
Trusted-Maximizers Entropy Search for Efficient Bayesian Optimization

Quoc Phong Nguyen¹

Zhaoxuan Wu^{3,4}

Bryan Kian Hsiang Low¹

Patrick Jaillet²

¹Department of Computer Science, National University of Singapore, Republic of Singapore

²Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, USA

³Institute of Data Science, National University of Singapore, Republic of Singapore

⁴NUSGS Integrative Sciences and Engineering Programme, National University of Singapore, Republic of Singapore

Abstract

Information-based *Bayesian optimization* (BO) algorithms have achieved state-of-the-art performance in optimizing a black-box objective function. However, they usually require several approximations or simplifying assumptions (without clearly understanding their effects on the BO performance) and/or their generalization to batch BO is computationally unwieldy, especially with an increasing batch size. To alleviate these issues, this paper presents a novel *trusted-maximizers entropy search* (TES) acquisition function: It measures how much an input query contributes to the information gain on the maximizer over a finite set of *trusted maximizers*, i.e., inputs optimizing functions that are sampled from the Gaussian process posterior belief of the objective function. Evaluating TES requires either only a stochastic approximation with sampling or a deterministic approximation with expectation propagation, both of which are investigated and empirically evaluated using synthetic benchmark objective functions and real-world optimization problems, e.g., hyperparameter tuning of a convolutional neural network and synthesizing ‘physically realizable’ faces to fool a black-box face recognition system. Though TES can naturally be generalized to a batch variant with either approximation, the latter is amenable to be scaled to a much larger batch size in our experiments.

1 INTRODUCTION

Bayesian optimization (BO) is an effective strategy for iteratively obtaining noisy outputs at input queries of an *objective function* to efficiently learn about the input locations of its *maximum* value, i.e., its *maximizers*, regardless of the function’s closed-form expression, derivatives, or convexity

[Brochu et al., 2010b, Shahriari et al., 2015]. It is particularly useful when the budget of function evaluations is limited due to time/monetary costs such as in the design [Brochu et al., 2010a] and machine learning model training [Snoek et al., 2012, Dai et al., 2019]. There has been a rapid growth in the number of works generalizing BO to nonmyopic BO [Ling et al., 2016], high-dimensional BO [Hoang et al., 2018], private outsourced BO [Kharkovskii et al., 2020], multi-fidelity BO [Zhang et al., 2017, 2019], recursive reasoning BO [Dai et al., 2020], and ranking BO [Nguyen et al., 2021b].

On the other hand, restricting a BO solution to only one input query at each iteration cannot take advantage of parallel function evaluations which are becoming increasingly attainable due to, for example, the rapid growth of computing resources’ affordability and availability these days. In such a scenario and the time constraint, the urge to obtain the noisy outputs at a batch of input queries, i.e., *batch Bayesian optimization*, becomes prominent. However, this task is challenging due to the extra modifications needed to account for correlation in a batch such, as seen in several existing works [Ginsbourger et al., 2010, Desautels et al., 2014, Shah and Ghahramani, 2015, González et al., 2016, Daxberger and Low, 2017]. It encourages the development of BO algorithms that can be naturally generalized to the batch scenario without making additional adjustments.

To approach this problem, this paper focuses on information-based BO acquisition functions. In particular, it is the information gain about the maximizer (or the maximum value) of the objective function through observing noisy function outputs at input queries. This definition is much the same in spite of the number of input queries at each BO iteration, e.g., a single query or a batch of queries. The essence of these acquisition functions is different from *upper confidence bound* (UCB) that UCB minimizes the *regret* (over function outputs at evaluated inputs), while information-based acquisitions aim to infer the maximizer’s location regardless of the function outputs at the evaluated inputs [Hennig and Schuler, 2012]. They also have a distinctive

feature from other heuristic acquisition functions such as probability of improvement [Kushner, 1964] and *expected improvement* (EI) [Mockus et al., 1978]. It is about capturing the knowledge/belief of the maximizer (or the maximum value) over the whole input space (a global measure), which cannot be done with local heuristics, as argued in the work of Hennig and Schuler [2012]. Empirically, these information-based acquisition functions such as *predictive entropy search* (PES) have shown notable success in optimizing black-box functions [Hernández-Lobato et al., 2014].

In this paper, we take a different approach in constructing the *trusted-maximizers entropy search* (TES) acquisition function that (1) intuitively conveys the notion of maximizer without directly evaluating the information gain of the global maximizer; and (2) is amenable to batch BO. The first property is made feasible by the observation: while being the global maximizer requires its function output to be not less than those evaluated at all other inputs in the function’s domain, it is much easier to eliminate an input as the global maximizer which requires the existence of only an input whose function output is larger. Hence, if we can construct a set of “promising” inputs to be the maximizer, refining this set by rejecting its inputs can effectively refine the belief of the maximizer. The second property comes from the fact that our approximation does not depend on the input queries, hence considering a batch of input queries does not increase the approximation’s complexity. Both stochastic and deterministic approximations are examined in the paper, and their empirical performances are illustrated using synthetic benchmark objective functions and real-world optimization problems, e.g., hyperparameter tuning of a convolutional neural network and synthesizing ‘physically realizable’ faces to fool a black-box face recognition system.

2 RELATED WORKS

Most of the existing information-based works are centered around an intuitive measure of the information gain of the maximizer of the objective function, e.g., *entropy search* (ES) [Hennig and Schuler, 2012] and PES [Hernández-Lobato et al., 2014]. While ES selects an input query by directly maximizing the information gain on the maximizer via reducing the maximizer’s entropy, it requires a series of approximations. On the other hand, PES exploits the symmetric property of mutual information to sidestep the approximations in ES, and yet, it requires computing a predictive belief that involves an intractable constraint. In our view, these approaches’ inherent difficulty is rooted in the modeling of the objective function’s global maximizer as an input whose function output is not less than those evaluated at other inputs in the domain. Imposing exactly this condition is hardly feasible as it involves the whole input space. Therefore, we would like to develop TES that does not require such a condition.

There are exceptions that focus on the information gain of the maximum value of the objective function such as *output-space predictive entropy search* [Hoffman and Ghahramani, 2015], *max-value entropy search* (MES) [Wang and Jegelka, 2017], *fast information-theoretic Bayesian optimization* [Ru et al., 2018], and *binary entropy search for maximum value prediction* [Nguyen et al., 2021a] which result in a simpler approximation, e.g., MES enjoys a closed-form expression. However, as admitted by MES’s authors, it is not intuitive that the information of the maximum value is a good search strategy for the maximizer.

Regarding the batch BO problem, there have been several works in the existing literature such as q -EI with a constant liar strategy (QEI-CL) [Ginsbourger et al., 2010], Gaussian process batch upper confidence bound (BUCB) [Desautels et al., 2014], QEI [Marmin et al., 2015], batch BO via local penalization (BBO-LP) [González et al., 2016], and parallel predictive entropy search (PPES) [Shah and Ghahramani, 2015]. However, QEI-CL, BUCB, and LP-EI optimize greedily each input one by one in the batch based on a heuristic value at input queries (QEI-CL), or purely based on the predictive variance to encourage exploration (BUCB), or based on a local penalty based on an estimation of the objective function’s Lipschitz constant (LP-EI). On the other hand, QEI and PPES are able to optimize the batch jointly, which is also a property of the proposed TES acquisition function in this paper. It is noted that only PPES is an acquisition function that is based on the information gain. While it is extended from the computationally efficient PES [Hernández-Lobato et al., 2014], the optimization of PPES requires a computationally expensive inner loop of *expectation propagation* (EP). To the best of our knowledge, there have not been any batch BO solution extended from MES. It is probably because a direct extension of MES to a batch variant requires computing the entropy of a truncated multivariate Gaussian random variable.

3 PROBLEM SETUP

We propose a BO algorithm for black-box optimization problems, where the analytical form or the derivative information of the objective function are unknown. Given a black-box function $f : \mathcal{X} \rightarrow \mathbb{R}$, we search for the global maximizer $\mathbf{x}^* \triangleq \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$ in an iterative manner. At each iteration, the BO algorithm is required to select input queries based on the posterior belief of f given all observations $\mathbf{y}_{\mathcal{D}} \triangleq (y_{\mathbf{x}'})_{\mathbf{x}' \in \mathcal{D}}$ at preceding input queries $\mathcal{D} \subset \mathcal{X}$ where $y_{\mathbf{x}'} \triangleq f(\mathbf{x}') + \epsilon$ and $\epsilon \sim \mathcal{N}(0, \sigma_n^2)$. These input queries are selected to improve the belief of \mathbf{x}^* as much as possible. Specifically, BO models the posterior belief of f by a GP with the function’s prior mean $\mathbb{E}[f(\mathbf{x})]$ and covariance $k_{\mathbf{x}\mathbf{x}'} \triangleq \text{cov}[f(\mathbf{x}), f(\mathbf{x}')] for all $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$. The covariance is constructed with the *squared exponential* (SE) kernel $k_{\mathbf{x}\mathbf{x}'} \triangleq \sigma_s^2 \exp(-0.5(\mathbf{x} - \mathbf{x}')^\top \Lambda^{-2}(\mathbf{x} - \mathbf{x}'))$$

where its hyperparameters include the length-scales $\Lambda \triangleq \text{diag}[\ell_1, \dots, \ell_d]$ and the signal variance σ_s^2 . The prior mean is usually set to zero to simplify the equations. Given \mathbf{y}_D , the predictive distributions of $f(\mathbf{x})$ and $y_{\mathbf{x}}$ are Gaussian distributions: $p(f(\mathbf{x})|\mathbf{y}_D) \sim \mathcal{N}(\mu_{\mathbf{x}}, \sigma_{\mathbf{x}}^2)$ and $p(y_{\mathbf{x}}|\mathbf{y}_D) \sim \mathcal{N}(\mu_{\mathbf{x}}, \sigma_{\mathbf{x}}^2 + \sigma_n^2)$, respectively, where

$$\begin{aligned} \mu_{\mathbf{x}} &\triangleq K_{\mathbf{x}D}(K_{DD} + \sigma_n^2 I)^{-1} \mathbf{y}_D \\ \sigma_{\mathbf{x}}^2 &\triangleq k_{\mathbf{x}\mathbf{x}} - K_{\mathbf{x}D}(K_{DD} + \sigma_n^2 I)^{-1} K_{D\mathbf{x}}, \end{aligned} \quad (1)$$

$K_{\mathbf{x}D} \triangleq (k_{\mathbf{x}\mathbf{x}'})_{\mathbf{x}' \in D}$, $K_{DD} \triangleq (k_{\mathbf{x}'\mathbf{x}''})_{\mathbf{x}', \mathbf{x}'' \in D}$, I is the identity matrix, and $K_{D\mathbf{x}} \triangleq K_{\mathbf{x}D}^\top$.

4 TRUSTED-MAXIMIZERS ENTROPY SEARCH

Information-based acquisition functions like ES [Henig and Schuler, 2012] and PES [Hernández-Lobato et al., 2014] have been designed to enable their BO algorithms to improve the (posterior) belief $p(\mathbf{x}^*|\mathbf{y}_D) \triangleq p(f(\mathbf{x}^*) = \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})|\mathbf{y}_D)$ of the global maximizer \mathbf{x}^* of the objective function f . In particular, ES can be directly used to measure the information gain on \mathbf{x}^* by selecting the next input query \mathbf{x} for evaluating f given the noisy observations \mathbf{y}_D from the previous BO iterations:

$$\alpha(\mathbf{y}_D, \mathbf{x}) \triangleq H(\mathbf{x}^*|\mathbf{y}_D) - \mathbb{E}_{p(y_{\mathbf{x}}|\mathbf{y}_D)}[H(\mathbf{x}^*|\mathbf{y}_D \cup \{\mathbf{x}\})]. \quad (2)$$

However, evaluating (2) requires a series of approximations. Furthermore, one of them incurs cubic time in the size of the discretized input domain and thus it becomes expensive with a large input domain (or risks (2) being approximated poorly). To sidestep these issues, PES exploits the symmetric property of mutual information to re-express (2) as

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

Intuitively, the selection of an input query \mathbf{x} to maximize (3) has to trade-off between exploration (hence inducing a large Gaussian predictive entropy $H(y_{\mathbf{x}}|\mathbf{y}_D)$) vs. exploitation of the current posterior belief $p(\mathbf{x}^*|\mathbf{y}_D)$ of \mathbf{x}^* to choose a nearby input \mathbf{x} of f to be evaluated (hence inducing a small expected predictive entropy $\mathbb{E}_{p(\mathbf{x}^*|\mathbf{y}_D)}[H(y_{\mathbf{x}}|\mathbf{y}_D, \mathbf{x}^*)]$) to yield a highly informative observation that in turn improves the posterior belief of \mathbf{x}^* . Computing the predictive belief $p(y_{\mathbf{x}}|\mathbf{y}_D, \mathbf{x}^*)$ necessary for the second entropy term in (3) requires incorporating the constraint of $f(\mathbf{x}^*) \geq f(\mathbf{x}) \forall \mathbf{x} \in \mathcal{X}$ which unfortunately is not tractable without imposing simplifying assumptions. It is not clear to what extent they would compromise the quality of the resulting approximation of the predictive belief $p(y_{\mathbf{x}}|\mathbf{y}_D, \mathbf{x}^*)$. The work of Wang and Jegelka [2017] has proposed using the information gain on the global maximum $f(\mathbf{x}^*)$ but admitted that it is not intuitive why it would necessarily be a good surrogate of that on the global maximizer \mathbf{x}^* . Therefore,

this paper still focuses on finding \mathbf{x}^* , albeit with a new perspective that is different from ES and PES, which we will discuss next.

Let $\mathcal{X}^* \subset \mathcal{X}$ denote a *finite* set of *trusted maximizers*, i.e., inputs that are more likely to be the global maximizer \mathbf{x}^* in the current BO iteration, that is, $p(\mathbf{x}^* \in \mathcal{X}^*|\mathbf{y}_D)$ is large. Let \mathbf{x}^* be a r.v. representing the maximizer of f over \mathcal{X}^* , that is, $p(\mathbf{x}^*|\mathbf{y}_D) \triangleq p(f(\mathbf{x}^*) = \max_{\mathbf{x} \in \mathcal{X}^*} f(\mathbf{x})|\mathbf{y}_D)$. To ease the exposition of our key idea here, we consider a unique maximizer of f over \mathcal{X}^* but do not assume this when we describe our technical approach later. By finding the maximizer \mathbf{x}^* of f over \mathcal{X}^* , it effectively eliminates the remaining input queries in $\mathcal{X}^* \setminus \{\mathbf{x}^*\}$ from being the global maximizer \mathbf{x}^* because $f(\mathbf{x}^*) \geq f(\mathbf{x}^*) > f(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{X}^* \setminus \{\mathbf{x}^*\}$. So, if $p(\mathbf{x}^* \in \mathcal{X}^*|\mathbf{y}_D)$ is large, then eliminating some input queries in \mathcal{X}^* would gain us useful information on the global maximizer \mathbf{x}^* . Furthermore, besides reducing the probability of each input query $\mathbf{x} \in \mathcal{X}^* \setminus \{\mathbf{x}^*\}$ being \mathbf{x}^* , their neighboring input queries also experience a reduced probability of being \mathbf{x}^* . We illustrate this with a simple example of input dimension $d = 1$ in Fig. 1: Let $\mathbf{x}^+ \in [0, 1.4]$ and $\mathbf{x}^- \in [1.6, 3]$. It can be observed that by conditioning on only $f(1) > f(2)$, the probability of $f(\mathbf{x}^+) > f(\mathbf{x}^-)$ in Fig. 1b (hence implying $f(\mathbf{x}^*) > f(\mathbf{x}^-)$) increases from that of 0.5 in Fig. 1a. In other words, the probability of $\mathbf{x}^- \in [1.6, 3]$ being \mathbf{x}^* reduces. The magnitude of reduction depends on the correlation between $f(2)$ and the outputs of f evaluated at the input queries in $[1.6, 3]$, which in turn depends on the length-scale of the SE kernel (see Fig. 1c).

So, if we can construct the set \mathcal{X}^* of trusted maximizers such that $p(\mathbf{x}^* \in \mathcal{X}^*|\mathbf{y}_D)$ is large, then finding the maximizer \mathbf{x}^* of f over \mathcal{X}^* would gain us useful information on the global maximizer \mathbf{x}^* since it reduces the probability of each input query $\mathbf{x} \in \mathcal{X}^* \setminus \{\mathbf{x}^*\}$ and its neighbors being \mathbf{x}^* , as explained above. We therefore form \mathcal{X}^* using samples of \mathbf{x}^* drawn from the posterior belief $p(\mathbf{x}^*|\mathbf{y}_D)$: To do this, we adopt the method of Hernández-Lobato et al. [2014] to first approximately sample $|\mathcal{X}^*|$ functions from the GP predictive belief given \mathbf{y}_D (1) and then compute the global maximizers of these sampled functions to form \mathcal{X}^* .

Given the set \mathcal{X}^* of trusted maximizers computed in the current BO iteration, our *trusted-maximizers entropy search* (TES) acquisition function is designed to enable its BO algorithm to find the maximizer \mathbf{x}^* of f over \mathcal{X}^* by improving the posterior belief $p(\mathbf{x}^*|\mathbf{y}_D)$. Specifically, TES measures the information gain on \mathbf{x}^* by selecting the next input query $\mathbf{x} \in \mathcal{X}$ for evaluating f given the noisy observations (D, \mathbf{y}_D) from the previous BO iterations:

$$\alpha^*(\mathbf{y}_D, \mathbf{x}) \triangleq H(\mathbf{x}^*|\mathbf{y}_D) - \mathbb{E}_{p(y_{\mathbf{x}}|\mathbf{y}_D)}[H(\mathbf{x}^*|\mathbf{y}_D \cup \{\mathbf{x}\})]. \quad (4)$$

Since the second entropy term in (4) requires the posterior belief $p(\mathbf{x}^*|\mathbf{y}_D \cup \{\mathbf{x}\})$ to be computed for prohibitively many different noisy outputs $y_{\mathbf{x}}$, we exploit the symmetric

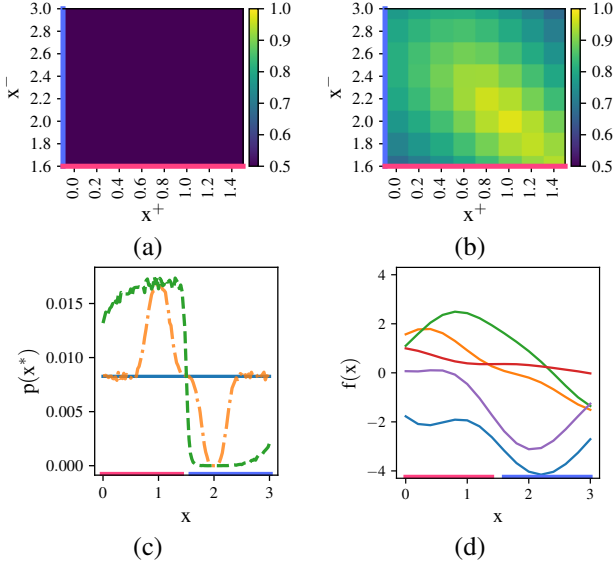


Figure 1: Probability of $f(\mathbf{x}^+) > f(\mathbf{x}^-)$ (a) before and (b) after conditioning on $f(1) > f(2)$ where $\mathbf{x}^+ \in [0, 1.4]$, $\mathbf{x}^- \in [1.6, 3]$, and the belief of f is modeled by a GP using a SE kernel with $\sigma_s^2 = 5$, $\sigma_n^2 = 10^{-9}$, and length-scale $\ell = 1$. (c) Probability of \mathbf{x}^* before (blue) and after conditioning on $f(1) > f(2)$ using the same GP prior with $\ell = 1$ (green) and $\ell = 0.2$ (orange). (d) 5 functions sampled from the same GP prior with $\ell = 1$ and satisfying $f(1) > f(2)$.

property of mutual information to re-express (4) as

$$\alpha^*(\mathbf{y}_D, \mathbf{x}) = H(y_{\mathbf{x}} | \mathbf{y}_D) - \sum_{\mathbf{x}^* \in \mathcal{X}^*} p(\mathbf{x}^* | \mathbf{y}_D) H(y_{\mathbf{x}} | \mathbf{y}_D, \mathbf{x}^*)$$

$$H(y_{\mathbf{x}} | \mathbf{y}_D, \mathbf{x}^*) \triangleq - \int p(y_{\mathbf{x}} | \mathbf{y}_D, \mathbf{x}^*) \log p(y_{\mathbf{x}} | \mathbf{y}_D, \mathbf{x}^*) dy_{\mathbf{x}} \quad (5)$$

where $p(\mathbf{x}^* | \mathbf{y}_D)$ can be evaluated efficiently (Appendix A).

Remark 1. TES (5) differs from PES (3) in measuring the information gain on the maximizer \mathbf{x}^* of f over a *finite* \mathcal{X}^* rather than the entire infinite input domain \mathcal{X} , which we have motivated earlier in this section. A practical implication is that PES needs to assume simplified constraints to achieve tractability while TES does not. For example, PES draws approximate posterior samples of \mathbf{x}^* but does not directly exploit the constraint that the global maximum is not less than the outputs of f evaluated at these samples. Instead, it assumes a simplified soft constraint (with Gaussian noise ϵ) that the output of f evaluated at each sample is not less than the largest noisy output observed in previous BO iterations. By observing that $p(y_{\mathbf{x}} | \mathbf{f}_{\mathcal{X}^*}, \mathbf{y}_D, \mathbf{x}^*) = p(y_{\mathbf{x}} | \mathbf{f}_{\mathcal{X}^*}, \mathbf{y}_D)$, the predictive belief $p(y_{\mathbf{x}} | \mathbf{y}_D, \mathbf{x}^*)$ in the $H(y_{\mathbf{x}} | \mathbf{y}_D, \mathbf{x}^*)$ term (5) can be derived by marginalizing out $\mathbf{f}_{\mathcal{X}^*} \triangleq (f(\mathbf{x}'))_{\mathbf{x}' \in \mathcal{X}^*}^\top$:

$$p(y_{\mathbf{x}} | \mathbf{y}_D, \mathbf{x}^*) = \int p(y_{\mathbf{x}} | \mathbf{f}_{\mathcal{X}^*}, \mathbf{y}_D) p(\mathbf{f}_{\mathcal{X}^*} | \mathbf{y}_D, \mathbf{x}^*) d\mathbf{f}_{\mathcal{X}^*}. \quad (6)$$

Since $p(\mathbf{f}_{\mathcal{X}^*} | \mathbf{y}_D, \mathbf{x}^*)$ cannot be evaluated in closed form, we approximate it with either (a) sampling whose approximation quality improves with the number of samples or (b) a Gaussian distribution to make $p(y_{\mathbf{x}} | \mathbf{y}_D, \mathbf{x}^*)$ Gaussian. These two approximation methods are discussed next.

Remark 2. Since $p(\mathbf{f}_{\mathcal{X}^*} | \mathbf{y}_D, \mathbf{x}^*)$ does not depend on the input query \mathbf{x} , its approximation remains unchanged when our TES acquisition function (5) is optimized with respect to the input query $\mathbf{x} \in \mathcal{X}$ in each BO iteration, that is, $\max_{\mathbf{x} \in \mathcal{X}} \alpha^*(\mathbf{y}_D, \mathbf{x})$. So, $p(\mathbf{f}_{\mathcal{X}^*} | \mathbf{y}_D, \mathbf{x}^*)$ is approximated only once in each BO iteration, which simplifies the optimization procedure and facilitates an efficient generalization to batch BO (Section 4.3). In contrast, the approximations needed to derive PPES and ES depend on \mathbf{x} and hence have to be repeated for every $\mathbf{x} \in \mathcal{X}$ to optimize PPES and ES.

We will discuss the evaluation of TES by both stochastic approximation with sampling, denoted as TES_{sp} , and deterministic approximation with *expectation propagation* (EP), denoted as TES_{ep} , as below.

4.1 STOCHASTIC APPROXIMATION WITH SAMPLING

Let $\mathcal{F}_{\mathbf{x}^*}$ denote a set of samples that are randomly drawn from $p(\mathbf{f}_{\mathcal{X}^*} | \mathbf{y}_D, \mathbf{x}^*)$. This can be obtained via rejection sampling by drawing samples from the GP predictive belief $p(\mathbf{f}_{\mathcal{X}^*} | \mathbf{y}_D)$ (1) and rejecting those samples of $\mathbf{f}_{\mathcal{X}^*}$ where $\max_{\mathbf{x}' \in \mathcal{X}^*} f(\mathbf{x}') > f(\mathbf{x}^*)$. However, if $p(\mathbf{x}^* | \mathbf{y}_D)$ is small, then such a sampling becomes inefficient due to a large number of rejected samples. Importance sampling can be used to avoid this issue. Let $\mathbf{f}_{\setminus \mathbf{x}^*} \triangleq (f(\mathbf{x}'))_{\mathbf{x}' \in \mathcal{X}^* \setminus \{\mathbf{x}^*\}}^\top$. Then,

$$p(\mathbf{f}_{\mathcal{X}^*} | \mathbf{y}_D, \mathbf{x}^*) \propto p(\mathbf{f}_{\mathcal{X}^*}, \mathbf{x}^* | \mathbf{y}_D)$$

$$= p(\mathbf{x}^* | \mathbf{f}_{\setminus \mathbf{x}^*}, \mathbf{y}_D) p(\mathbf{f}_{\setminus \mathbf{x}^*} | \mathbf{y}_D) p(f(\mathbf{x}^*) | \mathbf{f}_{\setminus \mathbf{x}^*}, \mathbf{y}_D, \mathbf{x}^*).$$

We first draw samples of $\mathbf{f}_{\mathcal{X}^*}$ from $p(\mathbf{f}_{\setminus \mathbf{x}^*} | \mathbf{y}_D)$ and then weight these samples by $p(\mathbf{x}^* | \mathbf{f}_{\setminus \mathbf{x}^*}, \mathbf{y}_D) = 1 - \Phi_{p(f(\mathbf{x}^*) | \mathbf{f}_{\setminus \mathbf{x}^*}, \mathbf{y}_D)}(f^+)$ where $f^+ \triangleq \max_{\mathbf{x}' \in \mathcal{X}^* \setminus \{\mathbf{x}^*\}} f(\mathbf{x}')$ and $\Phi_{p(f(\mathbf{x}^*) | \mathbf{f}_{\setminus \mathbf{x}^*}, \mathbf{y}_D)}(f^+)$ is the *cumulative distribution function* (c.d.f.) of the GP predictive belief $p(f(\mathbf{x}^*) | \mathbf{f}_{\setminus \mathbf{x}^*}, \mathbf{y}_D)$ evaluated at f^+ . The details of importance sampling are given in Appendix B.

Alternatively, if the number of samples to be drawn from $p(\mathbf{f}_{\mathcal{X}^*} | \mathbf{y}_D, \mathbf{x}^*)$ is allowed to vary for different \mathbf{x}^* , then we can simply draw samples from the GP predictive belief $p(\mathbf{f}_{\mathcal{X}^*} | \mathbf{y}_D)$ and group them based on their maximizers. This may result in no sample for some \mathbf{x}^* , especially if $p(\mathbf{x}^* | \mathbf{y}_D)$ is small. In this case, we simply update \mathcal{X}^* to omit such \mathbf{x}^* . It follows from Remark 2 that the construction of $\mathcal{F}_{\mathbf{x}^*}$ does not depend on the input query \mathbf{x} and is hence performed only once in each BO iteration.

Given $\mathcal{F}_{\mathbf{x}^*}$, $p(y_{\mathbf{x}} | \mathbf{y}_D, \mathbf{x}^*)$ in (6) can be approximated by

$$q_{\text{sp}}(y_{\mathbf{x}} | \mathbf{y}_D, \mathbf{x}^*) \triangleq |\mathcal{F}_{\mathbf{x}^*}|^{-1} \sum_{\mathbf{f}_{\mathcal{X}^*} \in \mathcal{F}_{\mathbf{x}^*}} p(y_{\mathbf{x}} | \mathbf{f}_{\mathcal{X}^*}, \mathbf{y}_D)$$

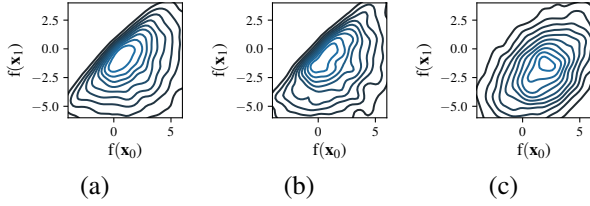


Figure 2: $p(f(\mathbf{x}_0), f(\mathbf{x}_1) | f(\mathbf{x}_0) > f(\mathbf{x}_1))$ produced by (a) rejection sampling, (b) importance sampling, and (c) EP where $p(f(\mathbf{x}_0), f(\mathbf{x}_1))$ is a standard bivariate Gaussian.

where the GP predictive belief $p(y_{\mathbf{x}} | \mathbf{f}_{\mathcal{X}^*}, \mathbf{y}_{\mathcal{D}})$ is derived in Appendix C. Then, TES_{sp} can be evaluated as follows

$$\alpha_{\text{sp}}^*(\mathbf{y}_{\mathcal{D}}, \mathbf{x}) \triangleq |\mathcal{F}_{\mathbf{x}^*}|^{-1} \sum_{\mathbf{x}^* \in \mathcal{X}^*} p(\mathbf{x}^* | \mathbf{y}_{\mathcal{D}}) \sum_{\mathbf{f}_{\mathcal{X}^*} \in \mathcal{F}_{\mathbf{x}^*}} \mathbb{E}_{p(y_{\mathbf{x}} | \mathbf{f}_{\mathcal{X}^*}, \mathbf{y}_{\mathcal{D}})} (\log q_{\text{sp}}(y_{\mathbf{x}} | \mathbf{y}_{\mathcal{D}}, \mathbf{x}^*) - \log q_{\text{sp}}(y_{\mathbf{x}} | \mathbf{y}_{\mathcal{D}}))$$

where $q_{\text{sp}}(y_{\mathbf{x}} | \mathbf{y}_{\mathcal{D}}) \triangleq \sum_{\mathbf{x}^* \in \mathcal{X}^*} p(\mathbf{x}^* | \mathbf{y}_{\mathcal{D}}) q_{\text{sp}}(y_{\mathbf{x}} | \mathbf{y}_{\mathcal{D}}, \mathbf{x}^*)$. It is amenable to stochastic optimization with respect to \mathbf{x} via stochastic gradients which are derived using mini-batches of $y_{\mathbf{x}}$ randomly sampled from the GP predictive belief $p(y_{\mathbf{x}} | \mathbf{f}_{\mathcal{X}^*}, \mathbf{y}_{\mathcal{D}})$ for all $\mathbf{f}_{\mathcal{X}^*} \in \mathcal{F}_{\mathbf{x}^*}$ ¹.

Given \mathcal{X}^* , the time complexity to evaluate $\alpha_{\text{sp}}^*(\mathbf{y}_{\mathcal{D}}, \mathbf{x})$ is $\mathcal{O}(|\mathcal{D}|^2 |\mathcal{X}^*| + |\mathcal{F}_{\mathbf{x}^*}| |\mathcal{X}^*|^2 + |\mathcal{X}^*|^3 + (|\mathcal{X}^*| + |\mathcal{D}|)^3)$ where $\mathcal{O}(|\mathcal{D}|^2 |\mathcal{X}^*|)$ is the time complexity of the GP prediction at \mathcal{X}^* given $\mathbf{y}_{\mathcal{D}}$, $\mathcal{O}(|\mathcal{F}_{\mathbf{x}^*}| |\mathcal{X}^*|^2 + |\mathcal{X}^*|^3)$ is the time complexity of sampling $\mathcal{F}_{\mathbf{x}^*}$, and $\mathcal{O}((|\mathcal{X}^*| + |\mathcal{D}|)^3)$ is the time complexity of evaluating $p(y_{\mathbf{x}} | \mathbf{f}_{\mathcal{X}^*}, \mathbf{y}_{\mathcal{D}})$.

4.2 DETERMINISTIC APPROXIMATION WITH EXPECTATION PROPAGATION (EP)

Alternatively, we can approximate $p(\mathbf{f}_{\mathcal{X}^*} | \mathbf{y}_{\mathcal{D}}, \mathbf{x}^*)$ with a Gaussian distribution so that $p(y_{\mathbf{x}} | \mathbf{y}_{\mathcal{D}}, \mathbf{x}^*)$ (6) can be evaluated in closed form. Compared with the above sampling method, its approximation cannot be refined due to its restriction to a Gaussian (see Fig. 2c). Nevertheless, it is time- and memory-efficient. Its approximation is performed by matching the moments: The mean and covariance of the Gaussian distribution are the same as that of $p(\mathbf{f}_{\mathcal{X}^*} | \mathbf{y}_{\mathcal{D}}, \mathbf{x}^*)$ being approximated. To do this, one can compute the empirical mean and covariance of the samples in $\mathcal{F}_{\mathbf{x}^*}$. An alternative without the sampling of $p(\mathbf{f}_{\mathcal{X}^*} | \mathbf{y}_{\mathcal{D}}, \mathbf{x}^*)$ is EP [Minka, 2001] which approximates $p(\mathbf{f}_{\mathcal{X}^*} | \mathbf{y}_{\mathcal{D}}, \mathbf{x}^*)$ with a Gaussian $q_{\text{ep}}(\mathbf{f}_{\mathcal{X}^*} | \mathbf{y}_{\mathcal{D}}, \mathbf{x}^*) \triangleq \mathcal{N}(\boldsymbol{\mu}_{\text{ep}}, \boldsymbol{\Sigma}_{\text{ep}})$, as detailed in Appendix D. It follows from Remark 2 that $q_{\text{ep}}(\mathbf{f}_{\mathcal{X}^*} | \mathbf{y}_{\mathcal{D}}, \mathbf{x}^*)$ does not depend on the input query \mathbf{x} and is hence computed only once in each BO iteration. Consequently, $p(y_{\mathbf{x}} | \mathbf{y}_{\mathcal{D}}, \mathbf{x}^*)$ (6) can be approximated in closed

¹Reparameterization trick is used [Kingma and Welling, 2013]

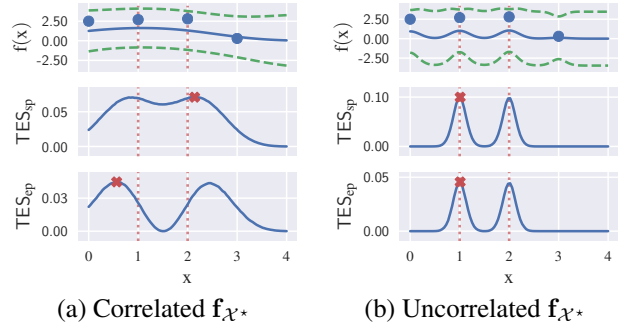


Figure 3: TES_{sp} and TES_{ep} for (a) correlated $\mathbf{f}_{\mathcal{X}^*}$ by varying the length-scale and (b) uncorrelated $\mathbf{f}_{\mathcal{X}^*}$. The top plots show the GP posterior mean as a solid blue line, uncertainty (GP posterior variance) as dashed green lines, and data points as blue points. The dotted red lines show the input positions of \mathcal{X}^* . The red crosses indicate the maximizers of acquisition functions.

form with a Gaussian:

$$q_{\text{ep}}(y_{\mathbf{x}} | \mathbf{y}_{\mathcal{D}}, \mathbf{x}^*) \triangleq \mathcal{N}(\mathbf{a}^\top \boldsymbol{\mu}_{\text{ep}} + b, \sigma_{\mathbf{x} | \mathcal{X}^*}^2 + \mathbf{a}^\top \boldsymbol{\Sigma}_{\text{ep}} \mathbf{a} + \sigma_n^2) \quad (7)$$

where the posterior variance $\sigma_{\mathbf{x} | \mathcal{X}^*}^2$ of the GP predictive belief $p(f(\mathbf{x}) | \mathbf{f}_{\mathcal{X}^*}, \mathbf{y}_{\mathcal{D}})$, \mathbf{a} , and b are given in Appendix E. Then, TES_{ep} can be evaluated as follows

$$\alpha_{\text{ep}}^*(\mathbf{y}_{\mathcal{D}}, \mathbf{x}) \triangleq \sum_{\mathbf{x}^* \in \mathcal{X}^*} p(\mathbf{x}^* | \mathbf{y}_{\mathcal{D}}) \mathbb{E}_{q_{\text{ep}}(y_{\mathbf{x}} | \mathbf{y}_{\mathcal{D}}, \mathbf{x}^*)} (\log q_{\text{ep}}(y_{\mathbf{x}} | \mathbf{y}_{\mathcal{D}}, \mathbf{x}^*) - \log q_{\text{ep}}(y_{\mathbf{x}} | \mathbf{y}_{\mathcal{D}}))$$

where $q_{\text{ep}}(y_{\mathbf{x}} | \mathbf{y}_{\mathcal{D}}) \triangleq \sum_{\mathbf{x}^* \in \mathcal{X}^*} p(\mathbf{x}^* | \mathbf{y}_{\mathcal{D}}) q_{\text{ep}}(y_{\mathbf{x}} | \mathbf{y}_{\mathcal{D}}, \mathbf{x}^*)$. Hence, it is amenable to stochastic optimization like the optimization of $\alpha_{\text{sp}}^*(\mathbf{y}_{\mathcal{D}}, \mathbf{x})$.

Given \mathcal{X}^* , the time complexity of evaluating $\alpha_{\text{ep}}^*(\mathbf{y}_{\mathcal{D}}, \mathbf{x})$ is mainly due to the EP procedure which incurs $\mathcal{O}(n_{\text{ep}} |\mathcal{X}^*| |\mathcal{D}|^3)$ where n_{ep} is the number of EP iterations.

Fig. 3b shows different acquisition functions in a synthetic example. Recall that TES measures the information gain on $\mathbf{x}^* \in \mathcal{X}^*$ by selecting the next input query $\mathbf{x} \in \mathcal{X}$ for evaluating f . So, for any \mathbf{x} that is far from \mathcal{X}^* , $f(\mathbf{x})$ has little correlation with $\mathbf{f}_{\mathcal{X}^*}$ and its TES value should be close to 0. This observation holds for TES_{sp} and TES_{ep} in Fig. 3b. Furthermore, TES_{sp} and TES_{ep} select an input in \mathcal{X}^* when $\mathbf{f}_{\mathcal{X}^*}$ are not strongly correlated, as shown in Fig. 3b; red crosses are input queries \mathbf{x} that optimize the acquisition functions. When there is a strong correlation within $\mathbf{f}_{\mathcal{X}^*}$, TES_{sp} and TES_{ep} can exploit the correlation to select an informative input query \mathbf{x} around \mathcal{X}^* , as shown in Fig. 3a. As a result, to optimize TES, we should initialize the optimization with inputs in \mathcal{X}^* . Another observation is that TES_{ep} exploits more than TES_{sp} , as shown in Appendix F.

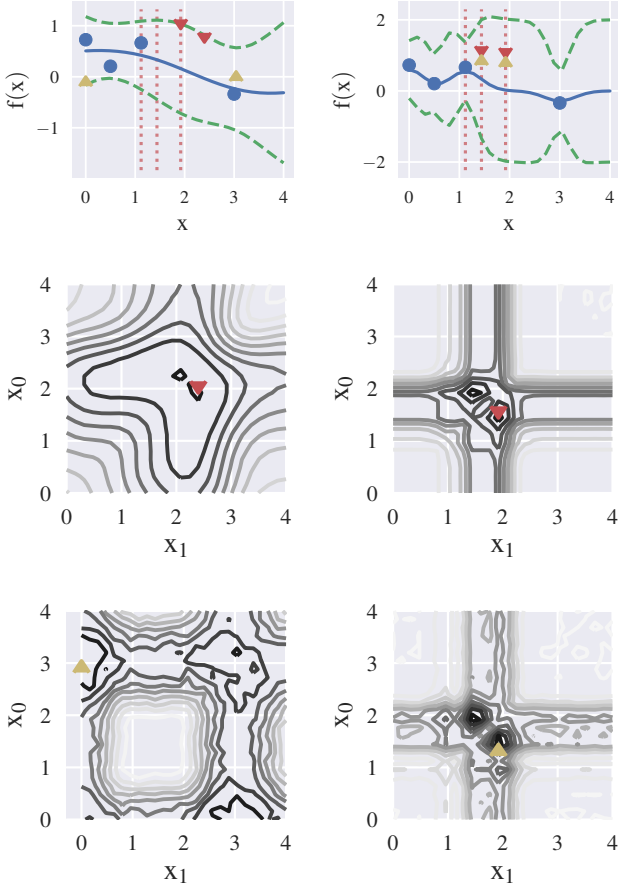


Figure 4: TES_{sp} and TES_{ep} with batches $\mathcal{B} = \{x_0, x_1\}$ for correlated \mathcal{X}^* (left column) and uncorrelated \mathcal{X}^* (right column) by varying the length-scale. The top, middle, and bottom plots are GP posteriors, TES_{sp} , and TES_{ep} , respectively. The notations are similar to Fig. 3. Yellow caret-ups and red caret-downs denote the input query \mathbf{x} for TES_{ep} and TES_{sp} , respectively. The bottom 4 plots are not perfectly symmetrical about $x_0 = x_1$ due to sampling.

4.3 GENERALIZATION TO BATCH BO

In batch BO where a batch of inputs, denoted as \mathcal{B} , is selected at each BO iteration, TES is defined as the amount of information gain about \mathbf{x}^* through observing the noisy observations $\mathbf{y}_{\mathcal{B}}$ of the batch \mathcal{B} :

$$\alpha^*(\mathbf{y}_{\mathcal{D}}, \mathcal{B}) = H(\mathbf{y}_{\mathcal{B}}|\mathbf{y}_{\mathcal{D}}) - \mathbb{E}_{p(\mathbf{x}^*|\mathbf{y}_{\mathcal{D}})}[H(\mathbf{y}_{\mathcal{B}}|\mathbf{y}_{\mathcal{D}}, \mathbf{x}^*)]. \quad (8)$$

It can be approximated either stochastically with samples or deterministically with EP. The Gaussian predictive entropy $H(\mathbf{y}_{\mathcal{B}}|\mathbf{y}_{\mathcal{D}})$ naturally encourages the diversity of a batch, e.g., sampling of highly correlated observations in a batch (i.e., redundant information) is discouraged. This is different from several existing works such as BBO-LP [González et al., 2016] which requires a local penalization term based on the objective function’s Lipschitz constant in addition to the acquisition function. Increasing the batch

size from one (in the previous section) to multiple input queries changes neither the set of samples $\mathcal{F}_{\mathbf{x}^*}$ approximating $p(\mathbf{f}_{\mathcal{X}^*}|\mathbf{y}_{\mathcal{D}}, \mathbf{x}^*)$ in the stochastic approximation nor the approximate posterior $q_{\text{ep}}(\mathbf{f}_{\mathcal{X}^*}|\mathbf{y}_{\mathcal{D}}, \mathbf{x}^*)$ in the deterministic approximation. Both approximations are applied in the same manner as those in the previous section (i.e., only once prior to the acquisition optimization in each BO iteration). The only difference in batch BO for $|\mathcal{B}| > 1$ is that $\mathbf{y}_{\mathcal{B}}$ is a multivariate r.v., so $q_{\text{sp}}(\mathbf{y}_{\mathcal{B}}|\mathbf{y}_{\mathcal{D}}, \mathbf{x}^*)$ and $q_{\text{ep}}(\mathbf{y}_{\mathcal{B}}|\mathbf{y}_{\mathcal{D}}, \mathbf{x}^*)$ are a mixture of multivariate Gaussian distributions and a multivariate Gaussian distribution, respectively. Still, samples of $\mathbf{y}_{\mathcal{B}}$ from these distributions can be drawn efficiently for the stochastic optimization of the acquisition function.

Batch TES with more than one input in the batch still exhibits similar behavior to that with only an input in the batch. That is, when $\mathbf{f}_{\mathcal{X}^*}$ are not strongly correlated, a sensible strategy is to sample observations at inputs in \mathcal{X}^* to identify \mathbf{x}^* (the right column of Fig. 4). In contrast, when there is a strong correlation in $\mathbf{f}_{\mathcal{X}^*}$, TES can exploit the correlation to sample informative observations at inputs around \mathcal{X}^* (the left column of Fig. 4). Based on this observation, we also initialize the acquisition optimization with a batch of inputs in \mathcal{X}^* .

Note that $\alpha^*(\mathbf{y}_{\mathcal{D}}, \mathcal{B})$ is the information gain about \mathbf{x}^* through observing $\mathbf{y}_{\mathcal{B}}$. Thus, it is comparable between batches of different sizes. As illustrated in Appendix G, the size of \mathcal{X}^* should be at least the batch size to make the most out of the observations $\mathbf{y}_{\mathcal{B}}$. Besides, we can select the batch size adaptively at each BO iteration based on a trade-off between the increase in the information gain and the cost of an observation.

5 EXPERIMENTS

Several existing acquisition functions and TES are evaluated empirically on synthetic functions and real-world optimization problems. We first show the competitive performance of TES on simple synthetic functions in Sections 5.1 and 5.2. Subsequently in Section 5.3, we demonstrate that it can keep up the good performance in high-dimensional complex real-world problems.

The predictive performance of BO algorithms is measured by the *immediate regret* (IR), which is defined as $f^* - f(\bar{\mathbf{x}})$ where $\bar{\mathbf{x}}$ is the maximizer of the posterior mean given $\mathbf{y}_{\mathcal{D}}$, and f^* is the true global maximum of the objective function. To account for the randomness in the optimization and the sampling of noisy observations, experiments are repeated over several random runs. The natural logarithm of the average of IR over these random runs is reported. The box plots of IR are shown in Appendix H. All experiments are initialized with a training set of 2 data samples unless otherwise specified. These synthetic and the real-world functions are

each modeled as a sample of a GP whose kernel hyperparameters are learned using maximum likelihood estimation [Rasmussen and Williams, 2006]. Adam [Kingma and Ba, 2015] is used to optimize TES with 300 iterations. At each stochastic iteration of Adam, 1000 samples of $y_{\mathbf{x}}$ (or $y_{\mathcal{B}}$) are drawn. The experiments are performed on an Intel Xeon E5-2683 v4 CPU and an NVIDIA GTX 1080 GPU.

5.1 BATCH BO WITH $|\mathcal{B}| = 1$

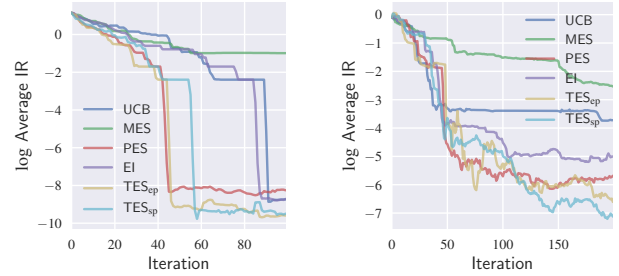
We first consider BO algorithms that select 1 input query in each BO iteration. We compare TES_{ep} and TES_{sp} with existing algorithms: EI [Mockus et al., 1978], UCB [Srinivas et al., 2010], MES [Wang and Jegelka, 2017], and PES [Hernández-Lobato et al., 2014]. There are 5 inputs in \mathcal{X}^* .

In the first experiment, we draw a function sample from a GP defined in $[0, 10]^2$ with hyperparameters $\sigma_s^2 = 2$, $l_1 = l_2 = 1$, which is named the *GP sampled function*. The noise variance to generate observations is $\sigma_n^2 = 10^{-4}$. The average results over 10 random runs are shown in Fig. 5a.: TES_{sp} and TES_{ep} outperform other acquisition functions by converging to a better global maximum. Surprisingly, TES_{ep} converges faster than TES_{sp} . It could be because it exploits more than TES_{sp} as mentioned above (illustrated in Appendix F). The performance of PES closely matches those of TES acquisition functions, while it takes longer for UCB and EI to converge. MES does not perform as well as other acquisition functions that concentrate on the information gain of the maximizers (TES and PES). Experimental results on Branin and Hartmann-3D [Lizotte, 2008] are in Appendix H.

The *log10P* is a function that models the phosphorus field of Broom’s Barn farm [Webster and Oliver, 2007] spatially distributed over a 1200m by 680m region discretized into a 31×18 grid of sampling locations. The logarithm to the base 10 values of the amount of phosphorus are recorded. Fig. 5b shows the average results over 5 random runs for the log10P with the input rescaled to $[0, 1]^2$. To generate a continuous function from the dataset, we fit a GP on the dataset. This results in a GP posterior with the dataset as training data. The hyperparameters are $\sigma_s^2 = 7.669 \times 10^{-2}$, $l_1 = 9.106 \times 10^{-2}$, $l_2 = 3.971 \times 10^{-2}$, and $\sigma_n^2 = 2.511 \times 10^{-2}$. The posterior mean of this GP is taken as the objective function to optimize. Note that the noise variance in this experiment is larger than that of the previous experiments. Hence, it is advantageous for information-based acquisition functions such as TES and PES as they tend to explore more than the others. Fig. 5b shows that TES_{sp} converges to a better maximum than others.

5.2 BATCH BO WITH $|\mathcal{B}| > 1$

In this section, we consider BO of batch size larger than 1. We compare our algorithms with QEI-CL [Ginsbourger



(a) The GP sampled function. (b) log10P.

Figure 5: BO with $|\mathcal{B}| = 1$.

et al., 2010], BUCB [Desautels et al., 2014], QEI [Marmin et al., 2015], and LP-EI, which is the BBO-LP [González et al., 2016] by plugging in EI.

Figs. 6a-c shows BO results with $|\mathcal{B}| = 3$ over 3 random experiments for the GP sampled function, Hartmann-4D, and log10P. The size of \mathcal{X}^* is set to 5. Overall, both TES_{sp} and TES_{ep} achieve good performance by converging to good maxima. The difference in the performance can be explained by the fact that TES can exploit the maximizer samples from the GP posterior to quickly explore the function domain.

To illustrate the scalability of TES_{ep} , we experiment with large batch sizes: $|\mathcal{B}| = 10, 20$ and 30 for the GP sampled function, and $|\mathcal{B}| = 10, 20$ and 40 for the log10P in Figs. 6d-i. The size of \mathcal{X}^* is set to be the same as that of \mathcal{B} . It is to make the most of the information in the batch as explained in Appendix G. In the GP sampled function experiments (Figs. 6d-f), though other acquisition functions converge faster than TES_{ep} , they seem to be stuck in sub-optimal maxima except for LP-EI in Fig. 6d. At the initial iterations with little observations, \mathcal{X}^* are likely to scatter over the whole input space, so TES_{ep} explores more than other acquisition functions. At the latter iterations with a sufficient number of observations, TES can exploit to discover a good maximum. In the log10P experiments with noisier observations (Figs. 6g-i), we observe that only QEI shows similar performance to our TES_{ep} , though our TES_{ep} outperforms for batch sizes of 20 and 40. It could be because both QEI and TES_{ep} jointly optimize the batch, which is an advantage compared with other greedy batch acquisition functions in these experiments. Another observation for TES_{ep} in Figs. 6d-i is that increasing the batch size improves the efficiency of BO with respect to the number of iterations.

5.3 OPTIMIZING COMPLEX REAL-WORLD PROBLEMS

In this section, we aim to illustrate that TES works well in real life. We apply our more scalable TES_{ep} algorithm to two real-world scenarios that can be regarded as black-box optimization problems. Other batch BO acquisition functions mentioned previously are evaluated as comparisons.

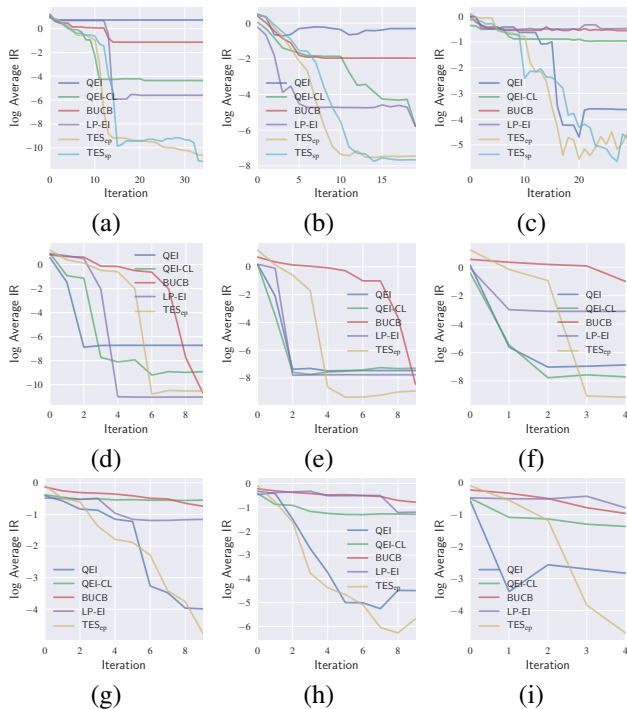
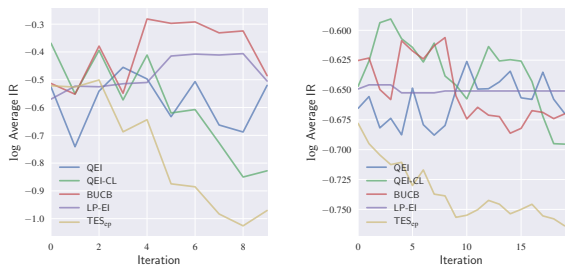


Figure 6: Batch BO: $|\mathcal{B}| = 3$ on (a) the GP sampled function, (b) Hartmann-4D, and (c) the log10P; \mathcal{B} of size (d) 10, (e) 20, and (f) 30 on the GP sampled function; and \mathcal{B} of size (g) 10, (h) 20, and (i) 40 on the log10P.

The first real-world experiment is motivated by the computationally expensive procedure of neural architecture search in convolutional neural networks (CNNs). We train a 2-layer CNN on the CIFAR-10 dataset and aim to efficiently search for the best combination of 7 hyperparameters: batch size, number of convolutional filters, number of dense units, dropout rate, L2 regularizer, learning rate, and learning rate decay. The bounds are $[16, 512]$, $[6, 256]$, $[128, 2048]$, $[0, 0.75]$, $[10^{-7}, 10^{-3}]$, $[10^{-7}, 10^{-1}]$, $[10^{-7}, 10^{-3}]$ respectively and the first three hyperparameters are integers. The IR is calculated by 1 minus the test accuracy.

Fig. 7a shows BO results with $|\mathcal{B}| = 20$ over 5 random experiments for CIFAR-10. The initial training set consists



(a) The CIFAR-10. (b) The face attack function.

Figure 7: BO results on real-world problems.

of 10 data samples. We observe greater fluctuations in the regret for all methods because the outputs of complex problems generally contain larger noises than those of synthetic functions. Overall, TES_{ep} still converges the fastest to the lowest regret. All other methods except QEI-CL are stuck in suboptimal maxima obtained at the beginning few iterations. This empirical experiment shows that TES_{ep} can explore multi-dimensional search spaces well in real life and find the best maximum among all the methods tested.

In the second real-world scenario, we synthesize ‘physically realizable’ faces to fool a black-box *face recognition system* (FRS) [Saha and Sim, 2020]. We formalize the problem as an optimization task, illustrating the potential privacy and security risks of such systems under malicious attacks. More precisely, by altering some parameters representing specific facial features, the attacker’s face is synthesized such that it minimizes the distance d to a victim’s face perceived by an FRS. The victim could be any person registered in the FRS, while the attacker is not. The lower the distance, the more likely the attacker could successfully fool and bypass the security of the system. Therefore, the performance metric is the distance d .

In our experiments, we use the Python *face_recognition* library [Geitgey, 2017] as the target black-box FRS and alter 8 parameters linked to facial features such as facial hair (e.g., beard, mustache), marks (e.g., forehead scar, cheek scar, mole), eyeglasses (e.g., squared, round) and expressions (e.g., smile). Note that bounds are imposed on the facial feature parameters to ensure more realistic appearances of the synthesized faces, i.e., ‘physically realizable’ faces. The multimodal discriminant analysis (MMDA) face synthesizer model described by Sim and Zhang [2015] is used. The library outputs a 128-dimension face encoding for each input face image and considers two faces as a pair of ‘match’ if the Euclidean distance between the two encodings is lower than a threshold $\theta = 0.5$. Given a gallery of 41 face images selected from the CMU Multi-PIE dataset [Gross et al., 2010] which consists of 36 males and 5 females, our objective is to modify an attacker’s face in Fig. 8a so that its distance to a registered face in the gallery is minimized. We hope the distance to go below θ for a successful break-in attack.

Fig. 7b shows BO results with $|\mathcal{B}| = 20$ over 5 random experiments for the face attack function. The initial training size is again 10. Similar to what we observed in the CIFAR-10 experiment, TES_{ep} achieves the lowest regret while other methods only obtain suboptimal maxima. TES_{ep} explores over the whole input space for about 10 iterations (or 200 samples) before turning to exploitation to discover a good maximum. Similar behavior is seen in Fig. 6c. Additionally, the box plots of IR are included in Appendix H.

We show the visualization of a typical result after iterations of optimization using BO with our TES_{ep} acquisition func-

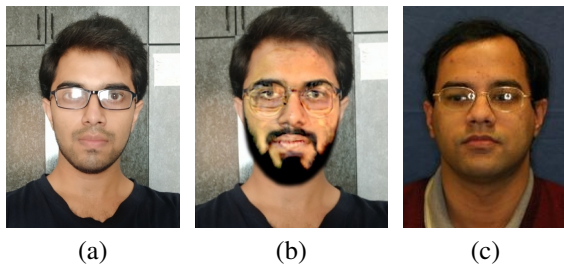


Figure 8: Visualization of a typical result. (a) is the original image of the attacker. (b) shows the synthesized image of the attacker. (c) is the ‘matched’ image in the gallery.

tion in Fig. 8. In this particular case, the attacker’s face is modified to Fig. 8b such that the distance to a registered face encoding in the gallery is optimized to 0.4534, implying that the privacy and security of the FRS have been compromised since the d is smaller than the threshold $\theta = 0.5$. Notably, the modified face is ‘physically realizable’, i.e., the attacker could grow a beard, change the glasses, put on some scars makeup and smile to fool the FRS. The victim’s face is shown in Fig. 8c. Therefore, our proposed algorithm is effective in such real-world optimization problems.

6 CONCLUSION

This paper presents *trusted-maximizers entropy search* (TES) for performing Bayesian optimization and its batch variant. Both stochastic approximation with sampling and deterministic approximation with a Gaussian distribution are investigated, which results in TES_{sp} and TES_{ep} , respectively. While TES_{sp} produces a better approximation of the posterior belief, TES_{ep} is less expensive to evaluate. Therefore, TES_{ep} can be scaled to BO with large batch sizes. TES_{sp} and TES_{ep} achieve favorable performance compared with existing acquisition functions in experiments with synthetic benchmark objective functions and complex real-world optimization problems.

Author Contributions

Quoc Phong Nguyen and Zhaoxuan Wu contributed equally to this paper.

Acknowledgements

This research is supported by the National Research Foundation, Singapore under its Strategic Capability Research Centres Funding Initiative and the Singapore Ministry of Education Academic Research Fund Tier 1. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

References

- E. Brochu, T. Brochu, and N. de Freitas. A Bayesian interactive optimization approach to procedural animation design. In *Proc. ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 103–112, 2010a.
- 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. *arXiv:1012.2599*, 2010b.
- Z. Dai, H. Yu, B. K. H. Low, and P. Jaillet. Bayesian optimization meets Bayesian optimal stopping. In *Proc. ICML*, pages 1496–1506, 2019.
- Z. Dai, Y. Chen, B. K. H. Low, P. Jaillet, and T. Ho. R2-B2: Recursive reasoning-based Bayesian optimization for no-regret learning in games. In *Proc. ICML*, 2020.
- E. A. Daxberger and B. K. H. Low. Distributed batch Gaussian process optimization. In *Proc. ICML*, pages 951–960, 2017.
- T. Desautels, A. Krause, and J. Burdick. Parallelizing exploration-exploitation tradeoffs in Gaussian process bandit optimization. *JMLR*, 15(1):3873–3923, 2014.
- A. Geitgey. Face recognition, 2017. URL https://github.com/ageitgey/face_recognition.
- D. Ginsbourger, R. Le Riche, and L. Carraro. Kriging is well-suited to parallelize optimization. In *Computational intelligence in expensive optimization problems*, pages 131–162. Springer, 2010.
- J. González, Z. Dai, P. Hennig, and N. Lawrence. Batch Bayesian optimization via local penalization. In *Proc. AISTATS*, pages 648–657, 2016.
- Ralph Gross, Iain Matthews, Jeffrey Cohn, Takeo Kanade, and Simon Baker. Multi-pie. *Image and vision computing*, 28(5):807–813, 2010.
- P. Hennig and C. J. Schuler. Entropy search for information-efficient global optimization. *JMLR*, pages 1809–1837, 2012.
- J. M. Hernández-Lobato, M. W. Hoffman, and Z. Ghahramani. Predictive entropy search for efficient global optimization of black-box functions. In *Proc. NIPS*, pages 918–926, 2014.
- T. N. Hoang, Q. M. Hoang, R. Ouyang, and B. K. H. Low. Decentralized high-dimensional Bayesian optimization with factor graphs. In *Proc. AAAI*, pages 3231–3238, 2018.
- M. W. Hoffman and Z. Ghahramani. Output-space predictive entropy search for flexible global optimization. In *NIPS workshop on Bayesian Optimization*, 2015.

- D. Kharkovskii, Z. Dai, and B. K. H. Low. Private out-sourced Bayesian optimization. In *Proc. ICML*, 2020.
- D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *Proc. ICLR*, 2015.
- D. P. Kingma and M. Welling. Auto-encoding variational Bayes. *arXiv:1312.6114*, 2013.
- H. J. Kushner. A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise. *Journal of basic engineering*, 86(1):97–106, 1964.
- C. K. Ling, B. K. H. Low, and P. Jaillet. Gaussian process planning with Lipschitz continuous reward functions: Towards unifying Bayesian optimization, active learning, and beyond. In *Proc. AAAI*, pages 1860–1866, 2016.
- D. J. Lizotte. *Practical Bayesian optimization*. PhD thesis, University of Alberta, Canada, 2008.
- S. Marmin, C. Chevalier, and D. Ginsbourger. Differentiating the multipoint expected improvement for optimal batch design. In *International Workshop on Machine Learning, Optimization and Big Data*, pages 37–48. Springer, 2015.
- T. P. Minka. *A family of algorithms for approximate Bayesian inference*. PhD thesis, Massachusetts Institute of Technology, 2001.
- J. Mockus, V. Tiesis, and A. Žilinskas. The application of Bayesian methods for seeking the extremum. In L. C. W. Dixon and G. P. Szegö, editors, *Towards Global Optimization 2*, pages 117–129. North-Holland Publishing Company, 1978.
- Q. P. Nguyen, B. K. H. Low, and P. Jaillet. An information-theoretic framework for unifying active learning problems. In *Proc. AAAI*, pages 9126–9134, 2021a.
- Q. P. Nguyen, S. Tay, B. K. H. Low, and P. Jaillet. Top-k ranking Bayesian optimization. In *Proc. AAAI*, pages 9135–9143, 2021b.
- C. E. Rasmussen and C. K. I. Williams. *Gaussian processes for machine learning*. MIT Press, 2006.
- B. Ru, M. McLeod, D. Granzio, and M. A. Osborne. Fast information-theoretic Bayesian optimisation. In *Proc. ICML*, pages 4381–4389, 2018.
- S. Saha and T. Sim. Is face recognition safe from realizable attacks? In *International Joint Conference on Biometrics*, 2020.
- A. Shah and Z. Ghahramani. Parallel predictive entropy search for batch global optimization of expensive objective functions. In *Proc. NIPS*, pages 3330–3338, 2015.
- B. Shahriari, K. Swersky, Z. Wang, R. Adams, and N. de Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1): 148–175, 2015.
- T. Sim and L. Zhang. Controllable face privacy. In *2015 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*, volume 04, pages 1–8, 2015.
- J. Snoek, H. Larochelle, and R. Adams. Practical bayesian optimization of machine learning algorithms. In *Proc. NIPS*, pages 2951–2959, 2012.
- N. Srinivas, A. Krause, S. Kakade, and M. Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proc. ICML*, pages 1015–1022, 2010.
- Z. Wang and S. Jegelka. Max-value entropy search for efficient Bayesian optimization. In *Proc. ICML*, pages 3627–3635, 2017.
- R. Webster and M. Oliver. *Geostatistics for environmental scientists*. John Wiley & Sons, 2007.
- Y. Zhang, T. N. Hoang, B. K. H. Low, and M. Kankanhalli. Information-based multi-fidelity Bayesian optimization. In *Proc. NeurIPS Workshop on Bayesian Optimization*, 2017.
- Y. Zhang, Z. Dai, and B. K. H. Low. Bayesian optimization with binary auxiliary information. In *Proc. UAI*, pages 1222–1232, 2019.