
Max-value Entropy Search for Efficient Bayesian Optimization

Zi Wang¹ Stefanie Jegelka¹

Abstract

Entropy Search (ES) and Predictive Entropy Search (PES) are popular and empirically successful Bayesian Optimization techniques. Both rely on a compelling information-theoretic motivation, and maximize the information gained about the $\arg \max$ of the unknown function; yet, both are plagued by the expensive computation for estimating entropies. We propose a new criterion, Max-value Entropy Search (MES), that instead uses the information about the maximum function value. We show relations of MES to other Bayesian optimization methods, and establish a regret bound. We observe that MES maintains or improves the good empirical performance of ES/PES, while tremendously lightening the computational burden. In particular, MES is much more robust to the number of samples used for computing the entropy, and hence more efficient for higher dimensional problems.

1. Introduction

Bayesian optimization (BO) has become a popular and effective way for black-box optimization of nonconvex, expensive functions in robotics, machine learning, computer vision, and many other areas of science and engineering (Brochu et al., 2009; Calandra et al., 2014; Krause & Ong, 2011; Lizotte et al., 2007; Snoek et al., 2012; Thornton et al., 2013; Wang et al., 2017). In BO, a prior is posed on the (unknown) objective function, and the uncertainty given by the associated posterior is the basis for an acquisition function that guides the selection of the next point to query the function. The selection of queries and hence the acquisition function is critical for the success of the method.

Different BO techniques differ in this acquisition function.

¹Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Massachusetts, USA. Correspondence to: Zi Wang <ziw@csail.mit.edu>, Stefanie Jegelka <stefje@csail.mit.edu>.

Among the most popular ones range the Gaussian process upper confidence bound (GP-UCB) (Auer, 2002; Srinivas et al., 2010), probability of improvement (PI) (Kushner, 1964), and expected improvement (EI) (Moćkus, 1974). Particularly successful recent additions are entropy search (ES) (Hennig & Schuler, 2012) and predictive entropy search (PES) (Hernández-Lobato et al., 2014), which aim to maximize the mutual information between the queried points and the location of the global optimum.

ES and PES are effective in the sense that they are query-efficient and identify a good point within competitively few iterations, but determining the next query point involves very expensive computations. As a result, these methods are most useful if the black-box function requires a lot of effort to evaluate, and are relatively slow otherwise. Moreover, they rely on estimating the entropy of the $\arg \max$ of the function. In high dimensions, this estimation demands a large number of samples from the input space, which can quickly become inefficient.

We propose a twist to the viewpoint of ES and PES that retains the information-theoretic motivation and empirically successful query-efficiency of those methods, but at a much reduced computational cost. The key insight is to replace the uncertainty about the $\arg \max$ with the uncertainty about the maximum function value. As a result, we refer to our new method as *Max-value Entropy Search (MES)*. As opposed to the $\arg \max$, the maximum function value lives in a one-dimensional space, which greatly facilitates the estimation of the mutual information via sampling. We explore two strategies to make the entropy estimation efficient: an approximation by a Gumbel distribution, and a Monte Carlo approach that uses random features.

Our contributions are as follows: (1) MES, a variant of the entropy search methods, which enjoys efficient computation and simple implementation; (2) an intuitive analysis which establishes the first connection between ES/PES and the previously proposed criteria GP-UCB, PI and EST (Wang et al., 2016), where the bridge is formed by MES; (3) a regret bound for a variant of MES, which, to our knowledge, is the first regret bound established for any variant of the entropy search methods; (4) an extension of MES to the high dimensional settings via additive Gaussian processes; and (5) empirical evaluations which demon-

strate that MES identifies good points as quickly or better than ES/PES, but is much more efficient and robust in estimating the mutual information, and therefore much faster than its input-space counterparts.

After acceptance of this work, we learned that Hoffman & Ghahramani (2015) independently arrived at the acquisition function in Eq. (5). Yet, our approximation (Eq. (6)) is different, and hence the actual acquisition function we evaluate and analyze is different.

2. Background

Our goal is to maximize a black-box function $f : \mathfrak{X} \rightarrow \mathbb{R}$ where $\mathfrak{X} \subset \mathbb{R}^d$ and \mathfrak{X} is compact. At time step t , we select point \mathbf{x}_t and observe a possibly noisy function evaluation $y_t = f(\mathbf{x}_t) + \epsilon_t$, where $\epsilon_t \sim \mathcal{N}(0, \sigma^2)$ are i.i.d. Gaussian variables. We use Gaussian processes (Rasmussen & Williams, 2006) to build a probabilistic model of the black-box function to be optimized. For high dimensional cases, we use a variant of the additive Gaussian process (Duvenaud et al., 2011; Kandasamy et al., 2015). For completeness, we here introduce some basics of GP and add-GP.

2.1. Gaussian Processes

Gaussian processes (GPs) are distributions over functions, and popular priors for Bayesian nonparametric regression. In a GP, any finite set of function values has a multivariate Gaussian distribution. A Gaussian process $GP(\mu, k)$ is fully specified by a mean function $\mu(\mathbf{x})$ and covariance (kernel) function $k(\mathbf{x}, \mathbf{x}')$. Let f be a function sampled from $GP(\mu, k)$. Given the observations $D_t = \{(\mathbf{x}_\tau, y_\tau)\}_{\tau=1}^t$, we obtain the posterior mean $\mu_t(\mathbf{x}) = \mathbf{k}_t(\mathbf{x})^\top (\mathbf{K}_t + \sigma^2 \mathbf{I})^{-1} \mathbf{y}_t$ and posterior covariance $k_t(\mathbf{x}, \mathbf{x}') = k(\mathbf{x}, \mathbf{x}') - \mathbf{k}_t(\mathbf{x})^\top (\mathbf{K}_t + \sigma^2 \mathbf{I})^{-1} \mathbf{k}_t(\mathbf{x}')$ of the function via the kernel matrix $\mathbf{K}_t = [k(\mathbf{x}_i, \mathbf{x}_j)]_{\mathbf{x}_i, \mathbf{x}_j \in D_t}$ and $\mathbf{k}_t(\mathbf{x}) = [k(\mathbf{x}_i, \mathbf{x})]_{\mathbf{x}_i \in D_t}$ (Rasmussen & Williams, 2006). The posterior variance is $\sigma_t^2(\mathbf{x}) = k_t(\mathbf{x}, \mathbf{x})$.

2.2. Additive Gaussian Processes

Additive Gaussian processes (add-GP) were proposed in (Duvenaud et al., 2011), and analyzed in the BO setting in (Kandasamy et al., 2015). Following the latter, we assume that the function f is a sum of independent functions sampled from Gaussian processes that are active on disjoint sets A_m of input dimensions. Precisely, $f(x) = \sum_{m=1}^M f^{(m)}(x^{A_m})$, with $A_i \cap A_j = \emptyset$ for all $i \neq j$, $|\cup_{i=1}^M A_i| = d$, and $f^{(m)} \sim GP(\mu^{(m)}, k^{(m)})$, for all $m \leq M$ ($M \leq d < \infty$). As a result of this decomposition, the function f is distributed according to $GP(\sum_{m=1}^M \mu^{(m)}, \sum_{m=1}^M k^{(m)})$. Given a set of noisy observations $D_t = \{(\mathbf{x}_\tau, y_\tau)\}_{\tau=1}^t$ where $y_\tau \sim \mathcal{N}(f(x_\tau), \sigma^2)$, the posterior mean and

covariance of the function component $f^{(m)}$ can be inferred as $\mu_t^{(m)}(\mathbf{x}) = \mathbf{k}_t^{(m)}(\mathbf{x})^\top (\mathbf{K}_t + \sigma^2 \mathbf{I})^{-1} \mathbf{y}_t$ and $k_t^{(m)}(\mathbf{x}, \mathbf{x}') = k^{(m)}(\mathbf{x}, \mathbf{x}') - \mathbf{k}_t^{(m)}(\mathbf{x})^\top (\mathbf{K}_t + \sigma^2 \mathbf{I})^{-1} \mathbf{k}_t^{(m)}(\mathbf{x}')$, where $\mathbf{k}_t^{(m)}(\mathbf{x}) = [k^{(m)}(\mathbf{x}_i, \mathbf{x})]_{\mathbf{x}_i \in D_t}$ and $\mathbf{K}_t = [\sum_{m=1}^M k^{(m)}(\mathbf{x}_i, \mathbf{x}_j)]_{\mathbf{x}_i, \mathbf{x}_j \in D_t}$. For simplicity, we use the shorthand $k^{(m)}(\mathbf{x}, \mathbf{x}') = k^{(m)}(\mathbf{x}^{A_m}, \mathbf{x}'^{A_m})$.

2.3. Evaluation Criteria

We use two types of evaluation criteria for BO, *simple regret* and *inference regret*. In each iteration, we choose to evaluate one input \mathbf{x}_t to “learn” where the arg max of the function is. The simple regret $r_T = \max_{\mathbf{x} \in \mathfrak{X}} f(\mathbf{x}) - \max_{t \in [1, T]} f(\mathbf{x}_t)$ measures the value of the best queried point so far. After all queries, we may infer an arg max of the function, which is usually chosen as $\tilde{\mathbf{x}}_T = \arg \max_{\mathbf{x} \in \mathfrak{X}} \mu_T(\mathbf{x})$ (Hennig & Schuler, 2012; Hernández-Lobato et al., 2014). We denote the inference regret as $R_T = \max_{\mathbf{x} \in \mathfrak{X}} f(\mathbf{x}) - f(\tilde{\mathbf{x}}_T)$ which characterizes how satisfying our inference of the arg max is.

3. Max-value Entropy Search

Entropy search methods use an information-theoretic perspective to select where to evaluate. They find a query point that maximizes the information about the location $\mathbf{x}_* = \arg \max_{\mathbf{x} \in \mathfrak{X}} f(\mathbf{x})$ whose value $y_* = f(\mathbf{x}_*)$ achieves the global maximum of the function f . Using the negative differential entropy of $p(\mathbf{x}_* | D_t)$ to characterize the uncertainty about \mathbf{x}_* , ES and PES use the acquisition functions

$$\alpha_t(x) = I(\{\mathbf{x}, y\}; \mathbf{x}_* | D_t) \quad (1)$$

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

$$= H(p(y | D_t, \mathbf{x})) - \mathbb{E}[H(p(y | D_t, \mathbf{x}, \mathbf{x}_*))]. \quad (3)$$

ES uses formulation (2), in which the expectation is over $p(y | D_t, \mathbf{x})$, while PES uses the equivalent, symmetric formulation (3), where the expectation is over $p(\mathbf{x}_* | D_t)$. Unfortunately, both $p(\mathbf{x}_* | D_t)$ and its entropy is analytically intractable and have to be approximated via expensive computations. Moreover, the optimum may not be unique, adding further complexity to this distribution.

We follow the same information-theoretic idea but propose a much cheaper and more robust objective to compute. Instead of measuring the information about the argmax \mathbf{x}_* , we use the information about the *maximum value* $y_* = f(\mathbf{x}_*)$. Our acquisition function is the gain in mutual information between the maximum y_* and the next point we query, which can be approximated analytically by evaluating the entropy of the predictive distribution:

$$\alpha_t(x) = I(\{\mathbf{x}, y\}; y_* | D_t) \quad (4)$$

$$= H(p(y | D_t, \mathbf{x})) - \mathbb{E}[H(p(y | D_t, \mathbf{x}, y_*))] \quad (5)$$

$$\approx \frac{1}{K} \sum_{y_* \in Y_*} \left[\frac{\gamma_{y_*}(\mathbf{x}) \psi(\gamma_{y_*}(\mathbf{x}))}{2\Psi(\gamma_{y_*}(\mathbf{x}))} - \log(\Psi(\gamma_{y_*}(\mathbf{x}))) \right] \quad (6)$$

where ψ is the probability density function and Ψ the cumulative density function of a normal distribution, and $\gamma_{y_*}(\mathbf{x}) = \frac{y_* - \mu_t(\mathbf{x})}{\sigma_t(\mathbf{x})}$. The expectation in Eq. (5) is over $p(y_*|D_n)$, which is approximated using Monte Carlo estimation by sampling a set of K function maxima. Notice that the probability in the first term $p(y|D_t, \mathbf{x})$ is a Gaussian distribution with mean $\mu_t(\mathbf{x})$ and variance $k_t(\mathbf{x}, \mathbf{x})$. The probability in the second term $p(y|D_n, \mathbf{x}, y_*)$ is a truncated Gaussian distribution: given y_* , the distribution of y needs to satisfy $y < y_*$. Importantly, while ES and PES rely on the expensive, d -dimensional distribution $p(\mathbf{x}_*|D_t)$, here, we use the one-dimensional $p(y_*|D_n)$, which is computationally much easier.

It may not be immediately intuitive that the *value* should bear sufficient information for a good search strategy. Yet, the empirical results in Section 5 will demonstrate that this strategy is typically at least as good as ES/PES. From a formal perspective, Wang et al. (2016) showed how an estimate of the maximum value implies a good search strategy (EST). Indeed, Lemma 3.1 will make the relation between EST and a simpler, degenerate version of MES explicit.

Hence, it remains to determine how to sample y_* . We propose two strategies: (1) sampling from an approximation via a Gumbel distribution; and (2) sampling functions from the posterior Gaussian distribution and maximizing the functions to obtain samples of y_* . We present the MES algorithm in Alg. 1.

3.1. Gumbel Sampling

The marginal distribution of $f(x)$ for any x is a one-dimensional Gaussian, and hence the distribution of y^* may be viewed as the maximum of an infinite collection of dependent Gaussian random variables. Since this distribution is difficult to compute, we make two simplifications. First, we replace the continuous set \mathcal{X} by a discrete (finite), dense subset $\hat{\mathcal{X}}$ of representative points. If we select $\hat{\mathcal{X}}$ to be an ϵ -cover of \mathcal{X} and the function f is Lipschitz continuous with constant L , then we obtain a valid upper bound on $f(\mathcal{X})$ by adding ϵL to any upper bound on $f(\hat{\mathcal{X}})$.

Second, we use a ‘‘mean field’’ approximation and treat the function values at the points in $\hat{\mathcal{X}}$ as independent. This approximation tends to over-estimate the maximum; this follows from Slepian’s lemma if $k(x, x') \geq 0$. Such upper bounds still lead to optimization strategies with vanishing regret, whereas lower bounds may not (Wang et al., 2016).

We sample from the approximation $\hat{p}(y_*|D_n)$ via its cumulative distribution function (CDF) $\widehat{\text{Pr}}[y_* < z] = \prod_{\mathbf{x} \in \hat{\mathcal{X}}} \Psi(\gamma_z(\mathbf{x}))$. That means we sample r uniformly from

Algorithm 1 Max-value Entropy Search (MES)

```

1: function MES ( $f, D_0$ )
2:   for  $t = 1, \dots, T$  do
3:      $\alpha_{t-1}(\cdot) \leftarrow \text{APPROX-MI}(D_{t-1})$ 
4:      $\mathbf{x}_t \leftarrow \arg \max_{\mathbf{x} \in \mathcal{X}} \alpha_{t-1}(\mathbf{x})$ 
5:      $y_t \leftarrow f(\mathbf{x}_t) + \epsilon_t, \epsilon_t \sim \mathcal{N}(0, \sigma^2)$ 
6:      $\mathcal{D}_t \leftarrow D_{t-1} \cup \{\mathbf{x}_t, y_t\}$ 
7:   end for
8: end function

9: function Approx-MI ( $D_t$ )
10:  if Sample with Gumbel then
11:    approximate  $\text{Pr}[\hat{y}_* < y]$  with  $\mathcal{G}(a, b)$ 
12:    sample a  $K$ -length vector  $\mathbf{r} \sim \text{Unif}([0, 1])$ 
13:     $\mathbf{y}_* \leftarrow a - b \log(-\log \mathbf{r})$ 
14:  else
15:    for  $i = 1, \dots, K$  do
16:      sample  $\tilde{f} \sim GP(\mu_t, k_t | D_t)$ 
17:       $y_{*(i)} \leftarrow \max_{\mathbf{x} \in \mathcal{X}} \tilde{f}(\mathbf{x})$ 
18:    end for
19:     $\mathbf{y}_* \leftarrow [y_{*(i)}]_{i=1}^K$ 
20:  end if
21:  return  $\alpha_t(\cdot)$  in Eq. (6)
22: end function

```

$[0, 1]$ and find z such that $\text{Pr}[y_* < z] = r$. A binary search for z to accuracy δ requires $O(\log \frac{1}{\delta})$ queries to the CDF, and each query takes $O(|\hat{\mathcal{X}}|) \approx O(n^d)$ time, so we obtain an overall time of $O(M|\hat{\mathcal{X}}| \log \frac{1}{\delta})$ for drawing M samples.

To sample more efficiently, we propose a $O(M + |\hat{\mathcal{X}}| \log \frac{1}{\delta})$ -time strategy, by approximating the CDF by a Gumbel distribution: $\widehat{\text{Pr}}[y_* < z] \approx \mathcal{G}(a, b) = e^{-e^{-\frac{z-a}{b}}}$. This choice is motivated by the Fisher-Tippett-Gnedenko theorem (Fisher, 1930), which states that the maximum of a set of i.i.d. Gaussian variables is asymptotically described by a Gumbel distribution (see the appendix for further details). This does not in general extend to non-i.i.d. Gaussian variables, but we nevertheless observe that in practice, this approach yields a good and fast approximation.

We sample from the Gumbel distribution via the Gumbel quantile function: we sample r uniformly from $[0, 1]$, and let the sample be $y = \mathcal{G}^{-1}(a, b) = a - b \log(-\log r)$. We set the appropriate Gumbel distribution parameters a and b by percentile matching and solve the two-variable linear equations $a - b \log(-\log r_1) = y_1$ and $a - b \log(-\log r_2) = y_2$, where $\text{Pr}[y_* < y_1] = r_1$ and $\text{Pr}[y_* < y_2] = r_2$. In practice, we use $r_1 = 0.25$ and $r_2 = 0.75$ so that the scale of the approximated Gumbel distribution is proportional to the interquartile range of the CDF $\widehat{\text{Pr}}[y_* < z]$.

3.2. Sampling y_* via Posterior Functions

For an alternative sampling strategy we follow (Hernández-Lobato et al., 2014): we draw functions from the posterior GP and then maximize each of the sampled functions. Given the observations $D_t = \{(\mathbf{x}_\tau, y_\tau)_{\tau=1}^t\}$, we can approximate the posterior Gaussian process using a 1-hidden-layer neural network $\tilde{f}(\mathbf{x}) = \mathbf{a}_t^\top \phi(\mathbf{x})$ where $\phi(\mathbf{x}) \in \mathbb{R}^D$ is a vector of feature functions (Neal, 1996; Rahimi et al., 2007) and the Gaussian weight $\mathbf{a}_t \in \mathbb{R}^D$ is distributed according to a multivariate Gaussian $\mathcal{N}(\boldsymbol{\nu}_t, \boldsymbol{\Sigma}_t)$.

Computing $\phi(\mathbf{x})$. By Bochner’s theorem (Rudin, 2011), the Fourier transform \hat{k} of a continuous and translation-invariant kernel k is guaranteed to be a probability distribution. Hence we can write the kernel of the GP to be $k(\mathbf{x}, \mathbf{x}') = \mathbb{E}_{\omega \sim \hat{k}(\omega)} [e^{i\omega^\top(\mathbf{x}-\mathbf{x}')}] = \mathbb{E}_{c \sim U[0, 2\pi]} \mathbb{E}_{\hat{k}} [2 \cos(\omega^\top \mathbf{x} + c) \cos(\omega^\top \mathbf{x}' + c)]$ and approximate the expectation by $k(\mathbf{x}, \mathbf{x}') \approx \phi^\top(\mathbf{x})\phi(\mathbf{x}')$ where $\phi_i(\mathbf{x}) = \sqrt{\frac{2}{D}} \cos(\omega_i^\top \mathbf{x} + c_i)$, $\omega_i \sim \hat{k}(\omega)$, and $c_i \sim U[0, 2\pi]$ for $i = 1, \dots, D$.

Computing $\boldsymbol{\nu}_t, \boldsymbol{\Sigma}_t$. By writing the GP as a random linear combination of feature functions $\mathbf{a}_t^\top \phi(\mathbf{x})$, we are defining the mean and covariance of the GP to be $\mu_t(\mathbf{x}) = \boldsymbol{\nu}_t^\top \phi(\mathbf{x})$ and $k(\mathbf{x}, \mathbf{x}') = \phi(\mathbf{x})^\top \boldsymbol{\Sigma}_t \phi(\mathbf{x}')$. Let $Z = [z_1, \dots, z_t] \in \mathbb{R}^{D \times t}$, where $z_\tau := \phi(\mathbf{x}_\tau) \in \mathbb{R}^D$. The GP posterior mean and covariance in Section 2.1 become $\mu_t(\mathbf{x}) = z^\top Z(Z^\top Z + \sigma^2 \mathbf{I})^{-1} \mathbf{y}_t$ and $k_t(\mathbf{x}, \mathbf{x}') = z^\top z' - z^\top Z(Z^\top Z + \sigma^2 \mathbf{I})^{-1} Z^\top z'$. Because $Z(Z^\top Z + \sigma^2 \mathbf{I})^{-1} = (ZZ^\top + \sigma^2 \mathbf{I})^{-1} Z$, we can simplify the above equations and obtain $\boldsymbol{\nu}_t = \sigma^{-2} \boldsymbol{\Sigma}_t Z_t^\top \mathbf{y}_t$ and $\boldsymbol{\Sigma}_t = (ZZ^\top \sigma^{-2} + \mathbf{I})^{-1}$.

To sample a function from this random 1-hidden-layer neural network, we sample $\tilde{\mathbf{a}}$ from $\mathcal{N}(\boldsymbol{\nu}_t, \boldsymbol{\Sigma}_t)$ and construct the sampled function $\tilde{f} = \tilde{\mathbf{a}}^\top \phi(\mathbf{x})$. Then we optimize \tilde{f} with respect to its input to get a sample of the maximum of the function $\max_{\mathbf{x} \in \mathcal{X}} \tilde{f}(\mathbf{x})$.

3.3. Relation to Other BO Methods

As a side effect, our new acquisition function draws connections between ES/PES and other popular BO methods. The connection between MES and ES/PES follows from the information-theoretic viewpoint; the following lemma makes the connections to other methods explicit.

Lemma 3.1. *The following methods are equivalent:*

1. MES, where we only use a single sample y_* for $\alpha_t(\mathbf{x})$;
2. EST with $m = y_*$;
3. GP-UCB with $\beta^{\frac{1}{2}} = \min_{\mathbf{x} \in \mathcal{X}} \frac{y_* - \mu_t(\mathbf{x})}{\sigma_t(\mathbf{x})}$;
4. PI with $\theta = y_*$.

This equivalence no longer holds if we use $M > 1$ samples of y_* in MES.

Proof. The equivalence among 2,3,4 is stated in Lemma 2.1 in (Wang et al., 2016). What remains to be shown is the equivalence between 1 and 2. When using a single y_* in MES, the next point to evaluate is chosen by maximizing $\alpha_t(\mathbf{x}) = \gamma_{y_*}(\mathbf{x}) \frac{\psi(\gamma_{y_*}(\mathbf{x}))}{2\Psi(\gamma_{y_*}(\mathbf{x}))} - \log(\Psi(\gamma_{y_*}(\mathbf{x})))$ and $\gamma_{y_*} = \frac{y_* - \mu_t(\mathbf{x})}{\sigma_t(\mathbf{x})}$. For EST with $m = y_*$, the next point to evaluate is chosen by minimizing $\gamma_{y_*}(\mathbf{x})$. Let us define a function $g(u) = u \frac{\psi(u)}{2\Psi(u)} - \log(\Psi(u))$. Clearly, $\alpha_t(\mathbf{x}) = g(\gamma_{y_*}(\mathbf{x}))$. Because $g(u)$ is a monotonically decreasing function, maximizing $g(\gamma_{y_*}(\mathbf{x}))$ is equivalent to minimizing $\gamma_{y_*}(\mathbf{x})$. Hence 1 and 2 are equivalent. \square

3.4. Regret Bound

The connection with EST directly leads to a bound on the simple regret of MES, when using only one sample of y_* . We prove Theorem 3.2 in the appendix.

Theorem 3.2 (Simple Regret Bound). *Let F be the cumulative probability distribution for the maximum of any function f sampled from $GP(\mu, k)$ over the compact search space $\mathcal{X} \subset \mathbb{R}^d$, where $k(\mathbf{x}, \mathbf{x}') \leq 1, \forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}$. Let $f_* = \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$ and $w = F(f_*) \in (0, 1)$, and assume the observation noise is iid $\mathcal{N}(0, \sigma)$. If in each iteration t , the query point is chosen as $\mathbf{x}_t = \arg \max_{\mathbf{x} \in \mathcal{X}} \gamma_{y_*^t}(\mathbf{x}) \frac{\psi(\gamma_{y_*^t}(\mathbf{x}))}{2\Psi(\gamma_{y_*^t}(\mathbf{x}))} - \log(\Psi(\gamma_{y_*^t}(\mathbf{x})))$, where $\gamma_{y_*^t}(\mathbf{x}) = \frac{y_*^t - \mu_t(\mathbf{x})}{\sigma_t(\mathbf{x})}$ and y_*^t is drawn from F , then with probability at least $1 - \delta$, in $T' = \sum_{i=1}^T \log_w \frac{\delta}{2\pi_i}$ number of iterations, the simple regret satisfies*

$$r_{T'} \leq \sqrt{\frac{C\rho_T}{T}} (\nu_{t^*} + \zeta_T) \quad (7)$$

where $C = 2/\log(1 + \sigma^{-2})$ and $\zeta_T = (2\log(\frac{\pi_T}{\delta}))^{\frac{1}{2}}$; π satisfies $\sum_{i=1}^T \pi_i^{-1} \leq 1$ and $\pi_t > 0$, and $t^* = \arg \max_t \nu_t$ with $\nu_t \triangleq \min_{\mathbf{x} \in \mathcal{X}, y_*^t > f_*} \gamma_{y_*^t}(\mathbf{x})$, and ρ_T is the maximum information gain of at most T selected points.

3.5. Model Adaptation

In practice we do not know the hyper-parameters of the GP, so we must adapt our GP model as we observe more data. A standard way to learn the GP hyper-parameters is to optimize the marginal data likelihood with respect to the hyper-parameters. As a full Bayesian treatment, we can also draw samples of the hyper-parameters using slice sampling (Vanhatalo et al., 2013), and then marginalize out the hyper-parameters in our acquisition function in Eq. (6). Namely, if we use E to denote the set of sampled settings for the GP hyper-parameters, our acquisition function becomes

$$\alpha_t(\mathbf{x}) = \sum_{\eta \in E} \sum_{y_* \in Y_*} \left[\frac{\gamma_{y_*}^\eta(\mathbf{x}) \psi(\gamma_{y_*}^\eta(\mathbf{x}))}{2\Psi(\gamma_{y_*}^\eta(\mathbf{x}))} - \log(\Psi(\gamma_{y_*}^\eta(\mathbf{x}))) \right],$$

where $\gamma_{y_*}^\eta(\mathbf{x}) = \frac{y_* - \mu_t^\eta(\mathbf{x})}{\sigma_t^\eta(\mathbf{x})}$ and the posterior inference on the mean function μ_t^η and σ_t^η depends on the GP hyper-parameter setting η . Similar approaches have been used in (Hernández-Lobato et al., 2014; Snoek et al., 2012).

4. High Dimensional MES with Add-GP

The high-dimensional input setting has been a challenge for many BO methods. We extend MES to this setting via additive Gaussian processes (Add-GP). In the past, Add-GP has been used and analyzed for GP-UCB (Kandasamy et al., 2015), which assumed the high dimensional black-box function is a summation of several disjoint lower dimensional functions. Utilizing this special additive structure, we overcome the statistical problem of having insufficient data to recover a complex function, and the difficulty of optimizing acquisition functions in high dimensions.

Since the function components $f^{(m)}$ are independent, we can maximize the mutual information between the input in the active dimensions A_m and maximum of $f^{(m)}$ for each component separately. Hence, we have a separate acquisition function for each component, where $y^{(m)}$ is the evaluation of $f^{(m)}$:

$$\alpha_t^{(m)}(\mathbf{x}) = I(\{\mathbf{x}^{A_m}, y^{(m)}\}; y_*^{(m)} | D_t) \quad (8)$$

$$= H(p(y^{(m)} | D_t, \mathbf{x}^{A_m})) - \mathbb{E}[H(p(y^{(m)} | D_t, \mathbf{x}^{A_m}, y_*^{(m)}))] \quad (9)$$

$$\approx \sum_{y_*^{(m)}} \gamma_{y_*^{(m)}}^{(m)}(\mathbf{x}) \frac{\psi(\gamma_{y_*^{(m)}}^{(m)}(\mathbf{x}))}{2\Psi(\gamma_{y_*^{(m)}}^{(m)}(\mathbf{x}))} - \log(\Psi(\gamma_{y_*^{(m)}}^{(m)}(\mathbf{x}))) \quad (10)$$

where $\gamma_{y_*^{(m)}}^{(m)}(\mathbf{x}) = \frac{y_*^{(m)} - \mu_t^{(m)}(\mathbf{x})}{\sigma_t^{(m)}(\mathbf{x})}$. Analogously to the non-additive case, we sample $y_*^{(m)}$, separately for each function component. We select the final x_t by choosing a sub-vector $x_t^{(m)} \in \arg \max_{\mathbf{x}^{(m)} \in A_m} \alpha_t^{(m)}(\mathbf{x}^{(m)})$ and concatenating the components.

Sampling $y_*^{(m)}$ with a Gumbel distribution. The Gumbel sampling from Section 3.1 directly extends to sampling $y_*^{(m)}$, approximately. We simply need to sample from the component-wise CDF $\widehat{\Pr}[y_*^{(m)} < z] = \prod_{\mathbf{x} \in \hat{\mathcal{X}}} \Psi(\gamma_y^{(m)}(\mathbf{x}))$, and use the same Gumbel approximation.

Sampling $y_*^{(m)}$ via posterior functions. The additive structure removes some connections on the input-to-hidden layer of our 1-hidden-layer neural network approximation $\hat{f}(\mathbf{x}) = \mathbf{a}_t^\top \phi(\mathbf{x})$. Namely, for each feature function ϕ there exists a unique group m such that ϕ is only active on \mathbf{x}^{A_m} , and $\phi(\mathbf{x}) = \sqrt{\frac{2}{D}} \cos(\omega^\top \mathbf{x}^{A_m} + c)$ where $\mathbb{R}^{|A_m|} \ni \omega \sim$

$\hat{\kappa}^{(m)}(\omega)$ and $c \sim U[0, 2\pi]$. Similar to the non-additive case, we may draw a posterior sample $\mathbf{a}_t \sim \mathcal{N}(\boldsymbol{\nu}_t, \boldsymbol{\Sigma}_t)$ where $\boldsymbol{\nu}_t = \sigma^{-2} \boldsymbol{\Sigma}_t Z_t^\top \mathbf{y}_t$ and $\boldsymbol{\Sigma}_t = (Z Z^\top \sigma^{-2} + \mathbf{I})^{-1}$. Let $B_m = \{i : \phi_i(\mathbf{x}) \text{ is active on } \mathbf{x}^{A_m}\}$. The posterior sample for the function component $f^{(m)}$ is $\hat{f}^{(m)}(\mathbf{x}) = (\mathbf{a}_t^{B_m})^\top \phi^{B_m}(\mathbf{x}^{A_m})$. Then we can maximize $\hat{f}^{(m)}$ to obtain a sample for $y_*^{(m)}$.

The algorithm for the additive max-value entropy search method (add-MES) is shown in Algorithm 2. The function APPROX-MI does the pre-computation for approximating the mutual information in a similar way as in Algorithm 1, except that it only acts on the active dimensions in the m -th group.

Algorithm 2 Additive Max-value Entropy Search

```

1: function Add-MES ( $f, D_0$ )
2:   for  $t = 1, \dots, T$  do
3:     for  $m = 1, \dots, M$  do
4:        $\alpha_{t-1}^{(m)}(\cdot) \leftarrow \text{APPROX-MI}(D_{t-1})$ 
5:        $\mathbf{x}_t^{A_m} \leftarrow \arg \max_{\mathbf{x}^{A_m} \in \mathcal{X}^{A_m}} \alpha_{t-1}^{(m)}(\mathbf{x})$ 
6:     end for
7:      $\mathbf{y}_t \leftarrow f(\mathbf{x}_t) + \epsilon_t, \epsilon_t \sim \mathcal{N}(0, \sigma^2)$ 
8:      $\mathcal{D}_t \leftarrow D_{t-1} \cup \{\mathbf{x}_t, \mathbf{y}_t\}$ 
9:   end for
10: end function
    
```

5. Experiments

In this section, we probe the empirical performance of MES and add-MES on a variety of tasks. Here, MES-G denotes MES with y_* sampled from the approximate Gumbel distribution, and MES-R denotes MES with y_* computed by maximizing a sampled function represented by random features. Following (Hennig & Schuler, 2012; Hernández-Lobato et al., 2014), we adopt the zero mean function and non-isotropic squared exponential kernel as the prior for the GP. We compare to methods from the entropy search family, i.e., ES and PES, and to other popular Bayesian optimization methods including GP-UCB (denoted by UCB), PI, EI and EST. The parameter for GP-UCB was set according to Theorem 2 in (Srinivas et al., 2010); the parameter for PI was set to be the observation noise σ . For the functions with unknown GP hyper-parameters, every 10 iterations, we learn the GP hyper-parameters using the same approach as was used by PES (Hernández-Lobato et al., 2014). For the high dimensional tasks, we follow (Kandasamy et al., 2015) and sample the additive structure/GP parameters with the highest data likelihood when they are unknown. We evaluate performance according to the simple regret and inference regret as defined in Section 2.3. We used the open source Matlab implementation of PES, ES and EST (Hennig & Schuler, 2012; Hernández-Lobato




Figure 1. (a) Inference regret; (b) simple regret. MES methods are much less sensitive to the number of maxima y_* sampled for the acquisition function (1, 10 or 100) than PES is to the number of argmaxes x_* .

Table 1. The runtime of selecting the next input. PES 100 is significantly slower than other methods. MES-G’s runtime is comparable to the fastest method EI while it performs better in terms of simple and inference regrets.

METHOD	TIME (S)	METHOD	TIME (S)
UCB	0.08 ± 0.05	PES 1	0.20 ± 0.06
PI	0.10 ± 0.02	MES-R 100	5.85 ± 0.86
EI	0.07 ± 0.03	MES-R 10	0.67 ± 0.11
EST	0.15 ± 0.02	MES-R 1	0.13 ± 0.03
ES	8.07 ± 3.02	MES-G 100	0.12 ± 0.02
PES 100	15.24 ± 4.44	MES-G 10	0.09 ± 0.02
PES 10	1.61 ± 0.50	MES-G 1	0.09 ± 0.03

et al., 2014; Wang et al., 2016). Our Matlab code and test functions are available at <https://github.com/zi-w/Max-value-Entropy-Search/>.

5.1. Synthetic Functions

We begin with a comparison on synthetic functions sampled from a 3-dimensional GP, to probe our conjecture that MES is much more robust to the number of y_* sampled to estimate the acquisition function than PES is to the number of x_* samples. For PES, we sample 100 (PES 100), 10 (PES 10) and 1 (PES 1) argmaxes for the acquisition function. Similarly, we sample 100, 10, 1 y_* values for MES-R and MES-G. We average the results on 100 functions sampled from the same Gaussian kernel with scale parameter 5.0 and bandwidth parameter 0.0625, and observation noise $\mathcal{N}(0, 0.01^2)$.

Figure 1 shows the simple and inference regrets. For both regret measures, PES is very sensitive to the the number of x_* sampled for the acquisition function: 100 samples lead to much better results than 10 or 1. In contrast, both MES-G and MES-R perform competitively even with 1 or 10 samples. Overall, MES-G is slightly better than MES-R, and both MES methods performed better than other ES methods. MES methods performed better than all other methods with respect to simple regret. For inference regret, MES methods performed similarly to EST, and much better than all other methods including PES and ES.

In Table 1, we show the runtime of selecting the next input per iteration¹ using GP-UCB, PI, EI, EST, ES, PES, MES-R and MES-G on the synthetic data with fixed GP hyper-parameters. For PES and MES-R, every x_* or y_* requires running an optimization sub-procedure, so their running time grows noticeably with the number of samples. MES-G avoids this optimization, and competes with the fastest methods EI and UCB.

In the following experiments, we set the number of x_* sampled for PES to be 200, and the number of y_* sampled for MES-R and MES-G to be 100 unless otherwise mentioned.

5.2. Optimization Test Functions

We test on three challenging optimization test functions: the 2-dimensional eggholder function, the 10-dimensional Shekel function and the 10-dimensional Michalewicz function. All of these functions have many local optima. We randomly sample 1000 points to learn a good GP hyper-parameter setting, and then run the BO methods with the same hyper-parameters. The first observation is the same for all methods. We repeat the experiments 10 times. The averaged simple regret is shown in the appendix, and the inference regret is shown in Table 2. On the 2-d eggholder function, PES was able to achieve better function values faster than all other methods, which verified the good performance of PES when sufficiently many x_* are sampled. However, for higher-dimensional test functions, the 10-d Shekel and 10-d Michalewicz function, MES methods performed much better than PES and ES, and MES-G performed better than all other methods.

5.3. Tuning Hyper-parameters for Neural Networks

Next, we experiment with Levenberg-Marquardt optimization for training a 1-hidden-layer neural network. The 4 parameters we tune with BO are the number of neurons, the damping factor μ , the μ -decrease factor, and the μ -increase factor. We test regression on the Boston housing dataset

¹All the timing experiments were run exclusively on an Intel(R) Xeon(R) CPU E5-2680 v4 @ 2.40GHz. The function evaluation time is excluded.

Table 2. Inference regret R_T for optimizing the eggholder function, Shekel function, and Michalewicz function.

METHOD	EGGHOLDER	SHEKEL	MICHALEWICZ
UCB	141.00 \pm 70.96	9.40 \pm 0.26	6.07 \pm 0.53
PI	52.04 \pm 39.03	6.64 \pm 2.00	4.97 \pm 0.39
EI	71.18 \pm 59.18	6.63 \pm 0.87	4.80 \pm 0.60
EST	55.84 \pm 24.85	5.57 \pm 2.56	5.33 \pm 0.46
ES	48.85 \pm 29.11	6.43 \pm 2.73	5.11 \pm 0.73
PES	37.94 \pm 26.05	8.73 \pm 0.67	5.17 \pm 0.74
MES-R	54.47 \pm 37.71	6.17 \pm 1.80	4.97 \pm 0.59
MES-G	46.56 \pm 27.05	5.45 \pm 2.07	4.49 \pm 0.51




Figure 2. Tuning hyper-parameters for training a neural network, (a) Boston housing dataset; (b) breast cancer dataset. MES methods perform better than other methods on (a), while for (b), MES-G, UCB, PES perform similarly and better than others.

and classification on the breast cancer dataset (Bache & Lichman, 2013). The experiments are repeated 20 times, and the neural network’s weight initialization and all other parameters are set to be the same to ensure a fair comparison. Both of the datasets were randomly split into train/validation/test sets. We initialize the observation set to have 10 random function evaluations which were set to be the same across all the methods. The averaged simple regret for the regression L2-loss on the validation set of the Boston housing dataset is shown in Fig. 2(a), and the classification accuracy on the validation set of the breast cancer dataset is shown in Fig. 2(b). For the classification problem on the breast cancer dataset, MES-G, PES and UCB achieved a similar simple regret. On the Boston housing dataset, MES methods achieved a lower simple regret. We also show the inference regrets for both datasets in Table 3.

5.4. Active Learning for Robot Pushing

We use BO to do active learning for the pre-image learning problem for pushing (Kaelbling & Lozano-Pérez, 2017). The function we optimize takes as input the pushing action of the robot, and outputs the distance of the pushed object to the goal location. We use BO to minimize the function in

 Table 3. Inference regret R_T for tuning neural network hyper-parameters on the Boston housing and breast cancer datasets.

METHOD	BOSTON	CANCER (%)
UCB	1.64 \pm 0.43	3.83 \pm 0.01
PI	2.15 \pm 0.99	4.40 \pm 0.01
EI	1.99 \pm 1.03	4.40 \pm 0.01
EST	1.65 \pm 0.57	3.93 \pm 0.01
ES	1.79 \pm 0.61	4.14 \pm 0.00
PES	1.52 \pm 0.32	3.84 \pm 0.01
MES-R	1.54 \pm 0.56	3.96 \pm 0.01
MES-G	1.51 \pm 0.61	3.83 \pm 0.01




Figure 4. BO for active data selection on two robot pushing tasks for minimizing the distance to a random goal with (a) 3-D actions and (b) 4-D actions. MES methods perform better than other methods on the 3-D function. For the 4-D function, MES methods converge faster to a good regret, while PI achieves lower regret in the very end.

order to find a good pre-image for pushing the object to the designated goal location. The first function we tested has a 3-dimensional input: robot location (r_x, r_y) and pushing duration t_r . We initialize the observation size to be one, the same across all methods. The second function has a 4-dimensional input: robot location and angle (r_x, r_y, r_θ), and pushing duration t_r . We initialize the observation to be 50 random points and set them the same for all the methods. We select 20 random goal locations for each function to test if BO can learn where to push for these locations. We show the simple regret in Fig. 4 and the inference regret in Table 4. MES methods performed on a par with or better than their competitors.

5.5. High Dimensional BO with Add-MES

In this section, we test our add-MES algorithm on high dimensional black-box function optimization problems. First we compare add-MES and add-GP-UCB (Kandasamy et al., 2015) on a set of synthetic additive functions with known additive structure and GP hyper-parameters. Each function component of the synthetic additive function is active on at most three input dimensions, and is sampled

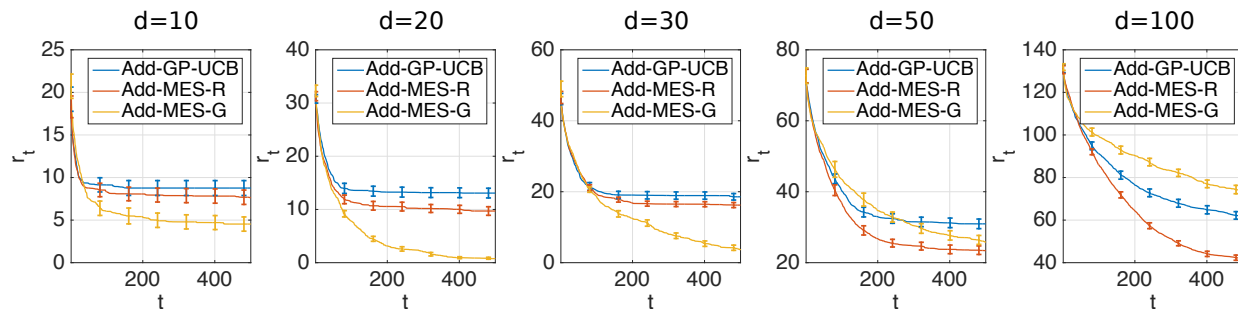


Figure 3. Simple regrets for add-GP-UCB and add-MES methods on the synthetic add-GP functions. Both add-MES methods outperform add-GP-UCB except for add-MES-G on the input dimension $d = 100$. Add-MES-G achieves the lowest simple regret when d is relatively low, while for higher d add-MES-R becomes better than add-MES-G.

Table 4. Inference regret R_T for action selection in robot pushing.

METHOD	3-D ACTION	4-D ACTION
UCB	1.10 ± 0.66	0.56 ± 0.44
PI	2.03 ± 1.77	0.16 ± 0.20
EI	1.89 ± 1.87	0.30 ± 0.33
EST	0.70 ± 0.90	0.24 ± 0.17
ES	0.62 ± 0.59	0.25 ± 0.20
PES	0.81 ± 1.27	0.38 ± 0.38
MES-R	0.61 ± 1.23	0.16 ± 0.10
MES-G	0.61 ± 1.26	0.24 ± 0.25

from a GP with zero mean and Gaussian kernel (bandwidth = 0.1 and scale = 5). For the parameter of add-GP-UCB, we follow (Kandasamy et al., 2015) and set $\beta_t^{(m)} = |A_m| \log 2t/5$. We set the number of $y_*^{(m)}$ sampled for each function component in add-MES-R and add-MES-G to be 1. We repeat each experiment for 50 times for each dimension setting. The results for simple regret are shown in Fig. 3. Add-MES methods perform much better than add-GP-UCB in terms of simple regret. Interestingly, add-MES-G works better in lower dimensional cases where $d = 10, 20, 30$, while add-MES-R outperforms both add-MES-G and add-GP-UCB for higher dimensions where $d = 50, 100$. In general, MES-G tends to overestimate the maximum of the function because of the independence assumption, and MES-R tends to underestimate the maximum of the function because of the imperfect global optimization of the posterior function samples. We conjecture that MES-R is better for settings where exploitation is preferred over exploration (e.g., not too many local optima), and MES-G works better if exploration is preferred.

To further verify the performance of add-MES in high dimensional problems, we test on two real-world high dimensional experiments. One is a function that returns the distance between a goal location and two objects being pushed by a robot which has 14 parameters². The other function re-

²We implemented the function in (Catto, 2011).

turns the walking speed of a planar bipedal robot, with 25 parameters to tune (Westervelt et al., 2007). In Fig. 5, we show the simple regrets achieved by add-GP-UCB and add-MES. Add-MES methods performed competitively compared to add-GP-UCB on both tasks.

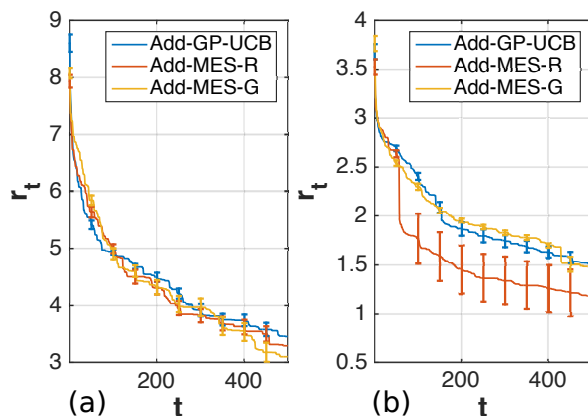


Figure 5. Simple regrets for add-GP-UCB and add-MES methods on (a) a robot pushing task with 14 parameters and (b) a planar bipedal walker optimization task with 25 parameters. Both MES methods perform competitively comparing to add-GP-UCB.

6. Conclusion

We proposed a new information-theoretic approach, max-value entropy search (MES), for optimizing expensive black-box functions. MES is competitive with or better than previous entropy search methods, but at a much lower computational cost. Via additive GPs, MES is adaptable to high-dimensional settings. We theoretically connected MES to other popular Bayesian optimization methods including entropy search, GP-UCB, PI, and EST, and showed a bound on the simple regret for a variant of MES. Empirically, MES performs well on a variety of tasks.

Acknowledgements

We thank Prof. Leslie Pack Kaelbling and Prof. Tomás Lozano-Pérez for discussions on active learning and Dr. William Huber for his solution to “Extreme Value Theory - Show: Normal to Gumbel” at `stats.stackexchange.com`, which leads to our Gumbel approximation in Section 3.1. We gratefully acknowledge support from NSF CAREER award 1553284, NSF grants 1420927 and 1523767, from ONR grant N00014-14-1-0486, and from ARO grant W911NF1410433. We thank MIT Supercloud and the Lincoln Laboratory Supercomputing Center for providing computational resources. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of our sponsors.

References

- Auer, Peter. Using confidence bounds for exploitation-exploration tradeoffs. *Journal of Machine Learning Research*, 3:397–422, 2002.
- Bache, Kevin and Lichman, Moshe. UCI machine learning repository. 2013.
- Brochu, Eric, Cora, Vlad M, and De Freitas, Nando. A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. Technical Report TR-2009-023, University of British Columbia, 2009.
- Calandra, Roberto, Seyfarth, André, Peters, Jan, and Deisenroth, Marc Peter. An experimental comparison of Bayesian optimization for bipedal locomotion. In *International Conference on Robotics and Automation (ICRA)*, 2014.
- Catto, Erin. Box2D, a 2D physics engine for games. <http://box2d.org>, 2011.
- Djolonga, Josip, Krause, Andreas, and Cevher, Volkan. High-dimensional Gaussian process bandits. In *Advances in Neural Information Processing Systems (NIPS)*, 2013.
- Duvenaud, David K, Nickisch, Hannes, and Rasmussen, Carl E. Additive Gaussian processes. In *Advances in Neural Information Processing Systems (NIPS)*, 2011.
- Fisher, Ronald Aylmer. *The genetical theory of natural selection: a complete variorum edition*. Oxford University Press, 1930.
- Hennig, Philipp and Schuler, Christian J. Entropy search for information-efficient global optimization. *Journal of Machine Learning Research*, 13:1809–1837, 2012.
- Hernández-Lobato, José Miguel, Hoffman, Matthew W, and Ghahramani, Zoubin. Predictive entropy search for efficient global optimization of black-box functions. In *Advances in Neural Information Processing Systems (NIPS)*, 2014.
- Hoffman, Matthew W and Ghahramani, Zoubin. Output-space predictive entropy search for flexible global optimization. In *NIPS workshop on Bayesian Optimization*, 2015.
- Kaelbling, Leslie Pack and Lozano-Pérez, Tomás. Learning composable models of primitive actions. In *International Conference on Robotics and Automation (ICRA)*, 2017.
- Kandasamy, Kirthevasan, Schneider, Jeff, and Poczos, Barnabas. High dimensional Bayesian optimisation and bandits via additive models. In *International Conference on Machine Learning (ICML)*, 2015.
- Kawaguchi, Kenji, Kaelbling, Leslie Pack, and Lozano-Pérez, Tomás. Bayesian optimization with exponential convergence. In *Advances in Neural Information Processing Systems (NIPS)*, 2015.
- Kawaguchi, Kenji, Maruyama, Yu, and Zheng, Xiaoyu. Global continuous optimization with error bound and fast convergence. *Journal of Artificial Intelligence Research*, 56(1):153–195, 2016.
- Krause, Andreas and Ong, Cheng S. Contextual Gaussian process bandit optimization. In *Advances in Neural Information Processing Systems (NIPS)*, 2011.
- Kushner, Harold J. A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise. *Journal of Fluids Engineering*, 86(1):97–106, 1964.
- Li, Chun-Liang, Kandasamy, Kirthevasan, Póczos, Barnabás, and Schneider, Jeff. High dimensional Bayesian optimization via restricted projection pursuit models. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2016.
- Lizotte, Daniel J, Wang, Tao, Bowling, Michael H, and Schuurmans, Dale. Automatic gait optimization with Gaussian process regression. In *International Conference on Artificial Intelligence (IJCAI)*, 2007.
- Massart, Pascal. *Concentration Inequalities and Model Selection*, volume 6. Springer, 2007.
- Močkus, J. On Bayesian methods for seeking the extremum. In *Optimization Techniques IFIP Technical Conference*, 1974.
- Neal, R.M. *Bayesian Learning for Neural networks*. Lecture Notes in Statistics 118. Springer, 1996.
- Rahimi, Ali, Recht, Benjamin, et al. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems (NIPS)*, 2007.
- Rasmussen, Carl Edward and Williams, Christopher KI. Gaussian processes for machine learning. *The MIT Press*, 2006.
- Rudin, Walter. *Fourier analysis on groups*. John Wiley & Sons, 2011.
- Slepian, David. The one-sided barrier problem for Gaussian noise. *Bell System Technical Journal*, 41(2):463–501, 1962.
- Snoek, Jasper, Larochelle, Hugo, and Adams, Ryan P. Practical Bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems (NIPS)*, 2012.
- Srinivas, Niranjan, Krause, Andreas, Kakade, Sham M, and Seeger, Matthias. Gaussian process optimization in the bandit setting: no regret and experimental design. In *International Conference on Machine Learning (ICML)*, 2010.

Thornton, Chris, Hutter, Frank, Hoos, Holger H, and Leyton-Brown, Kevin. Auto-WEKA: combined selection and hyperparameter optimization of classification algorithms. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2013.

Vanhatalo, Jarno, Riihimäki, Jaakko, Hartikainen, Jouni, Jylänki, Pasi, Tolvanen, Ville, and Vehtari, Aki. Gpstuff: Bayesian modeling with gaussian processes. *Journal of Machine Learning Research*, 14(Apr):1175–1179, 2013.

Von Mises, Richard. La distribution de la plus grande de n valeurs. *Rev. math. Union interbalcanique*, 1936.

Wang, Zi, Zhou, Bolei, and Jegelka, Stefanie. Optimization as estimation with Gaussian processes in bandit settings. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2016.

Wang, Zi, Jegelka, Stefanie, Kaelbling, Leslie Pack, and Lozano-Pérez, Tomás. Focused model-learning and planning for non-Gaussian continuous state-action systems. In *International Conference on Robotics and Automation (ICRA)*, 2017.

Wang, Ziyu, Zoghi, Masrour, Hutter, Frank, Matheson, David, and De Freitas, Nando. Bayesian optimization in high dimensions via random embeddings. In *International Conference on Artificial Intelligence (IJCAI)*, 2013.

Westervelt, Eric R, Grizzle, Jessy W, Chevallereau, Christine, Choi, Jun Ho, and Morris, Benjamin. *Feedback control of dynamic bipedal robot locomotion*, volume 28. CRC press, 2007.

A. Related work

Our work is largely inspired by the entropy search (ES) methods (Hennig & Schuler, 2012; Hernández-Lobato et al., 2014), which established the information-theoretic view of Bayesian optimization by evaluating the inputs that are most informative to the arg max of the function we are optimizing.

Our work is also closely related to probability of improvement (PI) (Kushner, 1964), expected improvement (EI) (Moćkus, 1974), and the BO algorithms using upper confidence bound to direct the search (Auer, 2002; Kawaguchi et al., 2015; 2016), such as GP-UCB (Srinivas et al., 2010). In (Wang et al., 2016), it was pointed out that GP-UCB and PI are closely related by exchanging the parameters. Indeed, all these algorithms build in the heuristic that the next evaluation point needs to be likely to achieve the maximum function value or have high probability of improving the current evaluations, which in turn, may also give more information on the function optima like how ES methods queries. These connections become clear as stated in Section 3.1 of our paper.

Finding these points that may have good values in high dimensional space is, however, very challenging. In the

past, high dimensional BO algorithms were developed under various assumptions such as the existence of a lower dimensional function structure (Djulonga et al., 2013; Wang et al., 2013), or an additive function structure where each component is only active on a lower manifold of the space (Li et al., 2016; Kandasamy et al., 2015). In this work, we show that our method also works well in high dimensions with the additive assumption made in (Kandasamy et al., 2015).

B. Using the Gumbel distribution to sample y_*

To sample the function maximum y_* , our first approach is to approximate the distribution for y^* and then sample from that distribution. We use independent Gaussians to approximate the correlated $f(\mathbf{x})$, $\forall \mathbf{x} \in \hat{\mathcal{X}}$ where $\hat{\mathcal{X}}$ is a discretization of the input search space \mathcal{X} (unless \mathcal{X} is discrete, in which case $\hat{\mathcal{X}} = \mathcal{X}$). A similar approach was adopted in (Wang et al., 2016). We can show that by assuming $\{f(\mathbf{x})\}_{\mathbf{x} \in \hat{\mathcal{X}}}$, our approximated distribution gives a distribution for an upperbound on $f(\mathbf{x})$.

Lemma B.1 (Slepian’s Comparison Lemma (Slepian, 1962; Massart, 2007)). *Let $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$ be two multivariate Gaussian random vectors with the same mean and variance, such that*

$$\mathbb{E}[\mathbf{v}_i \mathbf{v}_j] \leq \mathbb{E}[\mathbf{u}_i \mathbf{u}_j], \forall i, j.$$

Then for every y

$$\Pr[\sup_{i \in [1, n]} \mathbf{v}_i \leq y] \leq \Pr[\sup_{i \in [1, n]} \mathbf{u}_i \leq y].$$

By the Slepian’s lemma, if the covariance $k_i(\mathbf{x}, \mathbf{x}') \geq 0$, $\forall \mathbf{x}, \mathbf{x}' \in \hat{\mathcal{X}}$, using the independent assumption with give us a distribution on the upperbound \hat{y}_* of $f(\mathbf{x})$, $\Pr[\hat{y}_* < y] = \prod_{\mathbf{x} \in \hat{\mathcal{X}}} \Psi(\gamma_y(\mathbf{x}))$.

We then use the Gumbel distribution to approximate the distribution for the maximum of the function values for $\hat{\mathcal{X}}$, $\Pr[\hat{y}_* < y] = \prod_{\mathbf{x} \in \hat{\mathcal{X}}} \Psi(\gamma_y(\mathbf{x}))$. If for all $\mathbf{x} \in \hat{\mathcal{X}}$, $f(\mathbf{x})$ have the same mean and variance, the Gumbel approximation is in fact asymptotically correct by the Fisher-Tippett-Gnedenko theorem (Fisher, 1930).

Theorem B.2 (The Fisher-Tippett-Gnedenko Theorem (Fisher, 1930)). *Let $\{v_i\}_{i=1}^{\infty}$ be a sequence of independent and identically-distributed random variables, and $M_n = \max_{1 \leq i \leq n} v_i$. If there exist constants $a_n > 0, b_n \in \mathbb{R}$ and a non degenerate distribution function F such that $\lim_{n \rightarrow \infty} \Pr(\frac{M_n - b_n}{a_n} \leq x) = F(x)$, then the limit distribution F belongs to either the Gumbel, the Fréchet or the Weibull family.*

In particular, for i.i.d. Gaussians, the limit distribution of the maximum of them belongs to the Gumbel distri-

bution (Von Mises, 1936). Though the Fisher-Tippett-Gnedenko theorem does not hold for independent and differently distributed Gaussians, in practice we still find it useful in approximating $\Pr[\hat{y}_* < y]$. In Figure 6, we show an example of the result of the approximation for the distribution of the maximum of $f(\mathbf{x}) \sim GP(\mu_t, k_t) \forall \mathbf{x} \in \hat{\mathcal{X}}$ given 50 observed data points randomly selected from a function sample from a GP with 0 mean and Gaussian kernel.




Figure 6. An example of approximating the cumulative probability of the maximum of independent differently distributed Gaussians $\Pr[\hat{y}_* < y]$ (Exact) with a Gumbel distribution $\mathcal{G}(a, b)$ (Approx) via percentile matching.

C. Regret bounds

Based on the connection of MES to EST, we show the bound on the learning regret for MES with a point estimate for $\alpha(x)$.

Theorem 3.2 (Simple Regret Bound). *Let F be the cumulative probability distribution for the maximum of any function f sampled from $GP(\mu, k)$ over the compact search space $\mathcal{X} \subset \mathbb{R}^d$, where $k(\mathbf{x}, \mathbf{x}') \leq 1, \forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}$. Let $f_* = \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$ and $w = F(f_*) \in (0, 1)$, and assume the observation noise is iid $\mathcal{N}(0, \sigma)$. If in each iteration t , the query point is chosen as $\mathbf{x}_t = \arg \max_{\mathbf{x} \in \mathcal{X}} \gamma_{y_*^t}(\mathbf{x}) \frac{\psi(\gamma_{y_*^t}(\mathbf{x}))}{2\Psi(\gamma_{y_*^t}(\mathbf{x}))} - \log(\Psi(\gamma_{y_*^t}(\mathbf{x})))$, where $\gamma_{y_*^t}(\mathbf{x}) = \frac{y_*^t - \mu_t(\mathbf{x})}{\sigma_t(\mathbf{x})}$ and y_*^t is drawn from F , then with probability at least $1 - \delta$, in $T' = \sum_{i=1}^T \log_w \frac{\delta}{2\pi_i}$ number of iterations, the simple regret satisfies*

$$r_{T'} \leq \sqrt{\frac{C\rho_T}{T}}(\nu_{t^*} + \zeta_T) \quad (7)$$

where $C = 2/\log(1 + \sigma^{-2})$ and $\zeta_T = (2\log(\frac{\pi_T}{\delta}))^{\frac{1}{2}}$; π satisfies $\sum_{i=1}^T \pi_i^{-1} \leq 1$ and $\pi_t > 0$, and $t^* = \arg \max_t \nu_t$ with $\nu_t \triangleq \min_{\mathbf{x} \in \mathcal{X}, y_*^t > f_*} \gamma_{y_*^t}(\mathbf{x})$, and ρ_T is the maximum information gain of at most T selected points.

Before we continue to the proof, notice that if the function

upper bound \hat{y}_* is sampled using the approach described in Section 3.1 and $k_t(\mathbf{x}, \mathbf{x}') \geq 0, \forall \mathbf{x}, \mathbf{x}' \in \hat{\mathcal{X}}$, we may still get the regret guarantee by setting $y_* = \hat{y}_*$ (or $y_* = \hat{y}_* + \epsilon L$ if \mathcal{X} is continuous) since $\Pr[\max_{\hat{\mathcal{X}}} \leq y] \geq \Pr[\hat{y}_* < y]$. Moreover, Theorem 3.2 assumes y_* is sampled from a universal maximum distribution of functions from $GP(\mu, k)$, but it is not hard to see that if we have a distribution of maximums adapted from $GP(\mu_t, k_t)$, we can still get the same regret bound by setting $T' = \sum_{i=1}^T \log_{w_i} \frac{\delta}{2\pi_i}$, where $w_i = F_i(f_*)$ and F_i corresponds to the maximum distribution at an iteration where $y_* > f_*$. Next we introduce a few lemmas and then prove Theorem 3.2.

Lemma C.1 (Lemma 3.2 in (Wang et al., 2016)). *Pick $\delta \in (0, 1)$ and set $\zeta_t = (2\log(\frac{\pi_t}{2\delta}))^{\frac{1}{2}}$, where $\sum_{t=1}^T \pi_t^{-1} \leq 1, \pi_t > 0$. Then, it holds that $\Pr[\mu_{t-1}(\mathbf{x}_t) - f(\mathbf{x}_t) \leq \zeta_t \sigma_{t-1}(\mathbf{x}_t), \forall t \in [1, T]] \geq 1 - \delta$.*

Lemma C.2 (Lemma 3.3 in (Wang et al., 2016)). *If $\mu_{t-1}(\mathbf{x}_t) - f(\mathbf{x}_t) \leq \zeta_t \sigma_{t-1}(\mathbf{x}_t)$, the regret at time step t is upper bounded as $\tilde{r}_t \leq (\nu_t + \zeta_t) \sigma_{t-1}(\mathbf{x}_t)$, where $\nu_t \triangleq \min_{\mathbf{x} \in \mathcal{X}} \frac{\hat{\mu}_t - \mu_{t-1}(\mathbf{x})}{\sigma_{t-1}(\mathbf{x})}$, and $\hat{\mu}_t \geq \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$, $\forall t \in [1, T]$.*

Lemma C.3 (Lemma 5.3 in (Srinivas et al., 2010)). *The information gain for the points selected can be expressed in terms of the predictive variances. If $f_T = (f(\mathbf{x}_t)) \in \mathbb{R}^T$:*

$$I(\mathbf{y}_T; \mathbf{f}_T) = \frac{1}{2} \sum_{t=1}^T \log(1 + \sigma^{-2} \sigma_{t-1}^2(\mathbf{x}_t)).$$

Proof. (Theorem 3.2) By lemma 3.1 in our paper, we know that the theoretical results from EST (Wang et al., 2016) can be adapted to MES if $y_* \geq f_*$. The key question is when a sampled y_* that can satisfy this condition. Because the cumulative density $w = F(f_*) \in (0, 1)$ and y_*^t are independent samples from F , there exists at least one y_*^t that satisfies $y_*^t > f_*$ with probability at least $1 - w^{k_i}$ in k_i iterations.

Let $T' = \sum_{i=1}^T k_i$ be the total number of iterations. We split these iterations to T parts where each part have k_i iterations, $i = 1, \dots, T$. By union bound, with probability at least $1 - \sum_{i=1}^T w^{k_i}$, in all the T parts of iterations, we have at least one iteration t_i which samples $y_*^{t_i}$ satisfying $y_*^{t_i} > f_*, \forall i = 1, \dots, T$.

Let $\sum_{i=1}^T w^{k_i} = \frac{\delta}{2}$, we can set $k_i = \log_w \frac{\delta}{2\pi_i}$ for any $\sum_{i=1}^T (\pi_i)^{-1} = 1$. A convenient choice for π_i is $\pi_i = \frac{\pi^2 i^2}{6}$. Hence with probability at least $1 - \frac{\delta}{2}$, there exist a sampled $y_*^{t_i}$ satisfying $y_*^{t_i} > f_*, \forall i = 1, \dots, T$.

Now let $\zeta_{t_i} = (2\log \frac{\pi_{t_i}}{\delta})^{\frac{1}{2}}$. By Lemma C.1 and Lemma C.2, the immediate regret $r_{t_i} = f_* - f(\mathbf{x}_{t_i})$ can be bounded as

$$r_{t_i} \leq (\nu_{t_i} + \zeta_{t_i}) \sigma_{t_i-1}(\mathbf{x}_{t_i}).$$




Figure 7. (a) 2-D eggholder function; (b) 10-D Shekel function; (c) 10-D Michalewicz function. PES achieves lower regret on the 2-d function while MES-G performed better than other methods on the two 10-d optimization test functions.

Note that by assumption $0 \leq \sigma_{t_i-1}^2(\mathbf{x}_{t_i}) \leq 1$, so we have $\sigma_{t_i-1}^2 \leq \frac{\log(1+\sigma^{-2}\sigma_{t_i-1}^2(\mathbf{x}_{t_i}))}{\log(1+\sigma^{-2})}$. Then by Lemma C.3, we have $\sum_{i=1}^T \sigma_{t_i-1}^2(\mathbf{x}_{t_i}) \leq \frac{2}{\log(1+\sigma^{-2})} I(\mathbf{y}_T; \mathbf{f}_T)$ where $\mathbf{f}_T = (f(\mathbf{x}_{t_i}))_{i=1}^T \in \mathbb{R}^T$, $\mathbf{y}_T = (y_{t_i})_{i=1}^T \in \mathbb{R}^T$. From assumptions, we have $I(\mathbf{y}_T; \mathbf{f}_T) \leq \rho_T$. By Cauchy-Schwarz inequality, $\sum_{i=1}^T \sigma_{t_i-1}(\mathbf{x}_{t_i}) \leq \sqrt{T \sum_{i=1}^T \sigma_{t_i-1}^2(\mathbf{x}_{t_i})} \leq \sqrt{\frac{2T\rho_T}{\log(1+\sigma^{-2})}}$. It follows that with probability at least $1 - \delta$,

$$\sum_{i=1}^T r_{t_i} \leq (\nu_{t^*} + \zeta_T) \sqrt{\frac{2T\rho_T}{\log(1+\sigma^{-2})}}.$$

As a result, our learning regret is bounded as

$$r_{T'} \leq \frac{1}{T} \sum_{i=1}^T r_{t_i} \leq (\nu_{t^*} + \zeta_T) \sqrt{\frac{2\rho_T}{T \log(1+\sigma^{-2})}},$$

where $T' = \sum_{i=1}^T k_i = \sum_{i=1}^T \log_w \frac{\delta}{2\pi_i}$ is the total number of iterations. \square

At first sight, it might seem like MES with a point estimate does not have a converging rate as good as *EST* or *GP-UCB*. However, notice that $\min_{\mathbf{x} \in \mathcal{X}} \gamma_{y_1}(\mathbf{x}) < \min_{\mathbf{x} \in \mathcal{X}} \gamma_{y_2}(\mathbf{x})$ if $y_1 < y_2$, which decides the rate of convergence in Eq. 7. So if we use y_* that is too large, the regret bound could be worse. If we use y_* that is smaller than f_* , however, its value won't count towards the learning regret in our proof, so it is also bad for the regret upper bound. With no principled way of setting y_* since f_* is unknown. Our regret bound in Theorem 3.2 is a randomized trade-off between sampling large and small y_* .

For the regret bound in add-GP-MES, it should follow add-GP-UCB. However, because of some technical problems in the proofs of the regret bound for add-GP-UCB, we haven't been able to show a regret bound for add-GP-MES either. Nevertheless, from the experiments on high dimensional functions, the methods worked well in practice.

D. Experiments

In this section, we provide more details on our experiments.

Optimization test functions In Fig. 7, we show the simple regret comparing BO methods on the three challenging optimization test functions: the 2-D eggholder function, the 10-D Shekel function, and the 10-D Michalewicz function.

Choosing the additive decomposition We follow the approach in (Kandasamy et al., 2015), and sample 10000 random decompositions (at most 2 dimensions in each group) and pick the one with the best data likelihood based on 500 data points uniformly randomly sampled from the search space. The decomposition setting was fixed for all the 500 iterations of BO for a fair comparison.