on a so-called gravity assist maneuver. The goal is to plan a spacecrafts trajectory to a target planet while minimizing energy consumption. For this task gravitational effects from other planets are exploited in order to change the spacecraft’s momentum without the need for active steering, thus saving fuel. A visualization of this scenario is shown in Fig. 4b. The decision variables are 1) v_0 , the initial speed of the spacecraft and 2) α_0 , the initial angle of flight from the start position. The optimization objective for this task is given by $J(\alpha_0, v_0) = \log_{10}(d_{\text{target}} + \beta \cdot v_0)$, with d_{target} being the closest distance of the resulting trajectory to the target planet and β is a parameter that trades-off between the

two cost terms. The resulting cost function is depicted in Fig. 4a, where the markers correspond to different resulting trajectories shown in Fig. 4b. The input noise is set to $\sigma_v = 0.05$ and $\sigma_\alpha = 3^\circ$ for v_0 and α_0 , respectively. The results for NES-EP and standard BO (EI) are depicted in Fig. 4c. Using either of the acquisition functions, the broad local optimum at the upper right corner of the domain (see Fig. 4a) is quickly explored within the first function evaluations. After 10–15 function evaluations, standard BO finds the global optimum in the lower left corner and continues to exploit this region of the domain. However, the global optimum is sensitive to perturbations on α_0 and thus the inference regret stagnates. On the other hand, NES-EP reliably finds the robust optimum and continues to explore the vicinity. As a result, the inference regret is almost two orders of magnitude smaller compared to standard BO.

5 Conclusion

In this paper, we introduced a novel information-theoretic acquisition function for robust Bayesian optimization. Our method, Noisy-Input Entropy Search (NES), considers a probabilistic formulation of the robust objective and maximizes the information gain about the robust maximum value g^* . Evaluation of NES requires the computation of the GP’s predictive distribution conditioned on the robust maximum value. As this distribution is analytically intractable, we propose two approximation schemes. The first is based on rejection sampling and is exact in the limit of infinite samples, but computationally challenging. For the second approximation scheme we employ expectation propagation, which is computationally more efficient. NES outperforms existing methods from the literature on a range of benchmark problems. Finally, we demonstrated the practical importance of the proposed approach on a task from aerospace engineering where robustness is critical.

Acknowledgements

The research of Melanie N. Zeilinger was supported by the Swiss National Science Foundation under grant no. PP00P2 157601/1.

References

- Elodie Adida and Georgia Perakis. A robust optimization approach to dynamic pricing and inventory control with no backorders. *Mathematical Programming*, 107(1-2):97–129, 2006.
- Ibrahim A. Ahmad and Pi-Erh Lin. A nonparametric estimation of the entropy for absolutely continuous distributions. *IEEE Transactions on Information Theory*, 22(3):372–375, 1976.
- John-Alexander M. Assael, Ziyu Wang, Bobak Shahriari, and Nando de Freitas. Heteroscedastic treed Bayesian optimisation. *arXiv preprint:1410.7172*, 2014.
- Tamer Başar and Pierre Bernhard. *H-infinity optimal control and related minimax design problems: a dynamic game approach*. Springer Science & Business Media, 2008.
- Justin J. Beland and Prasanth B. Nair. Bayesian optimization under uncertainty. *NIPS Workshop on Bayesian Optimization*, 2017.
- Dimitris Bertsimas, Omid Nohadani, and Kwong Meng Teo. Robust optimization for unconstrained simulation-based problems. *Operations Research*, 58(1):161–178, 2010.
- Hans-Georg Beyer and Bernhard Sendhoff. Robust optimization—a comprehensive survey. *Computer Methods in Applied Mechanics and Engineering*, 196(33-34):3190–3218, 2007.
- Ilija Bogunovic, Jonathan Scarlett, Stefanie Jegelka, and Volkan Cevher. Adversarially robust optimization with Gaussian processes. In *Advances in Neural Information Processing Systems (NIPS)*, pages 5765–5775, 2018.
- Eric Brochu, Vlad M. Cora, and Nando de Freitas. A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint:1012.2599*, 2010.
- Roberto Calandra, André Seyfarth, Jan Peters, and Marc Peter Deisenroth. Bayesian optimization for learning gaits under uncertainty. *Annals of Mathematics and Artificial Intelligence*, 76(1-2):5–23, 2016.
- Robert S. Chen, Brendan Lucier, Yaron Singer, and Vasilis Syrgkanis. Robust optimization for non-convex objectives. In *Advances in Neural Information Processing Systems (NIPS)*, pages 4705–4714, 2017.
- Wei Chen, Janet K. Allen, Kwok-Leung Tsui, and Farrokh Mistree. A procedure for robust design: minimizing variations caused by noise factors and control factors. *Journal of Mechanical Design*, 118(4):478–485, 1996.
- Dennis D. Cox and Susan John. A statistical method for global optimization. In *IEEE Transactions on Systems, Man, and Cybernetics*, pages 1242–1246, 1992.
- Antoine Cully, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. Robots that can adapt like animals. *Nature*, 521(7553):503–507, 2015.
- Patrick Dallaire, Camille Besse, and Brahim Chaib-Draa. Learning Gaussian process models from uncertain data. In *Proceedings of the International Conference on Neural Information Processing (ICONIP)*, pages 433–440, 2009.
- GPY. GPY: A Gaussian process framework in python. <http://github.com/SheffieldML/GPY>, since 2012.
- Ryan-Rhys Griffiths and José Miguel Hernández-Lobato. Constrained Bayesian optimization for automatic chemical design. *arXiv preprint:1709.05501*, 2017.
- Perry Groot, Adriana Birlutiu, and Tom Heskes. Bayesian Monte Carlo for the global optimization of expensive functions. In *Proceedings of the European Conference on Artificial Intelligence (ECAI)*, pages 249–254, 2010.
- Philipp Hennig and Christian J. Schuler. Entropy search for information-efficient global optimization. *Journal of Machine Learning Research*, 13:1809–1837, 2012.
- Ralf Herbrich. On Gaussian expectation propagation. Technical report, Microsoft Research Cambridge, 2005.
- José Miguel Hernández-Lobato, Matthew W. Hoffman, and Zoubin Ghahramani. Predictive entropy search for efficient global optimization of black-box functions. In *Advances in Neural Information Processing Systems (NIPS)*, pages 918–926, 2014.
- Matthew W. Hoffman and Zoubin Ghahramani. Output-space predictive entropy search for flexible global optimization. In *NIPS Workshop on Bayesian Optimization*, 2015.
- James W. Jawitz. Moments of truncated continuous univariate distributions. *Advances in Water Resources*, 27(3):269–281, 2004.
- Simon J. Julier and Jeffrey K. Uhlmann. Unscented filtering and nonlinear estimation. *Proceedings of the IEEE*, 92(3):401–422, 2004.
- Harold J. Kushner. A new method of locating the maximum point of an arbitrary multipeak curve in

- the presence of noise. *Journal of Basic Engineering*, 86(1):97–106, 1964.
- Miguel Lázaro-Gredilla, Joaquin Quiñonero-Candela, Carl Edward Rasmussen, and Aníbal R. Figueiras-Vidal. Sparse spectrum Gaussian process regression. *Journal of Machine Learning Research*, 11:1865–1881, 2010.
- Ruben Martinez-Cantin, Kevin Tee, and Michael McCourt. Practical Bayesian optimization in the presence of outliers. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1722–1731, 2018.
- Thomas P. Minka. Expectation propagation for approximate Bayesian inference. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 362–369, 2001.
- Jonas Moćkus. On Bayesian methods for seeking the extremum. In *Optimization Techniques IFIP Technical Conference*, pages 400–404, 1975.
- José Nogueira, Ruben Martinez-Cantin, Alexandre Bernardino, and Lorenzo Jamone. Unscented Bayesian optimization for safe robot grasping. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1967–1972, 2016.
- Rafael Oliveira, Lionel Ott, and Fabio Ramos. Bayesian optimisation under uncertain inputs. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1177–1184, 2019.
- Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- Murray Rosenblatt. Remarks on some nonparametric estimates of a density function. *The Annals of Mathematical Statistics*, pages 832–837, 1956.
- Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P. Adams, and Nando de Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016.
- Jasper Snoek, Oren Rippel, Kevin Swersky, Ryan Kiros, Nadathur Satish, Narayanan Sundaram, Mostofa Patwary, Mr Prabhat, and Ryan Adams. Scalable Bayesian optimization using deep neural networks. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 2171–2180, 2015.
- Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2010.
- Matthew Tesch, Jeff Schneider, and Howie Choset. Adapting control policies for expensive systems to changing environments. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 357–364. IEEE, 2011.
- Saul Toscano-Palmerin and Peter I. Frazier. Bayesian optimization with expensive integrands. *arXiv preprint:1803.08661*, 2018.
- Zi Wang and Stefanie Jegelka. Max-value entropy search for efficient Bayesian optimization. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 3627–3635, 2017.