

---

# Predictive Entropy Search for Efficient Global Optimization of Black-box Functions

---

**José Miguel Hernández-Lobato**  
jmh233@cam.ac.uk  
University of Cambridge

**Matthew W. Hoffman**  
mwh30@cam.ac.uk  
University of Cambridge

**Zoubin Ghahramani**  
zoubin@eng.cam.ac.uk  
University of Cambridge

## Abstract

We propose a novel information-theoretic approach for Bayesian optimization called Predictive Entropy Search (PES). At each iteration, PES selects the next evaluation point that maximizes the expected information gained with respect to the global maximum. PES codifies this intractable acquisition function in terms of the expected reduction in the differential entropy of the predictive distribution. This reformulation allows PES to obtain approximations that are both more accurate and efficient than other alternatives such as Entropy Search (ES). Furthermore, PES can easily perform a fully Bayesian treatment of the model hyperparameters while ES cannot. We evaluate PES in both synthetic and real-world applications, including optimization problems in machine learning, finance, biotechnology, and robotics. We show that the increased accuracy of PES leads to significant gains in optimization performance.

## 1 Introduction

Bayesian optimization techniques form a successful approach for optimizing black-box functions [5]. The goal of these methods is to find the global maximizer of a nonlinear and generally non-convex function  $f$  whose derivatives are unavailable. Furthermore, the evaluations of  $f$  are usually corrupted by noise and the process that queries  $f$  can be computationally or economically very expensive. To address these challenges, Bayesian optimization devotes additional effort to modeling the unknown function  $f$  and its behavior. These additional computations aim to minimize the number of evaluations that are needed to find the global optima.

Optimization problems are widespread in science and engineering and as a result so are Bayesian approaches to this problem. Bayesian optimization has successfully been used in robotics to adjust the parameters of a robot's controller to maximize gait speed and smoothness [16] as well as parameter tuning for computer graphics [6]. Another example application in drug discovery is to find the chemical derivative of a particular molecule that best treats a given disease [20]. Finally, Bayesian optimization can also be used to find optimal hyper-parameter values for statistical [29] and machine learning techniques [24].

As described above, we are interested in finding the global maximizer  $\mathbf{x}_* = \arg \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$  of a function  $f$  over some bounded domain, typically  $\mathcal{X} \subset \mathbb{R}^d$ . We assume that  $f(\mathbf{x})$  can only be evaluated via queries to a black-box that provides noisy outputs of the form  $y_i \sim \mathcal{N}(f(\mathbf{x}_i), \sigma^2)$ . We note, however, that our framework can be extended to other non-Gaussian likelihoods. In this setting, we describe a sequential search algorithm that, after  $n$  iterations, proposes to evaluate  $f$  at some location  $\mathbf{x}_{n+1}$ . To make this decision the algorithm conditions on all previous observations  $\mathcal{D}_n = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\}$ . After  $N$  iterations the algorithm makes a final recommendation  $\tilde{\mathbf{x}}_N$  for the global maximizer of the latent function  $f$ .

We take a Bayesian approach to the problem described above and use a probabilistic model for the latent function  $f$  to guide the search and to select  $\tilde{\mathbf{x}}_N$ . In this work we use a zero-mean Gaussian

---

**Algorithm 1** Generic Bayesian optimization

---

**Input:** a black-box with unknown mean  $f$

- 1: **for**  $n = 1, \dots, N$  **do**
- 2:   select  $\mathbf{x}_n = \arg \max_{\mathbf{x} \in \mathcal{X}} \alpha_{n-1}(\mathbf{x})$
- 3:   query the black-box at  $\mathbf{x}_n$  to obtain  $y_n$
- 4:   augment data  $\mathcal{D}_n = \mathcal{D}_{n-1} \cup \{(\mathbf{x}_n, y_n)\}$
- 5: **end for**
- 6: **return**  $\tilde{\mathbf{x}}_N = \arg \max_{\mathbf{x} \in \mathcal{X}} \mu_N(\mathbf{x})$

---



---

**Algorithm 2** PES acquisition function

---

**Input:** a candidate  $\mathbf{x}$ ; data  $\mathcal{D}_n$

- 1: sample  $M$  hyperparameter values  $\{\psi^{(i)}\}$
- 2: **for**  $i = 1, \dots, M$  **do**
- 3:   sample  $f^{(i)} \sim p(f|\mathcal{D}_n, \phi, \psi^{(i)})$
- 4:   set  $\mathbf{x}_*^{(i)} \leftarrow \arg \max_{\mathbf{x} \in \mathcal{X}} f^{(i)}(\mathbf{x})$
- 5:   compute  $\mathbf{m}_0^{(i)}, \mathbf{V}_0^{(i)}$  and  $\tilde{\mathbf{m}}^{(i)}, \tilde{\mathbf{v}}^{(i)}$
- 6:   compute  $v_n^{(i)}(\mathbf{x})$  and  $v_n^{(i)}(\mathbf{x}|\mathbf{x}_*^{(i)})$
- 7: **end for**
- 8: **return**  $\alpha_n(\mathbf{x})$  as in (10)

---

precomputed

process (GP) prior for  $f$  [22]. This prior is specified by a positive-definite kernel function  $k(\mathbf{x}, \mathbf{x}')$ . Given any finite collection of points  $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ , the values of  $f$  at these points are jointly zero-mean Gaussian with covariance matrix  $\mathbf{K}_n$ , where  $[\mathbf{K}_n]_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$ . For the Gaussian likelihood described above, the vector of concatenated observations  $\mathbf{y}_n$  is also jointly Gaussian with zero-mean. Therefore, at any location  $\mathbf{x}$ , the latent function  $f(\mathbf{x})$  conditioned on past observations  $\mathcal{D}_n$  is then Gaussian with marginal mean  $\mu_n(\mathbf{x})$  and variance  $v_n(\mathbf{x})$  given by

$$\mu_n(\mathbf{x}) = \mathbf{k}_n(\mathbf{x})^\top (\mathbf{K}_n + \sigma^2 \mathbf{I})^{-1} \mathbf{y}_n, \quad v_n(\mathbf{x}) = k(\mathbf{x}, \mathbf{x}) - \mathbf{k}_n(\mathbf{x})^\top (\mathbf{K}_n + \sigma^2 \mathbf{I})^{-1} \mathbf{k}_n(\mathbf{x}), \quad (1)$$

where  $\mathbf{k}_n(\mathbf{x})$  is a vector of cross-covariance terms between  $\mathbf{x}$  and  $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ .

Bayesian optimization techniques use the above predictive distribution  $p(f(\mathbf{x})|\mathcal{D}_n)$  to guide the search for the global maximizer  $\mathbf{x}_*$ . In particular,  $p(f(\mathbf{x})|\mathcal{D}_n)$  is used during the computation of an acquisition function  $\alpha_n(\mathbf{x})$  that is optimized at each iteration to determine the next evaluation location  $\mathbf{x}_{n+1}$ . This process is shown in Algorithm 1. Intuitively, the acquisition function  $\alpha_n(\mathbf{x})$  should be high in areas where the maxima is most likely to lie given the current data. However,  $\alpha_n(\mathbf{x})$  should also encourage exploration of the search space to guarantee that the recommendation  $\tilde{\mathbf{x}}_N$  is a global optimum of  $f$ , not just a global optimum of the posterior mean. Several acquisition functions have been proposed in the literature. Some examples are the probability of improvement (PI) [14], the expected improvement (EI) [19, 13] or upper confidence bounds (UCB) [26]. Alternatively, one can combine several of these acquisition functions [10].

The acquisition functions described above are based on probabilistic measures of improvement (PI an EI) or on optimistic estimates of the latent function (UCB) which implicitly trade off between exploiting the posterior mean and exploring based on the uncertainty. An alternate approach, introduced by [28], proposes maximizing the expected posterior information gain about the global maximizer  $\mathbf{x}_*$  evaluated over a grid in the input space. A similar strategy was later employed by [9] which although it requires no such grid, instead relies on a difficult-to-evaluate approximation. In Section 2 we derive a rearrangement of this information-based acquisition function which leads to a more straightforward approximation that we call Predictive Entropy Search (PES). In Section 3 we show empirically that our approximation is more accurate than that of [9]. We evaluate this claim on both synthetic and real-world problems and further show that this leads to real gains in performance.

## 2 Predictive entropy search

We propose to follow the information-theoretic method for active data collection described in [17]. We are interested in maximizing information about the location  $\mathbf{x}_*$  of the global maximum, whose posterior distribution is  $p(\mathbf{x}_*|\mathcal{D}_n)$ . Our current information about  $\mathbf{x}_*$  can be measured in terms of the negative differential entropy of  $p(\mathbf{x}_*|\mathcal{D}_n)$ . Therefore, our strategy is to select  $\mathbf{x}_{n+1}$  which maximizes the expected reduction in this quantity. The corresponding acquisition function is

$$\alpha_n(\mathbf{x}) = \mathbb{H}[p(\mathbf{x}_*|\mathcal{D}_n)] - \mathbb{E}_{p(y|\mathcal{D}_n, \mathbf{x})}[\mathbb{H}[p(\mathbf{x}_*|\mathcal{D}_n \cup \{(\mathbf{x}, y)\})]], \quad (2)$$

where  $\mathbb{H}[p(\mathbf{x})] = -\int p(\mathbf{x}) \log p(\mathbf{x}) d\mathbf{x}$  represents the differential entropy of its argument and the expectation above is taken with respect to the posterior predictive distribution of  $y$  given  $\mathbf{x}$ . The exact evaluation of (2) is infeasible in practice. The main difficulties are i)  $p(\mathbf{x}_*|\mathcal{D}_n \cup \{(\mathbf{x}, y)\})$  must be computed for many different values of  $\mathbf{x}$  and  $y$  during the optimization of (2) and ii) the entropy computations themselves are not analytical. In practice, a direct evaluation of (2) is only

possible after performing many approximations [9]. To avoid this, we follow the approach described in [11] by noting that (2) can be equivalently written as the mutual information between  $\mathbf{x}_*$  and  $y$  given  $\mathcal{D}_n$ . Since the mutual information is a symmetric function,  $\alpha_n(\mathbf{x})$  can be rewritten as

$$\alpha_n(\mathbf{x}) = \mathbb{H}[p(y|\mathcal{D}_n, \mathbf{x})] - \mathbb{E}_{p(\mathbf{x}_*|\mathcal{D}_n)}[\mathbb{H}[p(y|\mathcal{D}_n, \mathbf{x}, \mathbf{x}_*)]], \quad (3)$$

where  $p(y|\mathcal{D}_n, \mathbf{x}, \mathbf{x}_*)$  is the posterior predictive distribution for  $y$  given the observed data  $\mathcal{D}_n$  and the location of the global maximizer of  $f$ . Intuitively, conditioning on the location  $\mathbf{x}_*$  pushes the posterior predictions up in locations around  $\mathbf{x}_*$  and down in regions away from  $\mathbf{x}_*$ . Note that, unlike the previous formulation, this objective is based on the entropies of predictive distributions, which are analytic or can be easily approximated, rather than on the entropies of distributions on  $\mathbf{x}_*$  whose approximation is more challenging.

The first term in (3) can be computed analytically using the posterior marginals for  $f(\mathbf{x})$  in (1), that is,  $\mathbb{H}[p(y|\mathcal{D}_n, \mathbf{x})] = 0.5 \log[2\pi e(v_n(\mathbf{x}) + \sigma^2)]$ , where we add  $\sigma^2$  to  $v_n(\mathbf{x})$  because  $y$  is obtained by adding Gaussian noise with variance  $\sigma^2$  to  $f(\mathbf{x})$ . The second term, on the other hand, must be approximated. We first approximate the expectation in (3) by averaging over samples  $\mathbf{x}_*^{(i)}$  drawn approximately from  $p(\mathbf{x}_*|\mathcal{D}_n)$ . For each of these samples, we then approximate the corresponding entropy function  $\mathbb{H}[p(y|\mathcal{D}_n, \mathbf{x}, \mathbf{x}_*^{(i)})]$  using expectation propagation [18]. The code for all these operations is publicly available at <http://jmhl.org>.

## 2.1 Sampling from the posterior over global maxima

In this section we show how to approximately sample from the conditional distribution of the global maximizer  $\mathbf{x}_*$  given the observed data  $\mathcal{D}_n$ , that is,

$$p(\mathbf{x}_*|\mathcal{D}_n) = p(f(\mathbf{x}_*) = \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})|\mathcal{D}_n). \quad (4)$$

If the domain  $\mathcal{X}$  is restricted to some finite set of  $m$  points, the latent function  $f$  takes the form of an  $m$ -dimensional vector  $\mathbf{f}$ . The probability that the  $i$ th element of  $\mathbf{f}$  is optimal can then be written as  $\int p(\mathbf{f}|\mathcal{D}_n) \prod_{j \leq m} \mathbb{I}[f_i \geq f_j] d\mathbf{f}$ . This suggests the following generative process: i) draw a sample from the posterior distribution  $p(\mathbf{f}|\mathcal{D}_n)$  and ii) return the index of the maximum element in the sampled vector. This process is known as Thompson sampling or probability matching when used as an arm-selection strategy in multi-armed bandits [8]. This same approach could be used for sampling the maximizer over a continuous domain  $\mathcal{X}$ . At first glance this would require constructing an infinite-dimensional object representing the function  $f$ . To avoid this, one could sequentially construct  $f$  while it is being optimized. However, evaluating such an  $f$  would ultimately have cost  $\mathcal{O}(m^3)$  where  $m$  is the number of function evaluations necessary to find the optimum. Instead, we propose to sample and optimize an analytic approximation to  $f$ . We will briefly derive this approximation below, but more detail is given in Appendix A.

Given a shift-invariant kernel  $k$ , Bochner’s theorem [4] asserts the existence of its Fourier dual  $s(\mathbf{w})$ , which is equal to the spectral density of  $k$ . Letting  $p(\mathbf{w}) = s(\mathbf{w})/\alpha$  be the associated normalized density, we can write the kernel as the expectation

$$k(\mathbf{x}, \mathbf{x}') = \alpha \mathbb{E}_{p(\mathbf{w})}[e^{-i\mathbf{w}^\top(\mathbf{x}-\mathbf{x}')}] = 2\alpha \mathbb{E}_{p(\mathbf{w}, b)}[\cos(\mathbf{w}^\top \mathbf{x} + b) \cos(\mathbf{w}^\top \mathbf{x}' + b)], \quad (5)$$

where  $b \sim \mathcal{U}[0, 2\pi]$ . Let  $\phi(\mathbf{x}) = \sqrt{2\alpha/m} \cos(\mathbf{W}\mathbf{x} + \mathbf{b})$  denote an  $m$ -dimensional feature mapping where  $\mathbf{W}$  and  $\mathbf{b}$  consist of  $m$  stacked samples from  $p(\mathbf{w}, b)$ . The kernel  $k$  can then be approximated by the inner product of these features,  $k(\mathbf{x}, \mathbf{x}') \approx \phi(\mathbf{x})^\top \phi(\mathbf{x}')$ . This approach was used by [21] as an approximation method in the context of kernel methods. The feature mapping  $\phi(\mathbf{x})$  allows us to approximate the Gaussian process prior for  $f$  with a linear model  $f(\mathbf{x}) = \phi(\mathbf{x})^\top \boldsymbol{\theta}$  where  $\boldsymbol{\theta} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  is a standard Gaussian. By conditioning on  $\mathcal{D}_n$ , the posterior for  $\boldsymbol{\theta}$  is also multivariate Gaussian,  $\boldsymbol{\theta}|\mathcal{D}_n \sim \mathcal{N}(\mathbf{A}^{-1}\boldsymbol{\Phi}^\top \mathbf{y}_n, \sigma^2 \mathbf{A}^{-1})$  where  $\mathbf{A} = \boldsymbol{\Phi}^\top \boldsymbol{\Phi} + \sigma^2 \mathbf{I}$  and  $\boldsymbol{\Phi}^\top = [\phi(\mathbf{x}_1) \dots \phi(\mathbf{x}_n)]$ .

Let  $\phi^{(i)}$  and  $\boldsymbol{\theta}^{(i)}$  be a random set of features and the corresponding posterior weights sampled both according to the generative process given above. They can then be used to construct the function  $f^{(i)}(\mathbf{x}) = \phi^{(i)}(\mathbf{x})^\top \boldsymbol{\theta}^{(i)}$ , which is an approximate posterior sample of  $f$ —albeit one with a finite parameterization. We can then maximize this function to obtain  $\mathbf{x}_*^{(i)} = \arg \max_{\mathbf{x} \in \mathcal{X}} f^{(i)}(\mathbf{x})$ , which is approximately distributed according to  $p(\mathbf{x}_*|\mathcal{D}_n)$ . Note that for early iterations when  $n < m$ , we can efficiently sample  $\boldsymbol{\theta}^{(i)}$  with cost  $\mathcal{O}(n^2 m)$  using the method described in Appendix B.2 of [23]. This allows us to use a large number of features in  $\phi^{(i)}(\mathbf{x})$ .

## 2.2 Approximating the predictive entropy

We now show how to approximate  $\mathbb{H}[p(y|\mathcal{D}_n, \mathbf{x}, \mathbf{x}_*)]$  in (3). Note that we can write the argument to  $\mathbb{H}$  in this expression as  $p(y|\mathcal{D}_n, \mathbf{x}, \mathbf{x}_*) = \int p(y|f(\mathbf{x}))p(f(\mathbf{x})|\mathcal{D}_n, \mathbf{x}_*)df(\mathbf{x})$ . Here  $p(f(\mathbf{x})|\mathcal{D}_n, \mathbf{x}_*)$  is the posterior distribution on  $f(\mathbf{x})$  given  $\mathcal{D}_n$  and the location  $\mathbf{x}_*$  of the global maximizer of  $f$ . When the likelihood  $p(y|f(\mathbf{x}))$  is Gaussian, we have that  $p(f(\mathbf{x})|\mathcal{D}_n)$  is analytically tractable since it is the predictive distribution of a Gaussian process. However, by further conditioning on the location  $\mathbf{x}_*$  of the global maximizer we are introducing additional constraints, namely that  $f(\mathbf{z}) \leq f(\mathbf{x}_*)$  for all  $\mathbf{z} \in \mathcal{X}$ . These constraints make  $p(f(\mathbf{x})|\mathcal{D}_n, \mathbf{x}_*)$  intractable. To circumvent this difficulty, we instead use the following simplified constraints:

- C1.  **$\mathbf{x}_*$  is a local maximum.** This is achieved by letting  $\nabla f(\mathbf{x}_*) = \mathbf{0}$  and ensuring that  $\nabla^2 f(\mathbf{x}_*)$  is negative definite. We further assume that the non-diagonal elements of  $\nabla^2 f(\mathbf{x}_*)$ , denoted by  $\text{upper}[\nabla^2 f(\mathbf{x}_*)]$ , are known, for example they could all be zero. This simplifies the negative-definite constraint. We denote by C1.1 the constraint given by  $\nabla f(\mathbf{x}_*) = \mathbf{0}$  and  $\text{upper}[\nabla^2 f(\mathbf{x}_*)] = \mathbf{0}$ . We denote by C1.2 the constraint that forces the elements of  $\text{diag}[\nabla^2 f(\mathbf{x}_*)]$  to be negative.
- C2.  **$f(\mathbf{x}_*)$  is larger than past observations.** We also assume that  $f(\mathbf{x}_*) \geq f(\mathbf{x}_i)$  for all  $i \leq n$ . However, we only observe  $f(\mathbf{x}_i)$  noisily via  $y_i$ . To avoid making inference on these latent function values, we approximate the above hard constraints with the soft constraint  $f(\mathbf{x}_*) > y_{\max} + \epsilon$ , where  $\epsilon \sim \mathcal{N}(0, \sigma^2)$  and  $y_{\max}$  is the largest  $y_i$  seen so far.
- C3.  **$f(\mathbf{x})$  is smaller than  $f(\mathbf{x}_*)$ .** This simplified constraint only conditions on the given  $\mathbf{x}$  rather than requiring  $f(\mathbf{x}_*) \leq f(\mathbf{z})$  for all  $\mathbf{z} \in \mathcal{X}$ .

We incorporate these simplified constraints into  $p(f(\mathbf{x})|\mathcal{D}_n)$  to approximate  $p(f(\mathbf{x})|\mathcal{D}_n, \mathbf{x}_*)$ . This is achieved by multiplying  $p(f(\mathbf{x})|\mathcal{D}_n)$  with specific factors that encode the above constraints. In what follows we briefly show how to construct these factors; more detail is given in Appendix B.

Consider the latent variable  $\mathbf{z} = [f(\mathbf{x}_*); \text{diag}[\nabla^2 f(\mathbf{x}_*)]]$ . To incorporate constraint C1.1 we can condition on the data and on the ‘‘observations’’ given by the constraints  $\nabla f(\mathbf{x}_*) = \mathbf{0}$  and  $\text{upper}[\nabla^2 f(\mathbf{x}_*)] = \mathbf{0}$ . Since  $f$  is distributed according to a GP, the joint distribution between  $\mathbf{z}$  and these observations is multivariate Gaussian. The covariance between the noisy observations  $\mathbf{y}_n$  and the extra noise-free derivative observations can be easily computed [25]. The resulting conditional distribution is also multivariate Gaussian with mean  $\mathbf{m}_0$  and covariance  $\mathbf{V}_0$ . These computations are similar to those performed in (1). Constraints C1.2 and C2 can then be incorporated by writing

$$p(\mathbf{z}|\mathcal{D}_n, \text{C1}, \text{C2}) \propto \Phi_{\sigma^2}(f(\mathbf{x}_*) - y_{\max}) \left[ \prod_{i=1}^d \mathbb{I}([\nabla^2 f(\mathbf{x}_*)]_{ii} \leq 0) \right] \mathcal{N}(\mathbf{z}|\mathbf{m}_0, \mathbf{V}_0), \quad (6)$$

where  $\Phi_{\sigma^2}$  is the cdf of a zero-mean Gaussian distribution with variance  $\sigma^2$ . The first new factor in this expression guarantees that  $f(\mathbf{x}_*) > y_{\max} + \epsilon$ , where we have marginalized  $\epsilon$  out, and the second set of factors guarantees that the entries in  $\text{diag}[\nabla^2 f(\mathbf{x}_*)]$  are negative.

Later integrals that make use of  $p(\mathbf{z}|\mathcal{D}_n, \text{C1}, \text{C2})$ , however, will not admit a closed-form expression. As a result we compute a Gaussian approximation  $q(\mathbf{z})$  to this distribution using Expectation Propagation (EP) [18]. The resulting algorithm is similar to the implementation of EP for binary classification with Gaussian processes [22]. EP approximates each non-Gaussian factor in (6) with a Gaussian factor whose mean and variance are  $\tilde{m}_i$  and  $\tilde{v}_i$ , respectively. The EP approximation can then be written as  $q(\mathbf{z}) \propto [\prod_{i=1}^{d+1} \mathcal{N}(z_i|\tilde{m}_i, \tilde{v}_i)]\mathcal{N}(\mathbf{z}|\mathbf{m}_0, \mathbf{V}_0)$ . Note that these computations have so far not depended on  $\mathbf{x}$ , so we can compute  $\{\mathbf{m}_0, \mathbf{V}_0, \tilde{\mathbf{m}}, \tilde{\mathbf{v}}\}$  once and store them for later use, where  $\tilde{\mathbf{m}} = (\tilde{m}_1, \dots, \tilde{m}_{d+1})$  and  $\tilde{\mathbf{v}} = (\tilde{v}_1, \dots, \tilde{v}_{d+1})$ .

We will now describe how to compute the predictive variance of some latent function value  $f(\mathbf{x})$  given these constraints. Let  $\mathbf{f} = [f(\mathbf{x}); f(\mathbf{x}_*)]$  be a vector given by the concatenation of the values of the latent function at  $\mathbf{x}$  and  $\mathbf{x}_*$ . The joint distribution between  $\mathbf{f}$ ,  $\mathbf{z}$ , the evaluations  $\mathbf{y}_n$  collected so far and the derivative ‘‘observations’’  $\nabla f(\mathbf{x}_*) = \mathbf{0}$  and  $\text{upper}[\nabla^2 f(\mathbf{x}_*)] = \mathbf{0}$  is multivariate Gaussian. Using  $q(\mathbf{z})$ , we then obtain the following approximation:

$$p(\mathbf{f}|\mathcal{D}_n, \text{C1}, \text{C2}) \approx \int p(\mathbf{f}|\mathbf{z}, \mathcal{D}_n, \text{C1.1})q(\mathbf{z})d\mathbf{z} = \mathcal{N}(\mathbf{f}|\mathbf{m}_{\mathbf{f}}, \mathbf{V}_{\mathbf{f}}). \quad (7)$$

Implicitly we are assuming above that  $\mathbf{f}$  depends on our observations and constraint C1.1, but is independent of C1.2 and C2 given  $\mathbf{z}$ . The computations necessary to obtain  $\mathbf{m}_{\mathbf{f}}$  and  $\mathbf{V}_{\mathbf{f}}$  are similar

to those used above and in (1). The required quantities are similar to the ones used by EP to make predictions in the Gaussian process binary classifier [22]. We can then incorporate C3 by multiplying  $\mathcal{N}(\mathbf{f}|\mathbf{m}_f, \mathbf{V}_f)$  with a factor that guarantees  $f(\mathbf{x}) < f(\mathbf{x}_*)$ . The predictive distribution for  $f(\mathbf{x})$  given  $\mathcal{D}_n$  and all the constraints can be approximated as

$$p(f(\mathbf{x})|\mathcal{D}_n, C1, C2, C3) \approx Z^{-1} \int \mathbb{I}(f_1 < f_2) \mathcal{N}(\mathbf{f}|\mathbf{m}_f, \mathbf{V}_f) df_2, \quad (8)$$

where  $Z$  is a normalization constant. The variance of the right hand side of (8) is given by

$$v_n(\mathbf{x}|\mathbf{x}_*) = [\mathbf{V}_f]_{1,1} - v^{-1}\beta(\beta + \alpha)\{[\mathbf{V}_f]_{1,1} - [\mathbf{V}_f]_{1,2}\}^2, \quad (9)$$

where  $v = [-1, 1]^T \mathbf{V}_f [-1, 1]$ ,  $\alpha = m/\sqrt{v}$ ,  $m = [-1, 1]^T \mathbf{m}_f$ ,  $\beta = \phi(\alpha)/\Phi(\alpha)$ , and  $\phi(\cdot)$  and  $\Phi(\cdot)$  are the standard Gaussian density function and cdf, respectively. By further approximating (8) by a Gaussian distribution with the same mean and variance we can write the entropy as  $H[p(y|\mathcal{D}_n, \mathbf{x}, \mathbf{x}_*)] \approx 0.5 \log[2\pi e(v_n(\mathbf{x}|\mathbf{x}_*) + \sigma^2)]$ .

The computation of (9) can be numerically unstable when  $s$  is very close to zero. This occurs when  $[\mathbf{V}_f]_{1,1}$  is very similar to  $[\mathbf{V}_f]_{1,2}$ . To avoid these numerical problems, we multiply  $[\mathbf{V}_f]_{1,2}$  by the largest  $0 \leq \kappa \leq 1$  that guarantees that  $s > 10^{-10}$ . This can be understood as slightly reducing the amount of dependence between  $f(\mathbf{x})$  and  $f(\mathbf{x}_*)$  when  $\mathbf{x}$  is very close to  $\mathbf{x}_*$ . Finally, fixing  $\text{upper}[\nabla^2 f(\mathbf{x}_*)]$  to be zero can also produce poor predictions when the actual  $f$  does not satisfy this constraint. To avoid this, we instead fix this quantity to  $\text{upper}[\nabla^2 f^{(i)}(\mathbf{x}_*)]$ , where  $f^{(i)}$  is the  $i$ th sample function optimized in Section 2.1 to sample  $\mathbf{x}_*^{(i)}$ .

### 2.3 Hyperparameter learning and the PES acquisition function

We now show how the previous approximations are integrated to compute the acquisition function used by predictive entropy search (PES). This acquisition function performs a formal treatment of the hyperparameters. Let  $\psi$  denote a vector of hyperparameters which includes any kernel parameters as well as the noise variance  $\sigma^2$ . Let  $p(\psi|\mathcal{D}_n) \propto p(\psi)p(\mathcal{D}_n|\psi)$  denote the posterior distribution over these parameters where  $p(\psi)$  is a hyperprior and  $p(\mathcal{D}_n|\psi)$  is the GP marginal likelihood. For a fully Bayesian treatment of  $\psi$  we must marginalize the acquisition function (3) with respect to this posterior. The corresponding integral has no analytic expression and must be approximated using Monte Carlo. This approach is also taken in [24].

We draw  $M$  samples  $\{\psi^{(i)}\}$  from  $p(\psi|\mathcal{D}_n)$  using slice sampling [27]. Let  $\mathbf{x}_*^{(i)}$  denote a sampled global maximizer drawn from  $p(\mathbf{x}_*|\mathcal{D}_n, \psi^{(i)})$  as described in Section 2.1. Furthermore, let  $v_n^{(i)}(\mathbf{x})$  and  $v_n^{(i)}(\mathbf{x}|\mathbf{x}_*^{(i)})$  denote the predictive variances computed as described in Section 2.2 when the model hyperparameters are fixed to  $\psi^{(i)}$ . We then write the marginalized acquisition function as

$$\alpha_n(\mathbf{x}) = \frac{1}{M} \sum_{i=1}^M \left\{ 0.5 \log[v_n^{(i)}(\mathbf{x}) + \sigma^2] - 0.5 \log[v_n^{(i)}(\mathbf{x}|\mathbf{x}_*^{(i)}) + \sigma^2] \right\}. \quad (10)$$

Note that PES is effectively marginalizing the original acquisition function (2) over  $p(\psi|\mathcal{D}_n)$ . This is a significant advantage with respect to other methods that optimize the same information-theoretic acquisition function but do not marginalize over the hyper-parameters. For example, the approach of [9] approximates (2) only for fixed  $\psi$ . The resulting approximation is computationally very expensive and recomputing it to average over multiple samples from  $p(\psi|\mathcal{D}_n)$  is infeasible in practice.

Algorithm 2 shows pseudo-code for computing the PES acquisition function. Note that most of the computations necessary for evaluating (10) can be done independently of the input  $\mathbf{x}$ , as noted in the pseudo-code. This initial cost is dominated by a matrix inversion necessary to pre-compute  $\mathbf{V}$  for each hyperparameter sample. The resulting complexity is  $\mathcal{O}[M(n+d+d(d-1)/2)^3]$ . This cost can be reduced to  $\mathcal{O}[M(n+d)^3]$  by ignoring the derivative observations imposed on  $\text{upper}[\nabla^2 f(\mathbf{x}_*)]$  by constraint C1.1. Nevertheless, in the problems that we consider  $d$  is very small (less than 20). After these precomputations are done, the evaluation of (10) is  $\mathcal{O}[M(n+d+d(d-1)/2)]$ .

## 3 Experiments

In our experiments, we use Gaussian process priors for  $f$  with squared-exponential kernels  $k(\mathbf{x}, \mathbf{x}') = \gamma^2 \exp\{-0.5 \sum_i (x_i - x'_i)^2 / \ell_i^2\}$ . The corresponding spectral density is zero-mean Gaussian with covariance given by  $\text{diag}\{\{\ell_i^{-2}\}\}$  and normalizing constant  $\alpha = \gamma^2$ . The model hyperparameters are  $\{\gamma, \ell_1, \dots, \ell_d, \sigma^2\}$ . We use broad, uninformative Gamma hyperpriors.

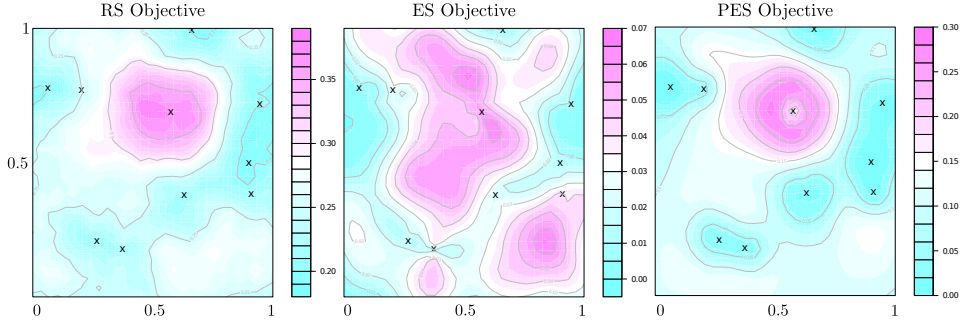


Figure 1: Comparison of different estimates of the objective function  $\alpha_n(\mathbf{x})$  given by (2). Left, ground truth obtained by the rejection sampling method RS. Middle, approximation produced by the ES method. Right, approximation produced by the proposed PES method. These plots show that the PES objective is much more similar to the RS ground truth than the ES objective.

First, we analyze the accuracy of PES in the task of approximating the differential entropy (2). We compare the PES approximation (10), with the approximation used by the entropy search (ES) method [9]. We also compare with the ground truth for (2) obtained using a rejection sampling (RS) algorithm based on (3). For this experiment we generate the data  $\mathcal{D}_n$  using an objective function  $f$  sampled from the Gaussian process prior as in [9]. The domain  $\mathcal{X}$  of  $f$  is fixed to be  $[0, 1]^2$  and data are generated using  $\gamma^2 = 1$ ,  $\sigma^2 = 10^{-6}$ , and  $\ell_i^2 = 0.1$ . To compute (10) we avoid sampling the hyperparameters and use the known values directly. We further fix  $M = 200$  and  $m = 1000$ .

The ground truth rejection sampling scheme works as follows. First,  $\mathcal{X}$  is discretized using a uniform grid. The expectation with respect to  $p(\mathbf{x}_*|\mathcal{D}_n)$  in (3) is then approximated using sampling. For this, we sample  $\mathbf{x}_*$  by evaluating a random sample from  $p(f|\mathcal{D}_n)$  on each grid cell and then selecting the cell with highest value. Given  $\mathbf{x}_*$ , we then approximate  $\mathbb{H}[p(y|\mathcal{D}_n, \mathbf{x}, \mathbf{x}_*)]$  by rejection sampling. We draw samples from  $p(f|\mathcal{D}_n)$  and reject those whose corresponding grid cell with highest value is not  $\mathbf{x}_*$ . Finally, we approximate  $\mathbb{H}[p(y|\mathcal{D}_n, \mathbf{x}, \mathbf{x}_*)]$  by first, adding zero-mean Gaussian noise with variance  $\sigma^2$  to the evaluations at  $\mathbf{x}$  of the functions not rejected during the previous step and second, we estimate the differential entropy of the resulting samples using kernels [1].

Figure 1 shows the objective functions produced by RS, ES and PES for a particular  $\mathcal{D}_n$  with 10 measurements whose locations are selected uniformly at random in  $[0, 1]^2$ . The locations of the collected measurements are displayed with an “x” in the plots. The particular objective function used to generate the measurements in  $\mathcal{D}_n$  is displayed in the left part of Figure 2. The plots in Figure 1 show that the PES approximation to (2) is more similar to the ground truth given by RS than the approximation produced by ES. In this figure we also see a discrepancy between RS and PES at locations near  $\mathbf{x} = (0.572, 0.687)$ . This difference is an artifact of the discretization used in RS. By zooming in and drawing many more samples we would see the same behavior in both plots.

We now evaluate the performance of PES in the task of finding the optimum of synthetic black-box objective functions. For this, we reproduce the within-model comparison experiment described in [9]. In this experiment we optimize objective functions defined in the 2-dimensional unit domain  $\mathcal{X} = [0, 1]^2$ . Each objective function is generated by first sampling 1024 function values from the GP prior assumed by PES, using the same  $\gamma^2$ ,  $\ell_i$  and  $\sigma^2$  as in the previous experiment. The objective function is then given by the resulting GP posterior mean. We generated a total of 1000 objective functions by following this procedure. The left plot in Figure 2 shows an example function.

In these experiments we compared the performance of PES with that of ES [9] and expected improvement (EI) [13], a widely used acquisition function in the Bayesian optimization literature. We again assume that the optimal hyper-parameter values are known to all methods. Predictive performance is then measured in terms of the immediate regret (IR)  $|f(\tilde{\mathbf{x}}_n) - f(\mathbf{x}_*)|$ , where  $\mathbf{x}_*$  is the known location of the global maximum and  $\tilde{\mathbf{x}}_n$  is the recommendation of each algorithm had we stopped at step  $n$ —for all methods this is given by the maximizer of the posterior mean. The right plot in Figure 2 shows the decimal logarithm of the median of the IR obtained by each method across the 1000 *different* objective functions. Confidence bands equal to one standard deviation are obtained using the bootstrap method. Note that while averaging these results is also interesting, corresponding to the expected performance averaged over the prior, here we report the median IR




Figure 2: Left, example of objective functions  $f$ . Right, median of the immediate regret (IR) for the methods PES, ES and EI in the experiments with synthetic objective functions.

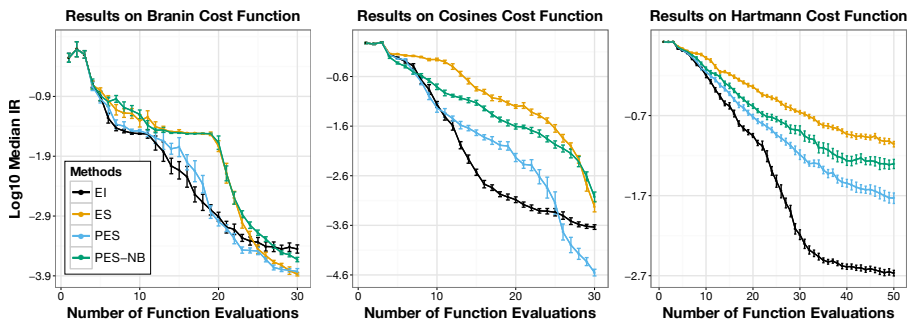


Figure 3: Median of the immediate regret (IR) for the methods EI, ES, PES and PES-NB in the experiments with well-known synthetic benchmark functions.

because the empirical distribution of IR values is very heavy-tailed. In this case, the median is more representative of the exact location of the bulk of the data. These results indicate that the best method in this setting is PES, which significantly outperforms ES and EI. The plot also shows that in this case ES is significantly better than EI.

We perform another series of experiments in which we optimize well-known synthetic benchmark functions including a mixture of cosines [2] and Branin-Hoo (both functions defined in  $[0, 1]^2$ ) as well as the Hartmann-6 (defined in  $[0, 1]^6$ ) [15]. In all instances, we fix the measurement noise to  $\sigma^2 = 10^{-3}$ . For both PES and EI we marginalize the hyperparameters  $\psi$  using the approach described in Section 2.3. ES, by contrast, cannot average its approximation of (2) over the posterior on  $\psi$ . Instead, ES works by fixing  $\psi$  to an estimate of its posterior mean (obtained using slice sampling) [27]. To evaluate the gains produced by the fully Bayesian treatment of  $\psi$  in PES, we also compare with a version of PES (PES-NB) which performs the same non-Bayesian (NB) treatment of  $\psi$  as ES. In PES-NB we use a single fixed hyperparameter as in previous sections with value given by the posterior mean of  $\psi$ . All the methods are initialized with three random measurements collected using latin hypercube sampling [5].

The plots in Figure 3 show the median IR obtained by each method on each function across 250 random initializations. Overall, PES is better than PES-NB and ES. Furthermore, PES-NB is also significantly better than ES in most of the cases. These results show that the fully Bayesian treatment of  $\psi$  in PES is advantageous and that PES can produce better approximations than ES. Note that PES performs better than EI in the Branin and cosines functions, while EI is significantly better on the Hartmann problem. This appears to be due to the fact that entropy-based strategies explore more aggressively which in higher-dimensional spaces takes more iterations. The Hartmann problem, however, is a relatively simple problem and as a result the comparatively more greedy behavior of EI does not result in significant adverse consequences. Note that the synthetic functions optimized in the previous experiment were much more multimodal than the ones considered here.

### 3.1 Experiments with real-world functions

We finally optimize different real-world cost functions. The first one (NNet) returns the predictive accuracy of a neural network on a random train/test partition of the Boston Housing dataset [3].



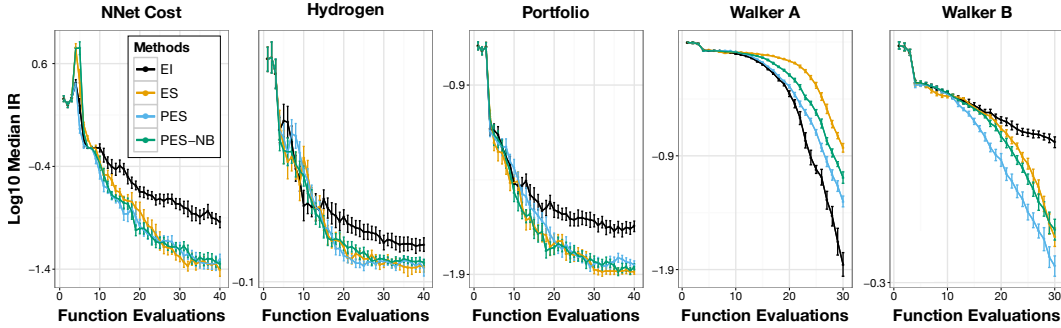


Figure 4: Median of the immediate regret (IR) for the methods PES, PES-NB, ES and EI in the experiments with non-analytic real-world cost functions.

The variables to optimize are the weight-decay parameter and the number of training iterations for the neural network. The second function (Hydrogen) returns the amount of hydrogen production of a particular bacteria in terms of the PH and Nitrogen levels of the growth medium [7]. The third one (Portfolio) returns the ratio of the mean and the standard deviation (the Sharpe ratio) of the 1-year ahead returns generated by simulations from a multivariate time-series model that is adjusted to the daily returns of stocks AXP, BA and HD. The time-series model is formed by univariate GARCH models connected with a Student’s  $t$  copula [12]. These three functions (NNet, Hydrogen and Portfolio) have as domain  $[0, 1]^2$ . Furthermore, in these examples, the ground truth function that we want to optimize is unknown and is only available through noisy measurements. To obtain a ground truth, we approximate each cost function as the predictive distribution of a GP that is adjusted to data sampled from the original function (1000 uniform samples for NNet and Portfolio and all the available data for Hydrogen [7]). Finally, we also consider another real-world function that returns the walking speed of a bipedal robot [30]. This function is defined in  $[0, 1]^8$  and its inputs are the parameters of the robot’s controller. In this case the ground truth function is noiseless and can be exactly evaluated through expensive numerical simulation. We consider two versions of this problem (Walker A) with zero-mean, additive noise of  $\sigma = 0.01$  and (Walker B) with  $\sigma = 0.1$ .

Figure 4 shows the median IR values obtained by each method on each function across 250 random initializations, except in Hydrogen where we used 500 due to its higher level of noise. Overall, PES, ES and PES-NB perform similarly in NNet, Hydrogen and Portfolio. EI performs rather poorly in these first three functions. This method seems to make excessively greedy decisions and fails to explore the search space enough. This strategy seems to be advantageous in Walker A, where EI obtains the best results. By contrast, PES, ES and PES-NB tend to explore more in this latter dataset. This leads to worse results than those of EI. Nevertheless, PES is significantly better than PES-NB and ES in both Walker datasets and better than EI in the noisier Walker B. In this case, the fully Bayesian treatment of hyper-parameters performed by PES produces improvements in performance.

## 4 Conclusions

We have proposed a novel information-theoretic approach for Bayesian optimization. Our method, predictive entropy search (PES), greedily maximizes the amount of one-step information on the location  $\mathbf{x}_*$  of the global maximum using its posterior differential entropy. Since this objective function is intractable, PES approximates the original objective using a reparameterization that measures entropy in the posterior predictive distribution of the function evaluations. PES produces more accurate approximations than Entropy Search (ES), a method based on the original, non-transformed acquisition function. Furthermore, PES can easily marginalize its approximation with respect to the posterior distribution of its hyper-parameters, while ES cannot. Experiments with synthetic and real-world functions show that PES often outperforms ES in terms of immediate regret. In these experiments, we also observe that PES often produces better results than expected improvement (EI), a popular heuristic for Bayesian optimization. EI often seems to make excessively greedy decisions, while PES tends to explore more. As a result, EI seems to perform better for simple objective functions while often getting stuck with noisier objectives or for functions with many modes.

**Acknowledgements** J.M.H.L acknowledges support from the Rafael del Pino Foundation.



## References

- [1] I. Ahmad and P.-E. Lin. A nonparametric estimation of the entropy for absolutely continuous distributions. *IEEE Transactions on Information Theory*, 22(3):372–375, 1976.
- [2] B. S. Anderson, A. W. Moore, and D. Cohn. A nonparametric approach to noisy and costly optimization. In *ICML*, pages 17–24, 2000.
- [3] K. Bache and M. Lichman. UCI machine learning repository, 2013.
- [4] S. Bochner. *Lectures on Fourier integrals*. Princeton University Press, 1959.
- [5] E. Brochu, V. M. Cora, and N. de Freitas. A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. Technical Report UBC TR-2009-23 and arXiv:1012.2599v1, Dept. of Computer Science, University of British Columbia, 2009.
- [6] E. Brochu, N. de Freitas, and A. Ghosh. Active preference learning with discrete choice data. In *NIPS*, pages 409–416, 2007.
- [7] E. H. Burrows, W.-K. Wong, X. Fern, F. W. R. Chaplen, and R. L. Ely. Optimization of ph and nitrogen for enhanced hydrogen production by *synechocystis* sp. pcc 6803 via statistical and machine learning methods. *Biotechnology Progress*, 25(4):1009–1017, 2009.
- [8] O. Chapelle and L. Li. An empirical evaluation of Thompson sampling. In *NIPS*, pages 2249–2257, 2011.
- [9] P. Hennig and C. J. Schuler. Entropy search for information-efficient global optimization. *Journal of Machine Learning Research*, 13, 2012.
- [10] M. W. Hoffman, E. Brochu, and N. de Freitas. Portfolio allocation for Bayesian optimization. In *UAI*, pages 327–336, 2011.
- [11] N. Houlsby, J. M. Hernández-Lobato, F. Huszar, and Z. Ghahramani. Collaborative Gaussian processes for preference learning. In *NIPS*, pages 2096–2104, 2012.
- [12] E. Jondeau and M. Rockinger. The copula-GARCH model of conditional dependencies: An international stock market application. *Journal of international money and finance*, 25(5):827–853, 2006.
- [13] D. R. Jones, M. Schonlau, and W. J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13(4):455–492, 1998.
- [14] H. Kushner. A new method of locating the maximum of an arbitrary multipeak curve in the presence of noise. *Journal of Basic Engineering*, 86, 1964.
- [15] D. Lizotte. *Practical Bayesian Optimization*. PhD thesis, University of Alberta, Canada, 2008.
- [16] D. Lizotte, T. Wang, M. Bowling, and D. Schuurmans. Automatic gait optimization with Gaussian process regression. In *IJCAI*, pages 944–949, 2007.
- [17] D. J. MacKay. Information-based objective functions for active data selection. *Neural Computation*, 4(4):590–604, 1992.
- [18] T. P. Minka. *A family of algorithms for approximate Bayesian inference*. PhD thesis, Massachusetts Institute of Technology, 2001.
- [19] J. Močkus, V. Tiesis, and A. Žilinskas. The application of Bayesian methods for seeking the extremum. In L. Dixon and G. Szego, editors, *Toward Global Optimization*, volume 2. Elsevier, 1978.
- [20] D. M. Negoescu, P. I. Frazier, and W. B. Powell. The knowledge-gradient algorithm for sequencing experiments in drug discovery. *INFORMS Journal on Computing*, 23(3):346–363, 2011.
- [21] A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *NIPS*, pages 1177–1184, 2007.
- [22] C. E. Rasmussen and C. K. Williams. *Gaussian processes for machine learning*. The MIT Press, 2006.
- [23] M. W. Seeger. Bayesian inference and optimal design for the sparse linear model. *Journal of Machine Learning Research*, 9:759–813, 2008.
- [24] J. Snoek, H. Larochelle, and R. P. Adams. Practical Bayesian optimization of machine learning algorithms. In *NIPS*, pages 2960–2968, 2012.
- [25] E. Solak, R. Murray-Smith, W. E. Leithead, D. J. Leith, and C. E. Rasmussen. Derivative observations in Gaussian process models of dynamic systems. In *NIPS*, pages 1057–1064, 2003.
- [26] N. Srinivas, A. Krause, S. M. Kakade, and M. Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *ICML*, pages 1015–1022, 2010.
- [27] J. Vanhatalo, J. Riihimäki, J. Hartikainen, P. Jylänki, V. Tolvanen, and A. Vehtari. Bayesian modeling with Gaussian processes using the matlab toolbox GPstuff (v3.3). *CoRR*, abs/1206.5754, 2012.
- [28] J. Villemonteix, E. Vazquez, and E. Walter. An informational approach to the global optimization of expensive-to-evaluate functions. *Journal of Global Optimization*, 44(4):509–534, 2009.
- [29] Z. Wang, S. Mohamed, and N. de Freitas. Adaptive Hamiltonian and Riemann Monte Carlo samplers. In *ICML*, 2013.
- [30] E. Westervelt and J. Grizzle. *Feedback Control of Dynamic Bipedal Robot Locomotion*. Control and Automation Series. CRC PressINC, 2007.

# Supplementary material for: Predictive Entropy Search for Efficient Global Optimization of Black-box Functions

José Miguel Hernández-Lobato, Matthew W. Hoffman and Zoubin Ghahramani

## A Details on approximating GP sample paths

In this section we give further details about the approach used in Section 2.1 to approximate a GP using random features. These random features can be used to approximate sample paths from the GP posterior. By optimizing these sample paths we obtain posterior samples over the global maxima  $\mathbf{x}_*$ . We derive in more detail the kernel approximation from (5). Formally, the theorem of [4] states

**Theorem 1** (Bochner’s theorem). *A continuous, shift-invariant kernel is positive definite if and only if it is the Fourier transform of a non-negative, finite measure.*

As a result given some kernel  $k(\mathbf{x}, \mathbf{x}') = k(\mathbf{x} - \mathbf{x}', \mathbf{0})$  there must exist an associated density  $s(\mathbf{w})$ , known as its *spectral density*, which is the Fourier dual of  $k$ . This can be written as

$$k(\mathbf{x}, \mathbf{x}') = \int e^{-i\mathbf{w}^\top(\mathbf{x}-\mathbf{x}')} s(\mathbf{w}) d\mathbf{w},$$

$$s(\mathbf{w}) = \frac{1}{(2\pi)^d} \int e^{i\mathbf{w}^\top \tau} k(\tau, \mathbf{0}) d\tau.$$

Further, we can treat this measure as a probability density  $p(\mathbf{w}) = s(\mathbf{w})/\alpha$  where  $\alpha = \int s(\mathbf{w}) d\mathbf{w}$  is the normalizing constant. Consequently, the kernel can be written as

$$k(\mathbf{x}, \mathbf{x}') = \alpha \mathbb{E}_{p(\mathbf{w})}[e^{-i\mathbf{w}^\top(\mathbf{x}-\mathbf{x}')}]$$

and due to the symmetry of  $p(\mathbf{w})$  [see 22] we can write the expectation as

$$= \alpha \mathbb{E}_{p(\mathbf{w})}[\frac{1}{2}(e^{-i\mathbf{w}^\top(\mathbf{x}-\mathbf{x}')} + e^{i\mathbf{w}^\top(\mathbf{x}-\mathbf{x}')})]$$

$$= \alpha \mathbb{E}_{p(\mathbf{w})}[\cos(\mathbf{w}^\top \mathbf{x} - \mathbf{w}^\top \mathbf{x}')].$$

We can then note that  $\int_0^{2\pi} \cos(a + 2b) db = 0$  for any constant offset  $a \in \mathbb{R}$ . As a result, for  $b$  uniformly distributed between 0 and  $2\pi$  we can write

$$= \alpha \mathbb{E}_{p(\mathbf{w})}[\cos(\mathbf{w}^\top \mathbf{x} - \mathbf{w}^\top \mathbf{x}') + \mathbb{E}_{p(b)}[\cos(\mathbf{w}^\top \mathbf{x} + \mathbf{w}^\top \mathbf{x}' + 2b)]]$$

$$= \alpha \mathbb{E}_{p(\mathbf{w}, b)}[\cos(\mathbf{w}^\top \mathbf{x} + b - \mathbf{w}^\top \mathbf{x}' - b) + \cos(\mathbf{w}^\top \mathbf{x} + b + \mathbf{w}^\top \mathbf{x}' + b)]$$

$$= 2\alpha \mathbb{E}_{p(\mathbf{w}, b)}[\cos(\mathbf{w}^\top \mathbf{x} + b) \cos(\mathbf{w}^\top \mathbf{x}' + b)].$$

The last equality can be derived from the sum of angles formula, which leads to the identity:  $2 \cos(x) \cos(y) = \cos(x - y) + \cos(x + y)$ . Finally, we can average over  $m$  weights and phases

$$= \frac{2\alpha}{m} \mathbb{E}_{p(\mathbf{W}, \mathbf{b})}[\cos(\mathbf{W}\mathbf{x} + \mathbf{b})^\top \cos(\mathbf{W}\mathbf{x}' + \mathbf{b})],$$

where  $[\mathbf{W}]_i \sim p(\mathbf{w})$  and  $[\mathbf{b}]_i \sim p(b)$  are stacked versions of the original random variables. The resulting quantity has the same expectation but results in a lower variance estimator. If we let  $\phi(\mathbf{x}) = \sqrt{2\alpha/m} \cos(\mathbf{W}\mathbf{x} + \mathbf{b})$  denote a random  $m$ -dimensional feature generated by this model we can also write the kernel as  $k(\mathbf{x}, \mathbf{x}') = \mathbb{E}_{p(\phi)}[\phi(\mathbf{x})^\top \phi(\mathbf{x}')].$

We now briefly show the equivalence between a Bayesian linear model using random features  $\phi$  and a GP with kernel  $k$ . Consider now a linear model  $f(\mathbf{x}) = \phi(\mathbf{x})^\top \boldsymbol{\theta}$  where  $\boldsymbol{\theta} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  has a standard Gaussian distribution and observations  $\mathcal{D}_n = \{(\mathbf{x}_i, y_i)\}_{i \leq n}$  of the form  $y_i \sim \mathcal{N}(f(\mathbf{x}_i), \sigma^2)$ . The posterior of  $\boldsymbol{\theta}$  given  $(\mathcal{D}_n, \phi)$  is also normal  $\mathcal{N}(\mathbf{m}, \mathbf{V})$  where

$$\mathbf{m} = (\boldsymbol{\Phi}^\top \boldsymbol{\Phi} + \sigma^2 \mathbf{I})^{-1} \boldsymbol{\Phi}^\top \mathbf{y},$$

$$\mathbf{V} = (\boldsymbol{\Phi}^\top \boldsymbol{\Phi} + \sigma^2 \mathbf{I})^{-1} \sigma^2,$$

and where  $[\Phi]_i = \phi(\mathbf{x}_i)$  and  $[\mathbf{y}]_i = y_i$  consist of the stacked features and observations respectively. We can also easily write the predictive distribution over  $f$  evaluated at a test point  $\mathbf{x}$ , which is Gaussian distributed with mean and variance given by

$$\begin{aligned}\mu_n(\mathbf{x}) &= \phi(\mathbf{x})^\top (\Phi^\top \Phi + \sigma^2 \mathbf{I})^{-1} \Phi^\top \mathbf{y}, \\ v_n(\mathbf{x}) &= \phi(\mathbf{x})^\top (\Phi^\top \Phi + \sigma^2 \mathbf{I})^{-1} \phi(\mathbf{x}) \sigma^2.\end{aligned}$$

By a simple application of the matrix-inversion lemma these quantities can be rewritten in terms which only make use of the inner products between features,

$$\begin{aligned}\mu_n(\mathbf{x}) &= \phi(\mathbf{x})^\top \Phi^\top (\Phi \Phi^\top + \sigma^2 \mathbf{I})^{-1} \mathbf{y}_{1:t}, \\ v_n(\mathbf{x}) &= \phi(\mathbf{x})^\top \phi(\mathbf{x}) - \phi(\mathbf{x})^\top \Phi^\top (\Phi \Phi^\top + \sigma^2 \mathbf{I})^{-1} \Phi \phi(\mathbf{x}),\end{aligned}$$

the expectations of which are equivalent to the kernel  $k$  and we obtain the same expressions as that in (1).

## B Details on approximating the predictive variance

We now provide further details on approximating the predictive variance  $v_n(\mathbf{x}|\mathbf{x}_*)$  of inputs  $\mathbf{x}$  given the position of the global optimizer  $\mathbf{x}_*$ . In particular we include all steps omitted in the presentation of Section 2.2.

### B.1 Incorporating the analytic latent constraints (C1.1)

We first turn to the random variables

$$\begin{aligned}\mathbf{z} &= [f(\mathbf{x}_*); \text{diag}[\nabla^2 f(\mathbf{x}_*)]], \\ \mathbf{c} &= [\mathbf{y}_n; \nabla f(\mathbf{x}_*); \text{upper}[\nabla^2 f(\mathbf{x}_*)]] = [\mathbf{y}_n; \mathbf{0}; \mathbf{0}].\end{aligned}$$

Here  $\mathbf{c}$  contains the random variables that we will condition on in order to enforce constraint C1.1. Given the input locations  $\mathbf{x}$  and  $\mathbf{x}_*$  we can construct a kernel matrix  $\mathbf{K}$  containing the covariance evaluated on the stacked vector  $[\mathbf{z}; \mathbf{c}]$ . We again refer to [25] in constructing this matrix which includes derivative observations, the computations of which are tedious but not overly complicated. Note also that the portions of  $\mathbf{K}$  which correspond to  $y_i$  will have an additional  $\sigma^2$  due to the observation noise. Next let  $\mathbf{K}_z$ ,  $\mathbf{K}_c$ , and  $\mathbf{K}_{zc}$  denote the corresponding diagonal and off-diagonal blocks of the kernel matrix. We can now condition on the observed values of  $\mathbf{c}$  to write

$$p(\mathbf{z}|\mathcal{D}_n, \text{C1.1}) = p(\mathbf{z}|\mathbf{c}) = \mathcal{N}(\mathbf{z}|\mathbf{m}_0, \mathbf{V}_0)$$

where  $\mathbf{m}_0 = \mathbf{K}_{zc} \mathbf{K}_c^{-1} \mathbf{c}$  and  $\mathbf{V}_0 = \mathbf{K}_z - \mathbf{K}_{zc} \mathbf{K}_c^{-1} \mathbf{K}_{zc}^\top$ .

### B.2 Incorporating the non-analytic latent constraints (C1.2 and C2)

The additional constraints C1.2 and C2 can be introduced explicitly as in (6), which takes the form of a single Gaussian factor and  $d + 1$  non-Gaussian factors

$$p(\mathbf{z}|\mathcal{D}_n, \text{C1}, \text{C2}) \propto \mathcal{N}(\mathbf{z}|\mathbf{m}_0, \mathbf{V}_0) \left[ \prod_{i=1}^{d+1} t_i(z_i) \right].$$

We approximate this distribution using a single multivariate Gaussian  $q(\mathbf{z})$  where each non-Gaussian factor is replaced by a Gaussian approximation  $\tilde{t}_i(z_i) = \mathcal{N}(z_i; \tilde{m}_i, \tilde{v}_i)$  such that

$$q(\mathbf{z}) = \mathcal{N}(\mathbf{z}|\mathbf{m}, \mathbf{V}) \propto \mathcal{N}(\mathbf{z}|\mathbf{m}_0, \mathbf{V}_0) \left[ \prod_{i=1}^{d+1} \mathcal{N}(z_i; \tilde{m}_i, \tilde{v}_i) \right]$$

where this approximation is parameterized by  $\mathbf{m} = \mathbf{V}[\tilde{\mathbf{V}}^{-1} \tilde{\mathbf{m}} + \mathbf{V}_0^{-1} \mathbf{m}_0]$  and  $\mathbf{V} = (\tilde{\mathbf{V}}^{-1} + \mathbf{V}_0^{-1})^{-1}$ . The parameters of the approximate factors are combined to form the vector  $[\tilde{\mathbf{m}}]_i = \tilde{m}_i$  and the diagonal matrix  $[\tilde{\mathbf{V}}]_{ii} = \tilde{v}_i$ .

To compute the approximate factors we use expectation propagation (EP). EP is a procedure that starts from some initial values for the approximate factors  $(\tilde{m}_i, \tilde{v}_i)$  and iteratively refines these quantities; here we initialize  $\tilde{m}_i = 0$  and  $\tilde{v}_i = \infty$  which corresponds to  $\mathbf{m} = \mathbf{m}_0$  and  $\mathbf{V} = \mathbf{V}_0$ . At each iteration, for every factor  $i$ , we remove the contribution of the  $i$ th approximate factor to form the *cavity* distribution  $q_{\setminus i}(\mathbf{z}) \propto q(\mathbf{z})/\tilde{t}_i(z_i)$ . Given the independent factors we consider here we can focus on each individual component  $q_{\setminus i}(z_i)$  separately with mean and variance

$$\begin{aligned}\bar{m}_i &= \bar{v}_i(m_i/v_{ii} - \tilde{m}_i/\tilde{v}_i), \\ \bar{v}_i &= (v_{ii}^{-1} - \tilde{v}_i^{-1})^{-1}.\end{aligned}$$

Let  $\hat{q}(z_i) \propto q_{\setminus i}(z_i)t_i(z_i)$  denote the *tilted* distribution where the  $i$ th approximate factor has been replaced by the corresponding real factor. EP proceeds by finding the approximation  $q_i$  that minimizes the KL-divergence  $D[\hat{q}_i||q_i]$  where  $q_i$  is restricted to be Gaussian. This amounts to matching the first two moments. Finally, by removing the influence of the cavity distribution and setting  $\tilde{t}_i(z_i) \propto q_i(z_i)/q_{\setminus i}(z_i)$  we can update the approximate factors. This can be performed using the same procedure which forms the cavity distribution.

For both sets of constraints used in this work the moments can easily be obtained by computing the log normalizing constant  $\bar{Z}_i = \log \int \mathcal{N}(z_i|\bar{m}_i, \bar{v}_i) t_i(z_i) dz_i$  and using the following identities:

$$\mathbb{E}_{\hat{q}}[z_i] = \bar{m}_i + \bar{v}_i \frac{\partial \bar{Z}_i}{\partial \bar{m}_i}, \quad \text{Var}_{\hat{q}}[z_i] = \bar{v}_i - \bar{v}_i^2 \left( \left[ \frac{\partial \bar{Z}_i}{\partial \bar{m}_i} \right]^2 - 2 \frac{\partial \bar{Z}_i}{\partial \bar{v}_i} \right). \quad (11)$$

We first show the updates for the parameters  $(\tilde{m}_i, \tilde{v}_i)$  of the factors corresponding to constraints on the diagonal Hessian, i.e. where  $t_i(z_i) = \mathbb{I}[z_i < 0]$ . Given the moments of the tilted distribution in (11), we can remove the contribution of the cavity distribution as above and write

$$\begin{aligned}\tilde{m}_i &\leftarrow \bar{m}_i + \kappa^{-1}, & \text{where } \alpha &= -\frac{\bar{m}_i}{\sqrt{\bar{v}_i}}, \\ \tilde{v}_i &\leftarrow \beta^{-1} - \bar{v}_i, & \beta &= \frac{\phi(\alpha)}{\Phi(\alpha)} \left[ \frac{\phi(\alpha)}{\Phi(\alpha)} + \alpha \right] \frac{1}{\bar{v}_i}, \\ & & \kappa &= -\left[ \frac{\phi(\alpha)}{\Phi(\alpha)} + \alpha \right] \frac{1}{\sqrt{\bar{v}_i}},\end{aligned}$$

where  $\phi$  and  $\Phi$  are the standard Gaussian density and cumulative distribution functions, respectively. For the final soft-maximum constraint,  $\Phi((z_i - y_{\max})/\sigma)$ , the moments can be calculated in a similar fashion. Using the same procedure as above we arrive at very similar updates:

$$\begin{aligned}\tilde{m}_i &\leftarrow \bar{m}_i + \kappa^{-1}, & \text{where } \alpha &= \frac{\bar{m}_i - y_{\max}}{\sqrt{\bar{v}_i + \sigma^2}}, \\ \tilde{v}_i &\leftarrow \beta^{-1} - \bar{v}_i, & \beta &= \frac{\phi(\alpha)}{\Phi(\alpha)} \left[ \frac{\phi(\alpha)}{\Phi(\alpha)} + \alpha \right] \frac{1}{\bar{v}_i + \sigma^2}, \\ & & \kappa &= \left[ \frac{\phi(\alpha)}{\Phi(\alpha)} + \alpha \right] \frac{1}{\sqrt{\bar{v}_i + \sigma^2}}.\end{aligned}$$

### B.3 Incorporating the prediction constraint (C3)

Given some test input  $\mathbf{x}$  we now turn to the problem of making predictions about  $f(\mathbf{x})$ . We again note that both the ‘‘prior’’ terms  $\mathbf{m}_0, \mathbf{V}_0$  and the EP factors,  $\tilde{\mathbf{m}}$  and  $\tilde{\mathbf{V}}$ , are independent of  $\mathbf{x}$  and can be precomputed once for later use at prediction time.

Let  $\mathbf{f} = [f(\mathbf{x}); f(\mathbf{x}_*)]$  be a vector given by the concatenation of the latent function at  $\mathbf{x}$  and  $\mathbf{x}_*$ . The distribution for  $\mathbf{f}$  given the first two constraints can be written as

$$p(\mathbf{f}|\mathcal{D}_n, C1, C2) \approx \int p(\mathbf{f}|\mathbf{z}, \mathbf{c}) q(\mathbf{z}) d\mathbf{z} = \mathcal{N}(\mathbf{f}|\mathbf{m}_f, \mathbf{V}_f). \quad (12)$$

By writing  $p(\mathbf{f}|\mathbf{z}, \mathbf{c})$  above we are assuming that  $\mathbf{f}$  is independent of C1.2 and C2 given  $\mathbf{z}$  and as a result the above is simply an integral over the product of two Gaussians. Let  $\mathbf{K}_{\dagger}$  be the cross-covariance matrix evaluated between  $\mathbf{f}$  and  $[\mathbf{z}; \mathbf{c}]$  and  $\mathbf{K}_f$  the covariance matrix associated with  $\mathbf{f}$ .

The posterior above will then be Gaussian with mean and variance

$$\begin{aligned}\mathbf{m}_f &= \mathbf{K}_\dagger[\mathbf{K} + \widetilde{\mathbf{W}}]^{-1}[\mathbf{c}; \widetilde{\mathbf{m}}] \\ \mathbf{V}_f &= \mathbf{K}_f - \mathbf{K}_\dagger[\mathbf{K} + \widetilde{\mathbf{W}}]^{-1}\mathbf{K}_\dagger^\top,\end{aligned}$$

where  $\widetilde{\mathbf{W}}$  is a block-diagonal matrix where the first block is zero and the second is  $\widetilde{\mathbf{V}}$  (note this matrix is also diagonal since  $\widetilde{\mathbf{V}}$  is diagonal).

The constraint that  $f(\mathbf{x}) < f(\mathbf{x}_*)$  can be incorporated by introducing an additional factor ensuring that  $\mathbf{c}^\top \mathbf{f}$  is positive for  $\mathbf{c} = [-1; 1]$ . The resulting density has a log normalizing constant given by

$$\begin{aligned}Z &= \log \int \mathcal{N}(\mathbf{f}|\mathbf{m}_f, \mathbf{V}_f) \mathbb{I}[\mathbf{c}^\top \mathbf{f} > 0] d\mathbf{f} \\ &= \log \int \mathcal{N}(a|m, v) \mathbb{I}[a > 0] da = \log \Phi(m/\sqrt{v}).\end{aligned}$$

where  $m = \mathbf{c}^\top \mathbf{m}_f$  and  $v = \mathbf{c}^\top \mathbf{V}_f \mathbf{c}$ . The derivatives of  $Z$  with respect to  $\mathbf{m}_f$  and  $\mathbf{V}_f$  can then be taken and a multivariate form of (11) can be used to obtain the first two moments [18]. In particular, the second moment is given by

$$\text{Var}[\mathbf{f}] = \mathbf{V}_f - v^{-1}\beta(\alpha + \beta)(\mathbf{V}_f \mathbf{c} \mathbf{c}^\top \mathbf{V}_f),$$

where  $\alpha = m/\sqrt{v}$  and  $\beta = \phi(\alpha)/\Phi(\alpha)$ . Finally, the first diagonal component of the resulting matrix corresponds to the marginal variance of  $f(\mathbf{x})$ , which can be simplified as in (9). This marginal variance can be used as the variance parameter in a Gaussian approximation to  $p(f(\mathbf{x})|\mathcal{D}_n, \mathbf{x}_*)$ . For the PES acquisition function we only need this variance parameter since the entropy of a Gaussian does not depend on its first moment.