
Private Outsourced Bayesian Optimization

Dmitrii Kharkovskii¹ Zhongxiang Dai¹ Bryan Kian Hsiang Low¹

Abstract

This paper presents the *private-outsourced-Gaussian process-upper confidence bound* (PO-GP-UCB) algorithm, which is the first algorithm for privacy-preserving *Bayesian optimization* (BO) in the *outsourced* setting with a provable performance guarantee. We consider the outsourced setting where the entity holding the dataset and the entity performing BO are represented by different parties, and the dataset cannot be released non-privately. For example, a hospital holds a dataset of sensitive medical records and outsources the BO task on this dataset to an industrial AI company. The key idea of our approach is to make the BO performance of our algorithm similar to that of non-private GP-UCB run using the original dataset, which is achieved by using a random projection-based transformation that preserves both privacy and the pairwise distances between inputs. Our main theoretical contribution is to show that a regret bound similar to that of the standard GP-UCB algorithm can be established for our PO-GP-UCB algorithm. We empirically evaluate the performance of our PO-GP-UCB algorithm with synthetic and real-world datasets.

1. Introduction

Bayesian optimization (BO) has become an increasingly popular method for optimizing highly complex black-box functions mainly due to its impressive sample efficiency. Such optimization problems appear frequently in various applications such as automated machine learning (ML), robotics, sensor networks, among others (Shahriari et al., 2016). However, despite its popularity, the classical setting of BO does not account for privacy issues, which arise due to the widespread use of ML models in applications dealing

with sensitive datasets such as health care (Yu et al., 2013), insurance (Chong et al., 2005) and fraud detection (Ngai et al., 2011). A natural solution is to apply the cryptographic framework of *differential privacy* (DP) (Dwork et al., 2016), which has become the state-of-the-art technique for private data release and has been widely adopted in ML (Sarwate & Chaudhuri, 2013).

To this end, a recent work (Kusner et al., 2015) proposed a DP variant of the *Gaussian process-upper confidence bound* (GP-UCB) algorithm, which is a well-known BO algorithm with theoretical performance guarantee (Srinivas et al., 2010). Kusner et al. (2015) consider the common BO task of hyperparameter tuning for ML models and introduce methods for privatizing the outputs of the GP-UCB algorithm (the optimal input hyperparameter setting found by the algorithm and the corresponding output measurement) by releasing them using standard DP mechanisms. However, in many scenarios, BO is performed in the *outsourced* setting, in which the entity holding the sensitive dataset (referred to as the *curator* hereafter) and the entity performing BO (referred to as the *modeler* hereafter) are represented by different parties with potentially conflicting interests. In recent years, such modelers (i.e., commercial companies) providing general-purpose optimization as a service have become increasingly prevalent, such as SigOpt which uses BO as a commercial service for black-box global optimization by providing query access to the users (Dewancker et al., 2016), and Google Cloud AutoML which offers the optimization of the architectures of neural networks as a cloud service. Unfortunately, the approach of Kusner et al. (2015) requires the modeler and the curator to be represented by the same entity and therefore both parties must have complete access to the sensitive dataset and full understanding of the BO algorithm, thus rendering it inapplicable in the outsourced setting. Some examples of such settings are given below:

(a) A hospital is trying to find out which patients are likely to be readmitted soon based on the result of an expensive medical test (Yu et al., 2013). Due to cost and time constraints, the hospital (curator) is only able to perform the test for a limited number of patients, and thus outsources the task of selecting candidate patients for testing to an industrial AI company (modeler). In this case, the inputs to BO are medical records of individual patients and the function to maximize (the output measurement) is the outcome of the

¹Department of Computer Science, National University of Singapore, Republic of Singapore. Correspondence to: Bryan Kian Hsiang Low <lowkh@comp.nus.edu.sg>.

medical test for different patients, which is used to assess the possibility of readmission. The hospital is unwilling to release the medical records, while the AI company does not want to share the details of their proprietary algorithm.

(b) A bank aims to identify the loan applicants with the highest return on investment and outsources the task to a financial AI consultancy. In this case, each input to BO is the data of a single loan applicant and the output measurement to be maximized is the return on investment for different applicants. The bank (curator) is unable to disclose the raw data of the loan applicants due to privacy and security concerns, whereas the AI consultancy (modeler) is unwilling to share the implementation of their selection strategy.

(c) A real estate agency attempts to locate the cheapest private properties in an urban city. Since evaluating every property requires sending an agent to the corresponding location, the agency (curator) outsources the task of selecting candidate properties for evaluation to an AI consultancy (modeler) to save resources. Each input to BO is a set of features representing a single property and the function to minimize (the output measurement) is the evaluated property price. The agency is unable to disclose the particulars of their customers due to legal implications, while the AI consultancy refuses to share their decision-making algorithm.

In all of these scenarios, the curator is unable to release the original dataset due to privacy concerns, and therefore has to provide a transformed privatized dataset to the modeler. Then, the modeler can perform BO (specifically, the GP-UCB algorithm) on the transformed dataset (the detailed problem setting is described in Section 2 and illustrated in Fig. 1). A natural choice for the privacy-preserving transformation is to use standard DP methods such as the Laplace or Gaussian mechanisms (Dwork & Roth, 2014). However, the theoretically guaranteed convergence of the GP-UCB algorithm (Srinivas et al., 2010) is only valid if it is run using the original dataset. Therefore, as a result of the privacy-preserving transformation required in the outsourced setting, it is unclear whether the theoretical guarantee of GP-UCB can be preserved and thus whether reliable performance can be delivered. This poses an interesting research question: *How do we design a privacy-preserving algorithm for outsourced BO with a provable performance guarantee?*

To address this challenge, we propose the *private-outsourced-Gaussian process-upper confidence bound* (PO-GP-UCB) algorithm (Section 3), which is the first algorithm for BO with DP in the outsourced setting with a provable performance guarantee¹. The key idea of our approach is to make the GP predictions and hence the BO performance

¹While the setting of the recent work of Nguyen et al. (2018) resembles ours, the authors use a self-proposed notion of privacy instead of the widely recognized DP and protect the privacy of only the output measurements.




Figure 1. Visual illustration of the problem setting.

of our algorithm similar to those of non-private GP-UCB run using the original dataset. To achieve this, instead of standard DP methods, we use a privacy-preserving transformation based on random projection (Johnson & Lindenstrauss, 1984), which approximately preserves the pairwise distances between inputs. We show that preserving the pairwise distances between inputs leads to preservation of the GP predictions and therefore the BO performance in the outsourced setting (compared with the standard setting of running non-private GP-UCB on the original dataset). Our main theoretical contribution is to show that a regret bound similar to that of the standard GP-UCB algorithm can be established for our PO-GP-UCB algorithm. We empirically evaluate the performance of our PO-GP-UCB algorithm with synthetic and real-world datasets (Section 4).

2. Notations and Preliminaries

Problem Setting. Privacy-preserving BO in the outsourced setting involves two parties: the *curator* who holds the sensitive dataset (e.g., a list of medical records), and the *modeler* who performs the outsourced BO on the transformed dataset provided by the curator (see Fig. 1 for a visual illustration of this setting). The curator holds the original dataset represented as a set $\mathcal{X} \subset \mathbb{R}^d$ formed by n d -dimensional inputs. The curator and the modeler intend to maximize an unknown expensive-to-evaluate objective function f defined over \mathcal{X} . At the beginning, the curator performs a privacy-preserving transformation of the original dataset \mathcal{X} to obtain a transformed dataset $\mathcal{Z} \subset \mathbb{R}^r$ formed by n r -dimensional inputs. As a result, every original input $x \in \mathcal{X}$ has an image, which is the corresponding transformed input $z \in \mathcal{Z}$. Then, the curator releases the transformed dataset \mathcal{Z} to the modeler, who can subsequently start to run the BO algorithm on \mathcal{Z} . In each iteration $t = 1, \dots, T$, the modeler selects a transformed input $z_t \in \mathcal{Z}$ to query and notifies the curator about the choice of z_t . Next, the curator identifies x_t which is the preimage of z_t under the privacy-preserving transformation², and then computes $f(x_t)$ to yield a noisy output measurement: $y_t \triangleq f(x_t) + \epsilon_{GP}$, in which $\epsilon_{GP} \sim \mathcal{N}(0, \sigma_n^2)$ is a zero-mean Gaussian noise with noise variance σ_n^2 . We

²We assume that \mathcal{X} and \mathcal{Z} describe the entire optimization domain, i.e., every $z_t \in \mathcal{Z}$ has a preimage $x_t \in \mathcal{X}$.

assume that y_t is unknown to the curator in advance and is computed only when requested by the modeler, which is reasonable in all motivating scenarios in Section 1. The curator then sends y_t to the modeler for performing the next iteration of BO. We have assumed that in contrast to the input x_t , the noisy output measurement y_t does not contain sensitive information and can thus be non-privately released. This assumption is reasonable in our setting, e.g., if y_t represents the outcome of a medical test, revealing y_t does not unveil the identity of the patient. We leave the extension of privately releasing y_t for future work (see Section 5).

Differential Privacy (DP). Differential privacy (Dwork et al., 2016) has become the state-of-the-art technique for private data release. DP is a cryptographic framework which provides rigorous mathematical guarantees on privacy, typically by adding some random noise during the execution of the data release algorithm. DP has been widely adopted by the ML community (see the work of Sarwate & Chaudhuri (2013) for a detailed survey). Intuitively, DP promises that changing a single input of the dataset imposes only a small change in the output of the data release algorithm, hence the output does not depend significantly on any individual input. As a result, an attacker is not able to tell if an input is changed in the dataset just by looking at the output of the data release algorithm. To define DP, we first need to introduce the notion of *neighboring* datasets. Following the prior works on DP (Blocki et al., 2012; Hardt & Roth, 2012), we define two neighboring datasets as those differing only in a single row (i.e., a single input) with the norm of the difference bounded by 1:

Definition 1. Let $\mathcal{X}, \mathcal{X}' \in \mathbb{R}^{n \times d}$ denote two datasets viewed as matrices³ with d -dimensional inputs $\{x_{(i)}\}_{i=1}^n$ and $\{x'_{(i)}\}_{i=1}^n$ as rows respectively. We call datasets \mathcal{X} and \mathcal{X}' *neighboring* if there exists an index $i^* \in 1, \dots, n$ such that $\|x_{(i^*)} - x'_{(i^*)}\| \leq 1$, and $\|x_{(j)} - x'_{(j)}\| = 0$ for any index $j \in 1, \dots, n, j \neq i^*$.

A randomized algorithm is differentially private if, for any two neighboring datasets, the distributions of the outputs of the algorithm calculated on these datasets are similar:

Definition 2. A randomized algorithm \mathcal{M} is (ϵ, δ) -differentially private for $\epsilon > 0$ and $\delta \in (0, 1)$ if, for all $O \subset \text{Range}(\mathcal{M})$ (where $\text{Range}(\mathcal{M})$ is the range of the outputs of the randomized algorithm \mathcal{M}) and for all neighboring datasets \mathcal{X} and \mathcal{X}' , we have that

$$P(\mathcal{M}(\mathcal{X}) \in O) \leq \exp(\epsilon) \cdot P(\mathcal{M}(\mathcal{X}') \in O) + \delta.$$

Note that the definition above is symmetric in terms of \mathcal{X} and \mathcal{X}' . The DP parameters ϵ, δ control the *privacy-utility*

³ We slightly abuse the notation and view the dataset \mathcal{X} (\mathcal{Z}) as an $n \times d$ ($n \times r$) matrix where each of the n rows corresponds to an original (transformed) input.

trade-off: The smaller they are, the tighter the privacy guarantee is, at the expense of lower accuracy due to increased amount of noise required to satisfy DP. The state-of-the-art works on the application of DP in ML (Abadi et al., 2016; Foulds et al., 2016; Papernot et al., 2017) use the values of ϵ in the single-digit range, while the value of δ is usually set to be smaller than $1/n$ (Dwork & Roth, 2014). Refer to the work of Dwork & Roth (2014) for more details about DP.

Bayesian Optimization (BO). We consider the problem of sequentially maximizing an unknown objective function $f : \mathcal{X} \rightarrow \mathbb{R}$, in which $\mathcal{X} \subset \mathbb{R}^d$ denotes a domain of d -dimensional inputs. We consider the domain to be discrete for simplicity. In the classical setting of BO, in each iteration $t = 1, \dots, T$, an unobserved input $x_t \in \mathcal{X}$ is selected to query by maximizing an *acquisition function* (AF), yielding a noisy output measurement $y_t \triangleq f(x_t) + \epsilon_{GP}$, in which $\epsilon_{GP} \sim \mathcal{N}(0, \sigma_n^2)$ is a zero-mean Gaussian noise with noise variance σ_n^2 . The AF should be designed to allow us to approach the global maximum $f(x^*)$ rapidly, in which $x^* \triangleq \operatorname{argmax}_{x \in \mathcal{X}} f(x)$. This can be achieved by minimizing a standard BO objective such as *regret*. The notion of regret intuitively refers to a loss in reward resulting from not knowing x^* beforehand. Formally, the *instantaneous regret* incurred in iteration t is defined as $r_t \triangleq f(x^*) - f(x_t)$. *Cumulative regret* is defined as the sum of all instantaneous regrets, i.e., $R_T \triangleq \sum_{t=1}^T r_t$, and *simple regret* is defined as the minimum among all instantaneous regrets, i.e., $S_T \triangleq \min_{t=1, \dots, T} r_t$. It is often desirable for a BO algorithm to be asymptotically *no-regret*, i.e., $\lim_{T \rightarrow \infty} S_T \leq \lim_{T \rightarrow \infty} R_T/T = 0$, which implies that convergence to the global maximum is guaranteed.

Gaussian Process (GP). In order to facilitate the design of the AF to minimize the regret, we model our belief of the unknown objective function f using a GP. Let $f(x)_{x \in \mathcal{X}}$ denote a GP, that is, every finite subset of $f(x)_{x \in \mathcal{X}}$ follows a multivariate Gaussian distribution (Rasmussen & Williams, 2006). Then, the GP is fully specified by its *prior* mean $\mu_x \triangleq \mathbb{E}[f(x)]$ and covariance $k_{xx'} \triangleq \operatorname{cov}[f(x), f(x')]$ for all $x, x' \in \mathcal{X}$. We assume that $k_{xx'}$ is defined by the commonly-used isotropic⁴ squared exponential covariance function $k_{xx'} \triangleq \sigma_y^2 \exp\{-0.5\|x - x'\|^2/l^2\}$, in which σ_y^2 is the signal variance controlling the intensity of output measurements and l is the length-scale controlling the correlation or “similarity” between output measurements. Furthermore, without loss of generality, we assume $\mu_x = 0$ and $k_{xx'} \leq 1$ for all $x, x' \in \mathcal{X}$. Given a column vector $\mathbf{y}_t \triangleq [y_i]_{i=1, \dots, t}^\top$ of noisy output measurements for some set

⁴ The non-isotropic squared exponential covariance function for $x, x' \in \mathcal{X}$ is defined as $k_{xx'} \triangleq \sigma_y^2 \exp\{(x - x')^\top \Gamma^{-2} (x - x')\}$, in which Γ is a diagonal matrix with length-scale components $[l_1, \dots, l_d]$. It can be easily transformed to an isotropic one by preprocessing the inputs, i.e., dividing each dimension of inputs x, x' by the respective length-scale component.

$\mathcal{X}_t \triangleq \{x_1, \dots, x_t\}$ of inputs after t iterations, the *posterior* predictive distribution of $f(x)$ at any input x is a Gaussian distribution with the following posterior mean and variance:

$$\begin{aligned} \mu_{t+1}(x) &\triangleq K_{x\mathcal{X}_t}(K_{\mathcal{X}_t\mathcal{X}_t} + \sigma_n^2 I)^{-1} \mathbf{y}_t \\ \sigma_{t+1}^2(x) &\triangleq k_{xx} - K_{x\mathcal{X}_t}(K_{\mathcal{X}_t\mathcal{X}_t} + \sigma_n^2 I)^{-1} K_{\mathcal{X}_t x}, \end{aligned} \quad (1)$$

in which $K_{x\mathcal{X}_t} \triangleq (k_{xx'})_{x' \in \mathcal{X}_t}$ is a row vector, $K_{\mathcal{X}_t x} \triangleq K_{x\mathcal{X}_t}^\top$, and $K_{\mathcal{X}_t\mathcal{X}_t} \triangleq (k_{x'x''})_{x', x'' \in \mathcal{X}_t}$.

Under the privacy-preserving transformation (Fig. 1), denote the image of the set \mathcal{X}_t as $\mathcal{Z}_t \triangleq \{z_1, \dots, z_t\}$, and the images of the original inputs x and x' as z and z' respectively. Then, the covariance function $k_{zz'}$ can be defined similarly as $k_{xx'}$, and thus we can define an analogue of the predictive distribution (1) for z and \mathcal{Z}_t (instead of x and \mathcal{X}_t), which we denote as $\tilde{\mu}_{t+1}(z)$ and $\tilde{\sigma}_{t+1}^2(z)$. We assume that the function f is sampled from a GP defined over the original domain \mathcal{X} with the covariance function $k_{xx'}$ and with known hyperparameters (σ_y^2 and l), and we use the same hyperparameters for the covariance function $k_{zz'}$.

The GP-UCB Algorithm. The AF adopted by the GP-UCB algorithm (Srinivas et al., 2010) is the *upper confidence bound* (UCB) of f induced by the posterior GP model. In each iteration t , an input $x_t \in \mathcal{X}$ is selected to query by trading off between (a) sampling close to an expected maximum (i.e., with large posterior mean $\mu_t(x_t)$) given the current GP belief (i.e., exploitation) vs. (b) sampling an input with high predictive uncertainty (i.e., with large posterior standard deviation $\sigma_t(x_t)$) to improve the GP belief of f over \mathcal{X} (i.e., exploration). Specifically, $x_t \triangleq \operatorname{argmax}_{x \in \mathcal{X}} \mu_t(x) + \beta_t^{1/2} \sigma_t(x)$, in which the parameter $\beta_t > 0$ is set to trade off between exploitation vs. exploration. A remarkable property of the GP-UCB algorithm shown by the work of Srinivas et al. (2010) is that it achieves *no regret* asymptotically if the parameters $\beta_t > 0$ are chosen properly.

A recent work by Kusner et al. (2015) proposed a DP variant of GP-UCB for hyperparameter tuning of ML models. However, as mentioned in Section 1, this approach implies that the modeler and curator are represented by the same entity and thus both parties have full access to the sensitive dataset and detailed knowledge of the BO algorithm. In our outsourced setting, in contrast, the modeler only has access to the transformed privatized dataset, while the curator is unaware of the details of the BO algorithm, as described in our motivating scenarios (Section 1).

3. Outsourced Bayesian Optimization

In our PO-GP-UCB algorithm, the curator needs to perform a privacy-preserving transformation of the original dataset $\mathcal{X} \subset \mathbb{R}^d$ and release the transformed dataset $\mathcal{Z} \subset \mathbb{R}^r$ to the modeler. Subsequently, the modeler runs BO (i.e., GP-UCB) using \mathcal{Z} . When performing the transformation, the

goal of the curator is two-fold: Firstly, the transformation has to be differentially private with given DP parameters ϵ, δ (Definition 2); secondly, the transformation should allow the modeler to obtain good BO performance on the transformed dataset (in a sense to be formalized later in this section).

3.1. Transformation via Random Projection

Good BO performance by the modeler (i.e., the second goal of the curator) can be achieved by making the GP predictions (1) (on which the performance of the BO algorithm depends) using the transformed dataset \mathcal{Z} close to those using the original dataset \mathcal{X} . To this end, we ensure that the distances between all pairs of inputs are approximately preserved after the transformation. This is motivated by the fact that the GP predictions (1) (hence the BO performance) depend on the inputs only through the value of covariance, which, in the case of isotropic covariance functions⁴, only depends on the pairwise distances between inputs. Consequently, by preserving the pairwise distances between inputs, the performance of the BO (GP-UCB) algorithm run by the modeler on \mathcal{Z} is made similar to that of the non-private GP-UCB algorithm run on the original dataset \mathcal{X} , for which theoretical convergence guarantee has been shown (Srinivas et al., 2010). As a result, the BO performance in the outsourced setting can be theoretically guaranteed (Section 3.3) and thus practically assured.

Therefore, to achieve both goals of the curator, we need to address the question as to *what transformation preserves both the pairwise distances between inputs and DP*. A natural approach is to add noise directly to the matrix of pairwise distances between the original inputs from \mathcal{X} using standard DP methods such as the Laplace or Gaussian mechanisms (Dwork & Roth, 2014). However, the resulting noisy distance matrix is not guaranteed to produce an invertible covariance matrix $K_{\mathcal{X}_t\mathcal{X}_t} + \sigma_n^2 I$, which is a requirement for the GP predictions (1). Instead, we perform the transformation through a technique based on random projection, which satisfies both goals of the curator. Firstly, random projection through random samples from standard normal distribution has been shown to preserve DP (Blocki et al., 2012). Secondly, as a result of the Johnson-Lindenstrauss lemma (Johnson & Lindenstrauss, 1984), random projection is also able to approximately preserve the pairwise distances between inputs, as shown in the following lemma:

Lemma 1. *Let $\nu \in (0, 1/2)$, $\mu \in (0, 1)$, $d \in \mathbb{N}$ and a set $\mathcal{X} \subset \mathbb{R}^d$ of n row vectors be given. Let $r \in \mathbb{N}$ and M be a $d \times r$ matrix whose entries are i.i.d. samples from $\mathcal{N}(0, 1)$. If $r \geq 8 \log(n^2/\mu)/\nu^2$, the probability of*

$$(1 - \nu) \|x - x'\|^2 \leq r^{-1} \|xM - x'M\|^2 \leq (1 + \nu) \|x - x'\|^2$$

for all $x, x' \in \mathcal{X}$ is at least $1 - \mu$.

Remark 1. r controls the dimension of the random projection, while ν and μ control the accuracy. Lemma 1 corrob-

orates the intuition that a smaller value of r leads to larger values of ν and μ , i.e., lower random projection accuracy.

The proof (Appendix B.1) consists of a union bound applied to the Johnson-Lindenstrauss lemma (Johnson & Lindenstrauss, 1984), which is a result from geometry stating that a set of points in a high-dimensional space can be embedded into a lower-dimensional space such that the pairwise distances between the points are nearly preserved.

3.2. The Curator Part

The curator part (Algorithm 1) of our PO-GP-UCB algorithm takes as input the original dataset \mathcal{X} viewed as an $n \times d$ matrix³, the DP parameters ϵ, δ (Definition 2) and the random projection parameter r (Lemma 1)⁵. To begin with, the curator subtracts the mean from each column of \mathcal{X} (line 2), and then picks a matrix M of samples from standard normal distribution $\mathcal{N}(0, 1)$ to perform random projection (line 3). Next, if the smallest singular value $\sigma_{\min}(\mathcal{X})$ of the centered dataset \mathcal{X} is not less than a threshold ω (calculated in line 5), the curator outputs the random projection $\mathcal{Z} \triangleq r^{-1/2}\mathcal{X}M$ of the centered dataset \mathcal{X} (line 7). Otherwise, the curator increases the singular values of the centered dataset \mathcal{X} (line 9) to obtain a new dataset $\tilde{\mathcal{X}}$ and outputs the random projection $\mathcal{Z} \triangleq r^{-1/2}\tilde{\mathcal{X}}M$ of the new dataset $\tilde{\mathcal{X}}$ (line 10). Lastly, the curator releases \mathcal{Z} to the modeler (line 11).

Algorithm 1 PO-GP-UCB (The curator part)

- 1: **Input:** $\mathcal{X}, \epsilon, \delta, r$
 - 2: $\mathcal{X} \leftarrow \mathcal{X} - \mathbf{1}\mathbf{1}^\top \mathcal{X}/n$ where $\mathbf{1}$ is a $n \times 1$ vector of 1's
 - 3: Pick a $d \times r$ matrix M of i.i.d. samples from $\mathcal{N}(0, 1)$
 - 4: Compute the SVD of $\mathcal{X} = U\Sigma V^\top$
 - 5: $\omega \leftarrow 16\sqrt{r} \log(2/\delta)\epsilon^{-1} \log(16r/\delta)$
 - 6: **if** $\sigma_{\min}(\mathcal{X}) \geq \omega$ **then**
 - 7: **return** $\mathcal{Z} \leftarrow r^{-1/2}\mathcal{X}M$
 - 8: **else**
 - 9: $\tilde{\mathcal{X}} \leftarrow U\sqrt{\Sigma^2 + \omega^2 I_{n \times d}}V^\top$ where Σ^2 ($I_{n \times d}$) is an $n \times d$ matrix whose main diagonal has squared singular values of \mathcal{X} (ones) in each coordinate and all other coordinates are 0
 - 10: **return** $\mathcal{Z} \leftarrow r^{-1/2}\tilde{\mathcal{X}}M$
 - 11: **end if**
 - 12: Release dataset \mathcal{Z} to the modeler
-

The fact that Algorithm 1 both preserves DP and approximately preserves the pairwise distances between inputs is stated in Theorems 1 and 2 below.

Theorem 1. *Algorithm 1 preserves (ϵ, δ) -DP.*

In the proof of Theorem 1 (Appendix B.2), all singular

⁵ Note that in Theorem 3, the parameter r is calculated based on specific values of the parameters μ and ν (Lemma 1) in order to achieve the performance guarantee. However, in practice, μ and ν are not required to specify the value of r for Algorithm 1.

values of the dataset \mathcal{X} are required to be not less than ω (calculated in line 5). This explains the necessity of line 9, where we increase the singular values of the dataset \mathcal{X} if $\sigma_{\min}(\mathcal{X}) < \omega$, to ensure that this requirement is satisfied.

Theorem 2. *Let a dataset $\mathcal{X} \subset \mathbb{R}^d$ be given. Let $\nu \in (0, 1/2)$, $\mu \in (0, 1)$ be given. Let $r \in \mathbb{N}$, such that $r \geq 8 \log(n^2/\mu)/\nu^2$. Then, the probability of*

$$(1 - \nu)\|x - x'\|^2 \leq \|z - z'\|^2 \leq (1 + \nu)C'\|x - x'\|^2$$

for all $x, x' \in \mathcal{X}$ and their images $z, z' \in \mathcal{Z}$ is at least $1 - \mu$, in which $C' \triangleq 1 + \mathbb{1}_{\sigma_{\min}(\mathcal{X}) < \omega} \omega^2 / \sigma_{\min}^2(\mathcal{X})$.

The proof (Appendix B.3) consists of bounding the change in distances between inputs due to the increase of the singular values of the dataset \mathcal{X} (line 9 of Algorithm 1) and applying Lemma 1. It can be observed from Theorem 2 that when $\sigma_{\min}(\mathcal{X}) \geq \omega$, $C' = 1$, hence Algorithm 1 approximately preserves the pairwise distances between inputs.

There are several important differences between our Algorithm 1 and the work of Blocki et al. (2012). Firstly, Algorithm 3 of Blocki et al. (2012) releases a DP estimate of the dataset covariance matrix, while our Algorithm 1 outputs a DP transformation of the original dataset. Secondly, Algorithm 3 of Blocki et al. (2012) does not have the “if/else” condition (line 6 of Algorithm 1) and always increases the singular values as in line 9 of Algorithm 1. In our case, however, if the singular values are increased due to the condition $\sigma_{\min}(\mathcal{X}) < \omega$ (i.e., the “else” clause, line 8 of Algorithm 1), the pairwise input distances of the dataset \mathcal{X} are no longer approximately preserved in \mathcal{Z} (Theorem 2), which results in a slightly different regret bound (see Theorem 3 and Remark 2 below). This requires us to introduce the “if/else” condition in Algorithm 1. We discuss these changes in greater detail in Appendix B.2.

3.3. The Modeler Part

The modeler part of our PO-GP-UCB algorithm (Algorithm 2) takes as input the transformed dataset $\mathcal{Z} \subset \mathbb{R}^r$ received from the curator as well as the GP-UCB parameter δ' , and runs the GP-UCB algorithm for T iterations on \mathcal{Z} . In each iteration t , the modeler selects the candidate transformed input z_t by maximizing the GP-UCB AF (line 4), and queries the curator for the corresponding noisy output measurement y_t (line 5). To perform such a query, the modeler can send the index (row) i_t of the selected transformed input z_t in the dataset \mathcal{Z} viewed as a matrix³ to the curator. The curator can then find the preimage x_t of z_t by looking into the same row i_t of the dataset \mathcal{X} viewed as a matrix³. After identifying x_t , the curator can compute $f(x_t)$ to yield a noisy output measurement $y_t \triangleq f(x_t) + \epsilon_{GP}$ and send it to the modeler. The modeler then updates the GP posterior belief (line 6) and proceeds to the next iteration $t + 1$.

In our theoretical analysis, we make the assumption of the *diagonal dominance* property of the covariance matrices, which was used by previous works on GP with DP (Smith et al., 2018) and active learning (Hoang et al., 2014b):

Definition 3. Let a dataset $\mathcal{X} \subset \mathbb{R}^d$ and a set $\mathcal{X}_0 \subseteq \mathcal{X}$ be given. The covariance matrix $K_{\mathcal{X}_0, \mathcal{X}_0}$ is said to be *diagonally dominant* if for any $x \in \mathcal{X}_0$

$$k_{xx} \geq (\sqrt{|\mathcal{X}_0| - 1} + 1) \sum_{x' \in \mathcal{X}_0 \setminus x} k_{xx'}.$$

Note that this assumption is adopted mainly for the theoretical analysis, and is thus not strictly required in order for our algorithm to deliver competitive practical performance (Section 4). Theorem 3 below presents the theoretical guarantee on the BO performance of our PO-GP-UCB algorithm run by the modeler (Algorithm 2).

Algorithm 2 PO-GP-UCB (The modeler part)

- 1: **Input:** \mathcal{Z}, δ', T
 - 2: **for** $t = 1, \dots, T$ **do**
 - 3: Set $\beta_t \leftarrow 2 \log(nt^2\pi^2/6\delta')$
 - 4: $z_t \leftarrow \operatorname{argmax}_{z \in \mathcal{Z}} \tilde{\mu}_t(z) + \beta_t^{1/2} \tilde{\sigma}_t(z)$
 - 5: Query the curator for y_t
 - 6: Update GP posterior belief: $\tilde{\mu}_{t+1}(z)$ and $\tilde{\sigma}_{t+1}(z)$
 - 7: **end for**
-

Theorem 3. Let $\varepsilon_{ucb} > 0$, $\delta_{ucb} \in (0, 1)$, $T \in \mathbb{N}$, DP parameters ϵ and δ , and a dataset $\mathcal{X} \subset \mathbb{R}^d$ be given. Let $d \triangleq \operatorname{diam}(\mathcal{X})/l$ where $\operatorname{diam}(\mathcal{X})$ is the diameter of \mathcal{X} and l is the GP length-scale. Suppose for all $t = 1, \dots, T$, $|y_t| \leq L$ and $K_{\mathcal{X}_{t-1}, \mathcal{X}_{t-1}}$ is diagonally dominant. Suppose $r \geq 8 \log(n^2/\mu)/\nu^2$ (Algorithm 1) where $\mu \triangleq \delta_{ucb}/2$ and $\nu \triangleq \min(\varepsilon_{ucb}/(2\sqrt{3}d^2L), 2/d^2, 1/2)$, and $\delta' \triangleq \delta_{ucb}/2$ (Algorithm 2). If $\sigma_{\min}(\mathcal{X}) \geq \omega$, then the simple regret S_T incurred by Algorithm 2 run by the modeler satisfies

$$S_T \leq (\varepsilon_{ucb}^2 + 24(C_2 + C_1\beta_T^{1/2})^2 \log T/T + 24/\log(1 + \sigma_n^{-2}) \cdot \beta_T \gamma_T/T)^{1/2}$$

with probability at least $1 - \delta_{ucb}$, in which γ_T is the maximum information gain⁶ on the function f from any set of noisy output measurements of size T , $C_1 \triangleq \mathcal{O}(\sigma_y \sqrt{\sigma_y^2 + \sigma_n^2}(\sigma_y^2/\sigma_n^2 + 1))$ and $C_2 \triangleq \mathcal{O}(\sigma_y^2/\sigma_n^2 \cdot L)$.

The key idea of the proof (Appendix B.5) is to ensure that every value of the GP-UCB AF computed on the transformed dataset \mathcal{Z} is close to the value of the corresponding GP-UCB AF computed on the original dataset \mathcal{X} . Consequently, the regret of the PO-GP-UCB algorithm run on \mathcal{Z} can be analyzed using similar techniques as those adopted in the

⁶Srinivas et al. (2010) has shown that $\gamma_T = \mathcal{O}((\log T)^{d+1})$ for the squared exponential kernel.

analysis of the non-private GP-UCB algorithm run on the original dataset \mathcal{X} (Srinivas et al., 2010), which leads to the regret bound shown in Theorem 3.

Remark 2. If $\sigma_{\min}(\mathcal{X}) < \omega$, a similar upper bound on the regret can be proved with the difference that ε_{ucb} specified by the curator is replaced by a different constant, which, unlike ε_{ucb} , cannot be set arbitrarily. This results from the fact that if $\sigma_{\min}(\mathcal{X}) < \omega$, Algorithm 1 increases the singular values of the dataset \mathcal{X} (see line 9). As a consequence, the pairwise distances between inputs are no longer approximately preserved after the transformation (see Theorem 2), resulting in a looser regret bound (see Remark 6 in Appendix B.5).

Remark 3. The presence of the constant ε_{ucb} makes the regret upper bound of PO-GP-UCB slightly different from that of the original GP-UCB algorithm. ε_{ucb} can be viewed as controlling the trade-off between utility (BO performance) and privacy preservation (see more detailed discussion in Section 3.4). In contrast, the only prior works on privacy-preserving BO by Kusner et al. (2015) and Nguyen et al. (2018) do not provide any regret bounds.

Remark 4. The upper bound on the simple regret S_T in Theorem 3 indirectly depends on the DP parameter ϵ : the bound holds when $\sigma_{\min}(\mathcal{X}) \geq \omega$, in which ω depends on ϵ (line 5 of Algorithm 1). Moreover, when $\sigma_{\min}(\mathcal{X}) < \omega$, ε_{ucb} (which appears in the regret bound) is replaced by a different constant, which depends on ϵ (see Remark 2).

3.4. Analysis and Discussion

Interestingly, our theoretical results are amenable to elegant interpretations regarding the privacy-utility trade-off.

The flexibility to tune the value of ω to satisfy the condition required by Theorem 3 (i.e., $\sigma_{\min}(\mathcal{X}) \geq \omega$) incurs an interesting trade-off. Specifically, if $\sigma_{\min}(\mathcal{X}) < \omega$, we can either (a) run PO-GP-UCB without modifying any parameter, or (b) reduce ω by tuning the algorithmic parameters to satisfy the condition $\sigma_{\min}(\mathcal{X}) \geq \omega$, both of which incur some costs. In case (a), the resulting regret bound is looser as explained in Remark 2, which might imply worse BO performance. In case (b), to reduce the value of ω , we can either (i) increase the DP parameters ϵ and δ which deteriorates the DP guarantee, or (ii) decrease the value of r . A smaller value of r implies larger values of μ and ν as required by Theorem 3 ($r \geq 8 \log(n^2/\mu)/\nu^2$) and thus larger values of ε_{ucb} and δ_{ucb} as seen in the definitions of μ and ν in Theorem 3. This consequently results in a worse regret upper bound (Theorem 3) and thus deteriorated BO performance. Therefore, the privacy-utility trade-off is involved in our strategy to deal with the scenario where $\sigma_{\min}(\mathcal{X}) < \omega$.

For a fixed value of ω such that $\sigma_{\min}(\mathcal{X}) \geq \omega$, the privacy-utility trade-off can also be identified and thus utilized to adjust the the algorithmic parameters: $\epsilon, \delta, \varepsilon_{ucb}$ and δ_{ucb} .




Figure 2. Simple regrets achieved by tested BO algorithms (with fixed r and different values of ϵ) vs. the number of iterations for (a) the synthetic GP dataset ($r = 10$), (b) loan applications dataset ($r = 15$), and (c) private property price dataset ($r = 15$).

Specifically, decreasing the values of the DP parameters ϵ and δ improves the privacy guarantee. However, in order to fix the value of ω (to ensure that the condition $\sigma_{\min}(\mathcal{X}) \geq \omega$ remains satisfied), the value of r needs to be reduced, which results in larger values of ϵ_{ucb} and δ_{ucb} and thus worse BO performance (as discussed in the previous paragraph). Similar analysis reveals that decreasing the values of ϵ_{ucb} and δ_{ucb} improves the BO performance, at the expense of looser privacy guarantee (i.e., larger required values of ϵ and δ). Furthermore, the role played by ω in Algorithm 1 provides a guideline on the practical design of the algorithm. In particular, for a fixed desirable level of privacy (i.e., fixed values of ϵ and δ), the value of r should be made as large as possible while still ensuring that the condition $\sigma_{\min}(\mathcal{X}) \geq \omega$ is satisfied, since larger r improves the BO performance until this condition is violated. This guideline will be exploited and validated in the experiments.

These insights regarding the privacy-utility trade-off serve as intuitive justifications of our PO-GP-UCB algorithm and provide useful guidelines for its practical deployment.

4. Experiments and Discussion

In this section, we empirically evaluate the performance of our PO-GP-UCB algorithm using four datasets including a synthetic GP dataset, a real-world loan applications dataset, a real-world property price dataset and, in Appendix A, the Branin-Hoo benchmark function. The performances (simple regrets) of our algorithm are compared with that of the non-private GP-UCB algorithm run using the original datasets (Srinivas et al., 2010). The original output measurements for both real-world datasets are log-transformed to remove skewness and extremity in order to stabilize the GP covariance structure. The GP hyperparameters are learned using maximum likelihood estimation (Rasmussen & Williams, 2006). All results are averaged over 50 random runs, each of which uses a different set of initializations for BO. Each random run uses an independent realization of the matrix M of i.i.d. samples from $\mathcal{N}(0, 1)$ for performing random projection (line 3 of Algorithm 1). We set the GP-

UCB parameter $\delta_{ucb} = 0.05$ (Theorem 3) and normalize the inputs to have a maximal norm of 25 in all experiments. Following the guidelines by the state-of-the-art works in DP (Dwork & Roth, 2014; Abadi et al., 2016; Foulds et al., 2016; Papernot et al., 2017), we fix the value of the DP parameter δ (Definition 2) to be smaller than $1/n$ in all experiments. Note that setting the values of the parameters μ , ν (Lemma 1) and the GP-UCB parameter ϵ_{ucb} (Theorem 3), as well as assuming the diagonal dominance of covariance matrices (Definition 3), is required only for our theoretical analysis and thus not necessary in the practical employment of our algorithm.

In every experiment that varies the value of the DP parameter ϵ (Definition 2) (Fig. 2), the PO-GP-UCB algorithm with the largest value of ϵ under consideration⁷ satisfies the condition $\sigma_{\min}(\mathcal{X}) \geq \omega$ (i.e., the “if” clause, line 6 of Algorithm 1), while the algorithms with all other values of ϵ under consideration satisfy the condition $\sigma_{\min}(\mathcal{X}) < \omega$ (i.e., the “else” clause, line 8 of Algorithm 1).

Synthetic GP dataset. The original inputs for this experiment are 2-dimensional vectors arranged into a uniform grid and discretized into a 100×100 input domain (i.e., $d = 2$ and $n = 10000$). The function to maximize is sampled from a GP with the GP hyperparameters $\mu_x = 0$, $l = 1.25$, $\sigma_y^2 = 1$ and $\sigma_n^2 = 10^{-5}$. We set the parameter $r = 10$ (Algorithm 1), DP parameter $\delta = 10^{-5}$ (Definition 2) and the GP-UCB parameter $T = 50$ for this experiment.

Fig. 2a shows the performances of PO-GP-UCB with different values of ϵ and that of non-private GP-UCB. It can be observed that smaller values of ϵ (tighter privacy guarantees) result in larger simple regret, which is consistent with the privacy-utility trade off. PO-GP-UCB with the largest value of $\epsilon = \exp(1.1)$ satisfying the condition $\sigma_{\min}(\mathcal{X}) \geq \omega$ achieves only $0.011\sigma_y$ more simple regret than non-private GP-UCB after 50 iterations. Interestingly, despite having

⁷Further increasing the value of ϵ will only decrease the value of ω (see line 5 of Algorithm 1), so the condition $\sigma_{\min}(\mathcal{X}) \geq \omega$ will remain satisfied. As a result, the dataset \mathcal{Z} returned by Algorithm 1 and hence the performance of PO-GP-UCB will stay the same.

a looser regret bound (see Remark 2), the PO-GP-UCB algorithm with some smaller values of ϵ satisfying the condition $\sigma_{\min}(\mathcal{X}) < \omega$ also only incurs slightly larger regret than non-private GP-UCB. In particular, PO-GP-UCB with $\epsilon = \exp(0.9)$ ($\epsilon = \exp(0.0)$) achieves only $0.069\sigma_y$ ($0.099\sigma_y$) more simple regret after 50 iterations. Therefore, our algorithm is able to achieve favorable performance with the values of ϵ in the single-digit range, which is consistent with the practice of the state-of-the-art works on the application of DP in ML (Abadi et al., 2016; Foulds et al., 2016; Papernot et al., 2017). This implies our algorithm’s practical capability of simultaneously achieving tight privacy guarantee and obtaining competitive BO performance.

We also investigate the impact of varying the value of the random projection parameter r on the performance of PO-GP-UCB. In particular, we consider 3 different values of DP parameter ϵ : $\epsilon = \exp(1.1)$, $\epsilon = \exp(1.3)$ and $\epsilon = \exp(1.5)$. We then fix the value of ϵ and vary the value of r . The largest value of r satisfying the condition $\sigma_{\min}(\mathcal{X}) \geq \omega$ is $r = 10$ for $\epsilon = \exp(1.1)$, $r = 15$ for $\epsilon = \exp(1.3)$ and $r = 20$ for $\epsilon = \exp(1.5)$. Tables 1, 2 and 3 reveal that the largest values of r satisfying the condition $\sigma_{\min}(\mathcal{X}) \geq \omega$ lead to the smallest simple regret after 50 iterations. Decreasing the value of r increases the simple regret, which agrees with our analysis in Section 3.4 (i.e., smaller r results in worse regret upper bound). On the other hand, increasing r such that the condition $\sigma_{\min}(\mathcal{X}) < \omega$ is satisfied also results in larger simple regret, which is again consistent with the analysis in Remark 2 stating that the regret upper bound becomes looser in this scenario. This experiment suggests that, in practice, for a fixed desirable privacy level (i.e., if the values of the DP parameters ϵ and δ are fixed), r should be chosen as the largest value satisfying the condition $\sigma_{\min}(\mathcal{X}) \geq \omega$.

Table 1. Simple regrets achieved by PO-GP-UCB with fixed $\epsilon = \exp(1.1)$ and different values of r after 50 iterations for the synthetic GP dataset. The largest value of r satisfying the condition $\sigma_{\min}(\mathcal{X}) \geq \omega$ is $r = 10$.

r	3	6	8	10	15	20
S_{50}	0.073	0.038	0.018	0.014	0.118	0.137

Table 2. Simple regrets achieved by PO-GP-UCB with fixed $\epsilon = \exp(1.3)$ and different values of r after 50 iterations for the synthetic GP dataset. The largest value of r satisfying the condition $\sigma_{\min}(\mathcal{X}) \geq \omega$ is $r = 15$.

r	3	9	12	15	20	30
S_{50}	0.091	0.009	0.019	0.008	0.127	0.134

Real-world loan applications dataset. A bank is selecting the loan applicants with the highest return on investment (ROI) and outsources the task to a financial AI consultancy. The inputs to BO are the data of 36000 loan applicants (we use the public data from <https://www.lendingclub.com/>),

Table 3. Simple regrets achieved by PO-GP-UCB with fixed $\epsilon = \exp(1.5)$ and different values of r after 50 iterations for the synthetic GP dataset. The largest value of r satisfying the condition $\sigma_{\min}(\mathcal{X}) \geq \omega$ is $r = 20$.

r	5	10	15	20	30	50
S_{50}	0.05	0.021	0.003	0.002	0.094	0.142

each consisting of three features: the total amount committed by investors for the loan, the interest rate on the loan and the annual income of the applicant (i.e., $n = 36000$ and $d = 3$). The function to maximize (the output measurement) is the ROI for an applicant. The original inputs are preprocessed to form an isotropic covariance function⁴. We set $r = 15$, $\delta = 10^{-5}$ and $T = 50$.

Fig. 2b presents the results of varying the value of ϵ . Similar to the synthetic GP dataset, after 50 iterations, the simple regret achieved by PO-GP-UCB with the largest value of $\epsilon = \exp(2.9)$ satisfying the condition $\sigma_{\min}(\mathcal{X}) \geq \omega$ is slightly larger (by $0.003\sigma_y$) than that achieved by non-private GP-UCB. Moreover, PO-GP-UCB with some values of ϵ in the single-digit range satisfying the condition $\sigma_{\min}(\mathcal{X}) < \omega$ shows marginally worse performance compared with non-private GP-UCB. In particular, after 50 iterations, $\epsilon = \exp(2.0)$ and $\epsilon = \exp(1.0)$ result in $0.019\sigma_y$ and $0.05\sigma_y$ more simple regret than non-private GP-UCB respectively.

We examine the effect of r on the performance of PO-GP-UCB, by fixing the value of DP parameter ϵ and changing r . We consider 3 different values of DP parameter ϵ : $\epsilon = \exp(2.7)$, $\epsilon = \exp(2.9)$ and $\epsilon = \exp(3.1)$. The largest value of r satisfying the condition $\sigma_{\min}(\mathcal{X}) \geq \omega$ is $r = 10$ for $\epsilon = \exp(2.7)$, $r = 15$ for $\epsilon = \exp(2.9)$ and $r = 20$ for $\epsilon = \exp(3.1)$. The results are presented in Tables 4, 5 and 6. PO-GP-UCB with the largest r satisfying the condition $\sigma_{\min}(\mathcal{X}) \geq \omega$ in general leads to the best performance, i.e., it achieves the smallest simple regret in Tables 4 and 5, and the second smallest simple regret in Table 6. Similar insights to the results of the synthetic GP dataset can also be drawn: reducing the value of r and increasing the value of r to satisfy the condition $\sigma_{\min}(\mathcal{X}) < \omega$ both result in larger simple regret, which again corroborates our theoretical analysis.

Table 4. Simple regrets achieved by PO-GP-UCB with fixed $\epsilon = \exp(2.7)$ and different values of r after 50 iterations for the real-world loan applications dataset. The largest value of r satisfying the condition $\sigma_{\min}(\mathcal{X}) \geq \omega$ is $r = 10$.

r	3	6	8	10	15	20
S_{50}	0.083	0.088	0.078	0.069	0.081	0.076

Real-world private property price dataset. A real estate agency is trying to locate the cheapest private properties and outsources the task of selecting the candidate properties to an AI consultancy. The original inputs are the

Table 5. Simple regrets achieved by PO-GP-UCB with fixed $\epsilon = \exp(2.9)$ and different values of r after 50 iterations for the real-world loan applications dataset. The largest value of r satisfying the condition $\sigma_{\min}(\mathcal{X}) \geq \omega$ is $r = 15$.

r	3	9	12	15	20	30
S_{50}	0.091	0.076	0.078	0.077	0.1	0.096

Table 6. Simple regrets achieved by PO-GP-UCB with fixed $\epsilon = \exp(3.1)$ and different values of r after 50 iterations for the real-world loan applications dataset. The largest value of r satisfying the condition $\sigma_{\min}(\mathcal{X}) \geq \omega$ is $r = 20$.

r	5	10	15	20	30	50
S_{50}	0.097	0.091	0.069	0.084	0.104	0.127

longitude/latitude coordinates of 2004 individual properties (i.e., $n = 2004$ and $d = 2$). We use the public data from <https://www.ura.gov.sg/realEstateIIWeb/transaction/search.action>. The function to minimize is the evaluated property price measured in dollars per square meter. We set $r = 15$, $\delta = 10^{-4}$ and $T = 100$.

The results of this experiment for different values of ϵ are displayed in Fig. 2c. Similar observations can be made that are consistent with the previous experiments. In particular, smaller values of ϵ (tighter privacy guarantees) generally lead to worse BO performance (larger simple regret); PO-GP-UCB with the largest value of $\epsilon = \exp(2.8)$ satisfying the condition $\sigma_{\min}(\mathcal{X}) \geq \omega$ incurs slightly larger simple regret ($0.051\sigma_y$) than non-private GP-UCB after 100 iterations; PO-GP-UCB with some values of ϵ in the single-digit range satisfying the condition $\sigma_{\min}(\mathcal{X}) < \omega$ exhibits small disadvantages compared with non-private GP-UCB after 100 iterations in terms of simple regrets: $\epsilon = \exp(1.0)$ and $\epsilon = \exp(0.5)$ result in $0.017\sigma_y$ and $0.082\sigma_y$ more simple regret respectively.

We again empirically inspect the impact of r on the performance of PO-GP-UCB in the same manner as the previous experiments: we fix the value of ϵ and vary the value of r . We consider 3 different values of DP parameter ϵ : $\epsilon = \exp(2.6)$, $\epsilon = \exp(2.8)$ and $\epsilon = \exp(3.0)$. The largest value of r satisfying the condition $\sigma_{\min}(\mathcal{X}) \geq \omega$ is $r = 10$ for $\epsilon = \exp(2.6)$, $r = 15$ for $\epsilon = \exp(2.8)$ and $r = 20$ for $\epsilon = \exp(3.0)$. Tables 7, 8 and 9 show that the smallest simple regret is achieved by the largest values of r satisfying the condition $\sigma_{\min}(\mathcal{X}) \geq \omega$. Similar to the previous experiments, smaller values of r and larger values of r that satisfy the condition $\sigma_{\min}(\mathcal{X}) < \omega$ both lead to larger simple regret, further validating the practicality of our guideline on the selection of r (Section 3.4).

5. Conclusion and Future Work

This paper describes PO-GP-UCB, which is the first algorithm for BO with DP in the outsourced setting with

Table 7. Simple regrets achieved by PO-GP-UCB with fixed $\epsilon = \exp(2.6)$ and different values of r after 100 iterations for the real-world property price dataset. The largest value of r satisfying the condition $\sigma_{\min}(\mathcal{X}) \geq \omega$ is $r = 10$.

r	3	6	8	10	15	20
S_{100}	0.682	0.516	0.495	0.485	0.485	0.493

Table 8. Simple regrets achieved by PO-GP-UCB with fixed $\epsilon = \exp(2.8)$ and different values of r after 100 iterations for the real-world property price dataset. The largest value of r satisfying the condition $\sigma_{\min}(\mathcal{X}) \geq \omega$ is $r = 15$.

r	3	9	12	15	20	30
S_{100}	0.567	0.553	0.479	0.453	0.493	0.52

Table 9. Simple regrets achieved by PO-GP-UCB with fixed $\epsilon = \exp(3.0)$ and different values of r after 100 iterations for the real-world property price dataset. The largest value of r satisfying the condition $\sigma_{\min}(\mathcal{X}) \geq \omega$ is $r = 20$.

r	5	10	15	20	30	50
S_{100}	0.591	0.523	0.486	0.482	0.489	0.488

theoretical performance guarantee. We prove the privacy-preserving property of our algorithm and show a theoretical upper bound on the regret. We use both synthetic and real-world experiments to show the empirical effectiveness of our algorithm, as well as its ability to achieve state-of-the-art privacy guarantees (in the single-digit range) and handle the privacy-utility trade-off. For future work, it would be interesting to investigate whether PO-GP-UCB can be extended for privately releasing the output measurements y_t . To this end, the work of Hall et al. (2013) which provides a way for DP release of functional data can potentially be applied. Another direction would be to investigate whether the work of Kenthapadi et al. (2013) on DP random projection can be used as a privacy-preserving mechanism in our outsourced BO framework to improve the privacy guarantee. We will consider generalizing PO-GP-UCB to non-myopic BO (Kharkovskii et al., 2020; Ling et al., 2016), batch BO (Daxberger & Low, 2017), high-dimensional BO (Hoang et al., 2018), and multi-fidelity BO (Zhang et al., 2017; 2019) settings and incorporating early stopping (Dai et al., 2019) and recursive reasoning (Dai et al., 2020). We will also consider our outsourced setting in the active learning context (Cao et al., 2013; Hoang et al., 2014a;b; Low et al., 2008; 2009; 2011; 2012; 2014a; Ouyang et al., 2014; Zhang et al., 2016). For applications with a huge budget of function evaluations, we like to couple PO-GP-UCB with the use of distributed/decentralized (Chen et al., 2012; 2013a;b; 2015; Hoang et al., 2016; 2019b;a; Low et al., 2015; Ouyang & Low, 2018) or online/stochastic (Hoang et al., 2015; 2017; Low et al., 2014b; Xu et al., 2014; Teng et al., 2020; Yu et al., 2019a;b) sparse GP models to represent the belief of the unknown objective function efficiently.

Acknowledgements

This research/project is supported by the National Research Foundation, Singapore under its Strategic Capability Research Centres Funding Initiative and its AI Singapore Programme (Award Number: AISG-GC-2019-002). 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

- Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., and Zhang, L. Deep learning with differential privacy. In *Proc. ACM CCS*, pp. 308–318, 2016.
- Blocki, J., Blum, A., Datta, A., and Sheffet, O. The Johnson-Lindenstrauss transform itself preserves differential privacy. In *Proc. FOCS*, pp. 410–419, 2012.
- Cao, N., Low, K. H., and Dolan, J. M. Multi-robot informative path planning for active sensing of environmental phenomena: A tale of two algorithms. In *Proc. AAMAS*, pp. 7–14, 2013.
- Chen, J., Low, K. H., Tan, C. K.-Y., Oran, A., Jaillet, P., Dolan, J. M., and Sukhatme, G. S. Decentralized data fusion and active sensing with mobile sensors for modeling and predicting spatiotemporal traffic phenomena. In *Proc. UAI*, pp. 163–173, 2012.
- Chen, J., Cao, N., Low, K. H., Ouyang, R., Tan, C. K.-Y., and Jaillet, P. Parallel Gaussian process regression with low-rank covariance matrix approximations. In *Proc. UAI*, pp. 152–161, 2013a.
- Chen, J., Low, K. H., and Tan, C. K.-Y. Gaussian process-based decentralized data fusion and active sensing for mobility-on-demand system. In *Proc. RSS*, 2013b.
- Chen, J., Low, K. H., Jaillet, P., and Yao, Y. Gaussian process decentralized data fusion and active sensing for spatiotemporal traffic modeling and prediction in mobility-on-demand systems. *IEEE Trans. Autom. Sci. Eng.*, 12: 901–921, 2015.
- Chong, M., Abraham, A., and Paprzycki, M. Traffic accident analysis using machine learning paradigms. *Informatica (Slovenia)*, 29(1):89–98, 2005.
- Dai, Z., Yu, H., Low, K. H., and Jaillet, P. Bayesian optimization meets Bayesian optimal stopping. In *Proc. ICML*, pp. 1496–1506, 2019.
- Dai, Z., Chen, Y., Low, K. H., Jaillet, P., and Ho, T.-H. R2-B2: Recursive reasoning-based Bayesian optimization for no-regret learning in games. In *Proc. ICML*, 2020.
- Daxberger, E. A. and Low, K. H. Distributed batch Gaussian process optimization. In *Proc. ICML*, pp. 951–960, 2017.
- Dewancker, I., McCourt, M., Clark, S., Hayes, P., Johnson, A., and Ke, G. Evaluation system for a Bayesian optimization service. arXiv:1605.06170, 2016.
- Dwork, C. and Roth, A. The algorithmic foundations of differential privacy. *Foundations and Trends[®] in Theoretical Computer Science*, 9(3-4):211–407, 2014.
- Dwork, C., McSherry, F., Nissim, K., and Smith, A. D. Calibrating noise to sensitivity in private data analysis. *J. Priv. Confidentiality*, 7(3):17–51, 2016.
- Foulds, J., Geumlek, J., Welling, M., and Chaudhuri, K. On the theory and practice of privacy-preserving Bayesian data analysis. In *Proc. UAI*, pp. 192–201, 2016.
- Hall, R., Rinaldo, A., and Wasserman, L. Differential privacy for functions and functional data. *JMLR*, 14(1): 703–727, 2013.
- Hardt, M. and Roth, A. Beating randomized response on incoherent matrices. In *Proc. STOC*, pp. 1255–1268, 2012.
- Hoang, Q. M., Hoang, T. N., and Low, K. H. A generalized stochastic variational Bayesian hyperparameter learning framework for sparse spectrum Gaussian process regression. In *Proc. AAAI*, pp. 2007–2014, 2017.
- Hoang, Q. M., Hoang, T. N., Low, K. H., and Kingsford, C. Collective model fusion for multiple black-box experts. In *Proc. ICML*, pp. 2742–2750, 2019a.
- Hoang, T. N., Low, K. H., Jaillet, P., and Kankanhalli, M. Active learning is planning: Nonmyopic ϵ -Bayes-optimal active learning of Gaussian processes. In *Proc. ECML/PKDD Nectar Track*, pp. 494–498, 2014a.
- Hoang, T. N., Low, K. H., Jaillet, P., and Kankanhalli, M. Nonmyopic ϵ -Bayes-optimal active learning of Gaussian processes. In *Proc. ICML*, pp. 739–747, 2014b.
- Hoang, T. N., Hoang, Q. M., and Low, K. H. A unifying framework of anytime sparse Gaussian process regression models with stochastic variational inference for big data. In *Proc. ICML*, pp. 569–578, 2015.
- Hoang, T. N., Hoang, Q. M., and Low, K. H. A distributed variational inference framework for unifying parallel sparse Gaussian process regression models. In *Proc. ICML*, pp. 382–391, 2016.
- Hoang, T. N., Hoang, Q. M., and Low, K. H. Decentralized high-dimensional Bayesian optimization with factor graphs. In *Proc. AAAI*, pp. 3231–3238, 2018.

- Hoang, T. N., Hoang, Q. M., Low, K. H., and How, J. P. Collective online learning of Gaussian processes in massive multi-agent systems. In *Proc. AAAI*, pp. 7850–7857, 2019b.
- Johnson, W. B. and Lindenstrauss, J. Extensions of Lipschitz maps into a Hilbert space. *Contemporary Mathematics*, 26(2):189–206, 1984.
- Kenthapadi, K., Korolova, A., Mironov, I., and Mishra, N. Privacy via the Johnson-Lindenstrauss transform. *J. Priv. Confidentiality*, 5(1):39–71, 2013.
- Kharkovskii, D., Ling, C. K., and Low, K. H. Nonmyopic Gaussian process optimization with macro-actions. In *Proc. AISTATS*, pp. 4593–4604, 2020.
- Kusner, M. J., Gardner, J. R., Garnett, R., and Weinberger, K. Q. Differentially private Bayesian Optimization. In *Proc. ICML*, pp. 918–927, 2015.
- Ling, C. K., Low, K. H., and Jaillet, P. Gaussian process planning with Lipschitz continuous reward functions: Towards unifying Bayesian optimization, active learning, and beyond. In *Proc. AAAI*, pp. 1860–1866, 2016.
- Low, K. H., Dolan, J. M., and Khosla, P. Adaptive multi-robot wide-area exploration and mapping. In *Proc. AAMAS*, pp. 23–30, 2008.
- Low, K. H., Dolan, J. M., and Khosla, P. Information-theoretic approach to efficient adaptive path planning for mobile robotic environmental sensing. In *Proc. ICAPS*, pp. 233–240, 2009.
- Low, K. H., Dolan, J. M., and Khosla, P. Active Markov information-theoretic path planning for robotic environmental sensing. In *Proc. AAMAS*, pp. 753–760, 2011.
- Low, K. H., Chen, J., Dolan, J. M., Chien, S., and Thompson, D. R. Decentralized active robotic exploration and mapping for probabilistic field classification in environmental sensing. In *Proc. AAMAS*, pp. 105–112, 2012.
- Low, K. H., Chen, J., Hoang, T. N., Xu, N., and Jaillet, P. Recent advances in scaling up Gaussian process predictive models for large spatiotemporal data. In Ravela, S. and Sandu, A. (eds.), *Dynamic Data-Driven Environmental Systems Science: First International Conference, DyDESS 2014*, pp. 167–181. LNCS 8964, Springer International Publishing, 2014a.
- Low, K. H., Xu, N., Chen, J., Lim, K. K., and Özgül, E. B. Generalized online sparse Gaussian processes with application to persistent mobile robot localization. In *Proc. ECML/PKDD Nectar Track*, pp. 499–503, 2014b.
- Low, K. H., Yu, J., Chen, J., and Jaillet, P. Parallel Gaussian process regression for big data: Low-rank representation meets Markov approximation. In *Proc. AAAI*, pp. 2821–2827, 2015.
- Ngai, E. W. T., Hu, Y., Wong, Y. H., Chen, Y., and Sun, X. The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature. *Decis. Support Syst.*, 50(3): 559–569, 2011.
- Nguyen, T. D., Gupta, S., Rana, S., and Venkatesh, S. A privacy preserving Bayesian Optimization with high efficiency. In *Proc. PAKDD*, pp. 543–555, 2018.
- Ouyang, R. and Low, K. H. Gaussian process decentralized data fusion meets transfer learning in large-scale distributed cooperative perception. In *Proc. AAAI*, pp. 3876–3883, 2018.
- Ouyang, R., Low, K. H., Chen, J., and Jaillet, P. Multi-robot active sensing of non-stationary Gaussian process-based environmental phenomena. In *Proc. AAMAS*, pp. 573–580, 2014.
- Papernot, N., Abadi, M., Erlingsson, U., Goodfellow, I., and Talwar, K. Semi-supervised knowledge transfer for deep learning from private training data. In *Proc. ICLR*, 2017.
- Rasmussen, C. E. and Williams, C. K. I. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- Sarwate, A. D. and Chaudhuri, K. Signal processing and machine learning with differential privacy: Algorithms and challenges for continuous data. *IEEE Signal Processing Magazine*, 30(5):86–94, 2013.
- Shahriari, B., Swersky, K., Wang, Z., Adams, R., and de Freitas, N. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016.
- Smith, M., Álvarez, M., Zwiessle, M., and Lawrence, N. D. Differentially private regression with Gaussian processes. In *Proc. AISTATS*, pp. 1195–1203, 2018.
- Srinivas, N., Krause, A., Kakade, S., and Seeger, M. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proc. ICML*, pp. 1015–1022, 2010.
- Stewart, G. W. and Sun, J. *Matrix Perturbation Theory*. Academic Press, 1990.
- Teng, T., Chen, J., Zhang, Y., and Low, K. H. Scalable variational bayesian kernel selection for sparse Gaussian process regression. In *Proc. AAAI*, pp. 5997–6004, 2020.

- Xu, N., Low, K. H., Chen, J., Lim, K. K., and Özgül, E. B. GP-Localize: Persistent mobile robot localization using online sparse Gaussian process observation model. In *Proc. AAAI*, pp. 2585–2592, 2014.
- Yu, H., Chen, Y., Dai, Z., Low, K. H., and Jaillet, P. Implicit posterior variational inference for deep Gaussian processes. In *Proc. NeurIPS*, pp. 14475–14486, 2019a.
- Yu, H., Hoang, T. N., Low, K. H., and Jaillet, P. Stochastic variational inference for Bayesian sparse Gaussian process regression. In *Proc. IJCNN*, 2019b.
- Yu, S., van Esbroeck, A., Farooq, F., Fung, G., Anand, V., and Krishnapuram, B. Predicting readmission risk with institution specific prediction models. In *Proc. IEEE ICHI*, pp. 415–420, 2013.
- Zhang, Y., Hoang, T. N., Low, K. H., and Kankanhalli, M. Near-optimal active learning of multi-output Gaussian processes. In *Proc. AAAI*, pp. 2351–2357, 2016.
- Zhang, Y., Hoang, T. N., Low, K. H., and Kankanhalli, M. Information-based multi-fidelity Bayesian optimization. In *Proc. NIPS Workshop on Bayesian Optimization*, 2017.
- Zhang, Y., Dai, Z., and Low, K. H. Bayesian optimization with binary auxiliary information. In *Proc. UAI*, 2019.