
Bayesian Optimization under Stochastic Delayed Feedback

Arun Verma^{*1} Zhongxiang Dai^{*1} Bryan Kian Hsiang Low¹

Abstract

Bayesian optimization (BO) is a widely-used sequential method for zeroth-order optimization of complex and expensive-to-compute black-box functions. The existing BO methods assume that the function evaluation (feedback) is available to the learner immediately or after a fixed delay. Such assumptions may not be practical in many real-life problems like online recommendations, clinical trials, and hyperparameter tuning where feedback is available after a random delay. To benefit from the experimental parallelization in these problems, the learner needs to start new function evaluations without waiting for delayed feedback. In this paper, we consider the BO under stochastic delayed feedback problem. We propose algorithms with sub-linear regret guarantees that efficiently address the dilemma of selecting new function queries while waiting for randomly delayed feedback. Building on our results, we also make novel contributions to batch BO and contextual Gaussian process bandits. Experiments on synthetic and real-life datasets verify the performance of our algorithms.

1. Introduction

Bayesian optimization (BO) (Shahriari et al., 2015; Frazier, 2018; Garnett, 2022) is a popular and widely-used sequential method for zeroth-order optimization of unknown black-box functions that are complex and expensive to compute. The existing BO methods assume that the function evaluation (feedback) is available to the learner immediately or after a fixed delay. However, these assumptions are impractical in many real-life problems where feedback is available after a *random delay*. To take advantage of the experimental parallelization in these problems, the learner needs to start

new function evaluations without waiting for the randomly delayed feedback. We refer to this new BO problem as ‘*Bayesian Optimization under Stochastic Delayed Feedback*’ (BO-SDF). In this paper, we propose algorithms that efficiently address the problem of selecting new function queries while waiting for randomly delayed feedback. Specifically, we answer the following question:

How to start a new function query when the observations of the past function queries are randomly delayed?

Many real-life applications can be cast as BO-SDF problems. For example, when performing clinical trials to discover new medicines, we need to optimize the amount of different raw materials in order to find their most effective composition using a small number of clinical trials (Chow & Chang, 2006), which is a natural application for BO due to its sample efficiency. However, the evaluation of the effectiveness of a medicine needs to take into account the side effects which are usually not observed immediately but are instead revealed *after a random period* due to the varying physiology of different patients. Therefore, a natural challenge here is how to use BO to choose the next composition for testing, while accounting for the stochastically delayed observations from some previous clinical trials. A similar challenge also arises when we aim to find the most effective dose of a new medicine through clinical trials (Takahashi & Suzuki, 2021).

Another motivating example arises from online product recommendation (Chapelle, 2014; Diemert et al., 2017). Most online platforms make recommendations in milliseconds, but the user’s response (i.e., watching a movie or buying a product) generally happens after a random time period which may range from hours to days or even weeks. Furthermore, the issue of randomly delayed feedback also plagues other common applications of BO such as hyperparameter tuning of machine learning (ML) models (Snoek et al., 2012) (e.g., the training time of a neural network depends on the number of layers, network width, among others) and material discovery (Ueno et al., 2016).

The closest BO setting to our BO-SDF is batch BO (Desautels et al., 2014; Daxberger & Low, 2017; Chowdhury & Gopalan, 2019) which also allows multiple queries to be performed on the black-box function in parallel. Both problems require choosing an input query to evaluate

^{*}Equal contribution ¹Department of Computer Science, National University of Singapore, Republic of Singapore. Correspondence to: Arun Verma <arunverma100@gmail.com>, Zhongxiang Dai <daiz9109@gmail.com>.

the black-box function while the observations of some previously selected inputs are not available (i.e., delayed). However, there are important differences because the delays are fixed in a certain way in batch BO but can be random in BO-SDF. Interestingly, batch BO can be considered as a special case of our BO-SDF problem and hence, our algorithms proposed for BO-SDF can be applied to batch BO while achieving important improvements (Section 4).

Following the practice of BO, in order to choose the input queries to quickly approach the global optima, we employ a *Gaussian process* (GP) (Rasmussen & Williams, 2006) to model the black-box function, which builds a posterior belief of the function using the history of function queries and their feedback. However, when constructing this posterior belief in BO-SDF problems, we are faced with the important question of *how to handle the randomly delayed feedback whose observations are not available*. To this end, we replace the unavailable delayed feedback by the minimum function value,¹ which we refer to as *censored feedback*. The use of censored feedback, interestingly, improves the exploration in BO-SDF problems (Section 3.1) and hence leads to better theoretical and empirical performances. With the posterior belief built using censored feedback, we propose algorithms using upper confidence bound (UCB) and Thompson sampling (TS), both of which enjoy sub-linear regret guarantees. Specifically, our contributions are as follows:

- We introduce and formalize the notion of *censored feedback* in the BO-SDF problem, and propose UCB- and TS-based algorithms with sub-linear regret guarantees (Section 3).
- Since batch BO is a special case of BO-SDF, we apply our proposed algorithms (Section 3) to batch BO and show that our algorithms enjoy tighter regret upper bounds than classic batch BO algorithms (Section 4). This gain is mainly attributed to the censored feedback, which leads to a better exploration policy.
- We extend our algorithms for the BO-SDF problem to contextual Gaussian process bandits (Krause & Ong, 2011) with stochastic delayed feedback in Section 5. This new contribution is itself of independent interest.
- Our experimental results in Section 6 validate the different performance aspects of our proposed algorithms on synthetic and real-life datasets.

2. Problem Setting

This paper considers the Bayesian optimization (BO) problem where the function evaluations (feedback) are

¹In many problems, the minimum value of a function is known, e.g., a user’s minimum response in online recommendation systems is ‘no click’ (i.e., 0) and the minimum accuracy for hyperparameter tuning of machine learning models is ‘0’.

available to the learner after a random delay. Let $\mathcal{Q} \subset \mathbb{R}^n$ be a finite domain of all function queries where $n \geq 1$.² The learner selects a new function query as soon as enough resources are available (or when experiment parallelization is possible). We denote the t -th query by $x_t \in \mathcal{Q}$. After the learner selects the query x_t to evaluate the unknown black-box function f , the environment generates a noisy function evaluation, denoted by feedback $y_t = f(x_t) + \varepsilon_t$. We assume that ε_t is an R -sub-Gaussian noise. The learner observes y_t only after a stochastic delay d_t , which is generated from an unknown distribution \mathcal{D} .

We refer to this new BO problem as ‘*Bayesian Optimization under Stochastic Delayed Feedback*’ (BO-SDF). The unknown function f , query space \mathcal{Q} , and unknown delay distribution \mathcal{D} identify an instance of the BO-SDF problem. The optimal query (x^*) has the maximum function value (global maximizer), i.e., $x^* \in \operatorname{argmax}_{x \in \mathcal{Q}} f(x)$. After selecting a query x_t , the learner incurs a penalty (or *instantaneous regret*) $p_t = f(x^*) - f(x_t)$.

Since the optimal function query is unknown, we sequentially estimate this using the available information of the selected queries and observed feedback. Our goal is to learn a sequential policy for selecting queries that finds the optimal query (or global maximizer) as quickly as possible. There are two common performance measures for evaluating a sequential policy. The first performance measure is *simple regret*. Let x_t be the t -th function query selected by the policy. Then, after observing T function evaluations, the simple regret is $r_T = f(x^*) - \max_{t \in \{1, \dots, T\}} f(x_t)$. Any good policy must have no regret, i.e., $\lim_{T \rightarrow \infty} r_T = 0$. The second performance measure of the policy is *cumulative regret* which is the sum of total penalties incurred by the learner. After observing T function evaluations, the cumulative regret of a policy is given by $\mathfrak{R}_T = \sum_{t=1}^T (f(x^*) - f(x_t))$. Any good policy should have sub-linear regret, i.e., $\lim_{T \rightarrow \infty} \mathfrak{R}_T/T = 0$. Even though simple regret and cumulative regret are different performance measures, having a policy with no regret or sub-linear regret implies that the policy will eventually converge to the optimal query (or global optimizer).

3. BO under Stochastic Delayed Feedback

A function estimator is the main component of any BO problem for achieving good performance. We use a Gaussian process (GP) as a surrogate for the posterior belief of the unknown function (Rasmussen & Williams, 2006; Srinivas et al., 2010). To deal with the delayed feedback, we will introduce the notion of *censored feedback* in GPs, in which the delayed feedback is replaced by the minimum

²We assume \mathcal{Q} to be finite, but it is straightforward to extend our theoretical results to a compact domain using a suitable discretization scheme (Li & Scarlett, 2022, Lemma 2).

function value. Finally, we will propose UCB- and TS-based algorithms with sub-linear regret guarantees.

3.1. Estimating Function using Gaussian Processes

We assume that the unknown function f belongs to the *reproducing kernel Hilbert space* (RKHS) associated with a kernel k , i.e., $\|f\|_k \leq \mathcal{B}_f$ where $\mathcal{B}_f > 0$. By following standard practice in the literature, we assume without loss of generality that $k(x, x') \leq 1$ for all $x, x' \in \mathcal{Q}$. We also assume that feedback is bounded, i.e., $|y_t| \leq \mathcal{B}_y$ for all $t \geq 1$. This assumption is not restrictive since any bounded function can be re-scaled to satisfy this boundedness requirement. For example, the validation accuracy of machine learning models in our hyperparameter tuning experiment in Section 6 is bounded in $[0, 1]$.

Censored Feedback. The learner has to efficiently exploit available information about queries and feedback to achieve a better performance. The estimated GP posterior belief has two parts, i.e., mean function and covariance function. The posterior covariance function only needs the function queries for building tighter confidence intervals, whereas the posterior mean function needs both queries and their feedback. Incorporating the queries with delayed feedback in updating the posterior mean function leads to a better estimate, but the question is *how to do it*. One possible solution is to replace the delayed feedback with some value, but then the question is *what should that value be*.

To pick the suitable value for replacing delayed feedback, we motivate ourselves from real-life applications like the online movies recommendations problem. When an online platform recommends a movie, it expects the following responses from the users – not watching the movie, watching some part of the movie, and watching the entire movie. The platform can map these responses to $[0, 1]$ where 0 is assigned for ‘not watching the movie’, 1 for ‘watching the entire movie’, and an appropriate value in $(0, 1)$ for ‘watching some part of the movie.’ There are two reasons for the delay in the user’s response: The user does not like the recommended movie or does not have time to watch the movie. Before the user starts watching the movie, the platform can replace the user’s delayed response with ‘not watching the movie’ (i.e., ‘0’) and update it later when more information is available. Therefore, we can replace the delayed feedback with the minimum function value.³ We refer to this replaced feedback as *censored feedback*. The idea of censored feedback is also used in other problems such as censoring delayed binary feedback (Vernade et al.,

³It is possible to replace the delayed feedback with other values. When it is replaced by the current GP posterior mean, it has the same effect as the *hallucination* in batch BO (more details in Section 4). Another possible value for replacing the delayed feedback is the k -nearest posterior means where $k > 0$.

2020) and censoring losses (Verma et al., 2019; 2021) or rewards (Verma & Hanawal, 2020) in online resource allocation problems.

When the minimum function value replaces the delayed feedback, the posterior mean around that query becomes small, consequently ensuring that the learner will not select the queries around the queries with delayed feedback before exploring the other parts of the query’s domain. Therefore, the censored feedback leads to better exploration than replacing delayed feedback with a larger value (e.g., current posterior mean). However, the queries with delayed feedback need to be stored so that the posterior mean can be updated when feedback is revealed. Due to several reasons (e.g., limited storage, making resources available to new queries), the learner cannot keep all queries with delayed feedback and replace the old queries with the latest ones.⁴ It can degrade the algorithm’s performance, as discussed in Section 3.3 and demonstrated by experiments in Section 6.

Let an *incomplete query* represent the function query with unobserved feedback and m be the number of incomplete queries the learner can store. We assume that the random delay for a query is the number of new queries that need to be started (i.e., number of iterations elapsed) before observing its feedback. Therefore, the delay d_t only takes values of non-negative integers.⁵ The censored feedback of the s -th query before selecting the t -th query is denoted by $\tilde{y}_{s,t} = y_s \mathbb{1}\{d_s \leq \min(m, t - s)\}$. That is, the learner censors the incomplete queries by assigning 0 to them.⁶ Now, the GP posterior mean and covariance functions can be expressed using available function queries and their censored feedback, as follows:

$$\begin{aligned} \mu_{t-1}(x) &= \mathbf{k}_{t-1}(x)^\top \mathbf{K}_{t,\lambda}^{-1} \tilde{\mathbf{y}}_{t-1}, \\ \sigma_{t-1}^2(x, x') &= k(x, x') - \mathbf{k}_{t-1}(x)^\top \mathbf{K}_{t,\lambda}^{-1} \mathbf{k}_{t-1}(x') \end{aligned} \quad (1)$$

where $\mathbf{K}_{t,\lambda} = \mathbf{K}_{t-1} + \lambda \mathbf{I}$, $\mathbf{k}_{t-1}(x) = (k(x, x_{t'}))_{t' \in [t-1]}^\top$, $\tilde{\mathbf{y}}_{t-1} = (\tilde{y}_{s,t})_{s \in [t-1]}^\top$, $\mathbf{K}_{t-1} = (k(x_{t'}, x_{t''}))_{t', t'' \in [t-1]}$, and λ is a regularization parameter that ensures $\mathbf{K}_{t,\lambda}$ is a positive definite matrix. We denote $\sigma_{t-1}^2(x) = \sigma_{t-1}^2(x, x)$.

⁴Our motivation for removing older queries comes from online recommendation problems where the chance of observing feedback diminishes with time. The learner can also use a more complicated scheme for query removal.

⁵We have introduced the delay based on *iterations* here to bring out the main ideas of censored feedback. Any suitable discrete distribution (e.g., Poisson) can model such delays. We have also introduced the censored feedback with *time-based* delay, which is more practical and can be modeled using a suitable continuous distribution (see Appendix A.4 for more details).

⁶We can make the minimum value of any bounded function 0 by adding to it the negative of the minimum function value. It does not change the optimizer since the optimizer of a function is invariant to the constant function shift. By doing this, we can represent censored feedback by an indicator function, making censored feedback easier to handle in theoretical analysis.

3.2. Algorithms for the BO-SDF problem

After having a posterior belief of the unknown function, the learner can decide which query needs to be selected for the subsequent function evaluation. Since the feedback is only observed for the selected query, the learner needs to deal with the exploration-exploitation trade-off. We use UCB- and TS-based algorithms for BO-SDF problems that handle the exploration-exploitation trade-off efficiently.

UCB-based Algorithm for BO-SDF problems. UCB is an effective and widely-used technique for dealing with the exploration-exploitation trade-off in various sequential decision-making problems (Auer et al., 2002; Garivier & Cappé, 2011). We propose a UCB-based algorithm named **GP-UCB-SDF** for the BO-SDF problems. It works as follows: When the learner is ready to start the t -th function query, he firstly updates the GP posterior mean and standard deviation defined in Eq. (1) using available censored noisy feedback. Then, he selects the input x_t for the next function evaluation by maximizing the following UCB value:

$$x_t = \operatorname{argmax}_{x \in \mathcal{Q}} (\mu_{t-1} + \nu_t \sigma_{t-1}(x)) \quad (2)$$

where $\nu_t = \mathcal{B}_y \sum_{s=t-m}^{t-1} \sigma_{t-1}(x_s) + \beta_t$, $\beta_t = \mathcal{B}_f + (R + \mathcal{B}_y) \sqrt{2(\gamma_{t-1} + 1 + \log(2/\delta))}$, $\delta \in (0, 1)$, and γ_t is the maximum information gain about the function f from any set of t function queries (Srinivas et al., 2010).

GP-UCB-SDF UCB-based Algorithm for BO-SDF

- 1: **Input:** $\lambda > 0, m \in \mathbb{N}$
 - 2: **for** $t = 1, 2, 3, \dots$ **do**
 - 3: Update GP posterior defined in Eq. (1) using available censored noisy feedback
 - 4: Select input x_t using Eq. (2) and compute f at x_t
 - 5: Keep observing the noisy feedback until the starting of next function evaluation
 - 6: **end for**
-

After selecting the input x_t , the learner queries the unknown function $f(x_t)$ and the environment generates a noisy feedback y_t with an associated delay $d_t \sim \mathcal{D}$. The learner observes y_t only after a delay of d_t iff $d_t \leq m$. Before starting a new function evaluation, the learner keeps observing the noisy feedback of past queries. The same process is repeated for the subsequent function evaluations.

TS-based Algorithm for BO-SDF problems. In contrast to the deterministic UCB, TS (Thompson, 1933; Agrawal & Goyal, 2012; 2013a; Kaufmann et al., 2012) is based on Bayesian updates that select inputs according to their probability of being the best input. Many works have shown that TS is empirically superior to UCB-based algorithms (Chapelle & Li, 2011; Chowdhury & Gopalan, 2019). We name our TS-based algorithm **GP-TS-SDF** which works

similarly to **GP-UCB-SDF**, except that the t -th function query is selected as follows:

$$x_t = \operatorname{argmax}_{x \in \mathcal{Q}} f_t(x) \quad (3)$$

where $f_t \sim \mathcal{GP}(\mu_{t-1}(\cdot), \nu_t^2 \sigma_{t-1}^2(\cdot))$.

GP-TS-SDF TS-based Algorithm for BO-SDF

- 1: **Input:** $\lambda > 0, m \in \mathbb{N}$
 - 2: **for** $t = 1, 2, 3, \dots$ **do**
 - 3: Update GP posterior defined in Eq. (1) using available censored noisy feedback
 - 4: Select input x_t using Eq. (3) and compute f at x_t
 - 5: Keep observing the noisy feedback until the starting of next function evaluation
 - 6: **end for**
-

3.3. Regret Analysis

Let $\rho_m = \mathbb{P}(d_s \leq m)$ for $s \geq 1$ be the probability of observing delayed feedback within the next m iterations. The following result gives the confidence bounds of GP posterior mean function and is important to our subsequent regret analysis:

Theorem 1 (Confidence Ellipsoid). *Let $x \in \mathcal{Q}$, $\lambda > 1$, and $t \geq 1$. Then, with probability at least $1 - \delta$,*

$$|\mu_{t-1}(x) - \rho_m f(x)| \leq \nu_t \sigma_{t-1}(x).$$

The detailed proofs of Theorem 1 and other results below are given in Appendix A. Now, we state the regret upper bounds of our proposed **GP-UCB-SDF** and **GP-TS-SDF**:

Theorem 2 (**GP-UCB-SDF**). *Let $C_1 = \sqrt{2/\log(1 + \lambda^{-1})}$. Then, with probability at least $1 - \delta$,*

$$\mathfrak{R}_T \leq \frac{2}{\rho_m} \left(C_1 \beta_T \sqrt{T \gamma_T} + m \mathcal{B}_y C_1^2 \gamma_T \right).$$

As expected, the regret inversely depends on ρ_m , i.e., the probability of observing delayed feedback within the next m iterations. If we ignore the logarithmic factors and constants in Theorem 2, then the regret of **GP-UCB-SDF** is upper bounded by $\tilde{O}\left(\rho_m^{-1} \left(\gamma_T \sqrt{T} + m \gamma_T\right)\right)$.

Theorem 3 (**GP-TS-SDF**). *With probability at least $1 - \delta$,*

$$\mathfrak{R}_T = \tilde{O}\left(\frac{1}{\rho_m} \left(\sqrt{T \gamma_T} (\sqrt{\gamma_T} + 1) + m (\gamma_T + \sqrt{T})\right)\right).$$

The regret bounds in both Theorems 2 and 3 are *sub-linear* for the commonly used squared exponential (SE) kernel, for which $\gamma_T = O(\log^{n+1}(T))$ (Srinivas et al., 2010).

By following the common practice in the BO literature (Contal et al., 2013; Kandasamy et al., 2018), we can

upper-bound the simple regret by the average cumulative regret, i.e., $r_T \leq \mathfrak{R}_T/T$.

Discussion on the trade-off regarding m . A large value of m ensures that fewer queries are discarded and favors the regret bound (i.e., via the term $1/\rho_m$) by making ρ_m large. However, a large m also means that our algorithm allows larger delays, which consequently forces us to enlarge the width of the confidence interval (Theorem 1) to account for them. It makes the regret bound worse, which is also reflected by the linear dependence on m in the second term of the regret upper bounds (Theorems 2 and 3). As m is directly related to the storage requirement of input queries whose feedback are yet to be observed, a large value of m is constrained by the resources available to the learner.

Input-dependent random delays. In some applications, the random delay can depend on the function query: for example, a larger number of hidden layers increases the training time of a neural network. Our regret bounds still hold with input-dependent random delays by redefining ρ_m for T queries as $\rho_m = \min_{t \in [T]} \mathbb{P}\{d_t \leq m\}$.

Limitations of censored feedback. Censoring, i.e., replacing values of incomplete queries with the minimum function values, leads to an aggressive exploration strategy. However, censoring is required to ensure that our BO method does not unnecessarily explore the incomplete queries or queries near them. Moreover, we can reduce the effect of this aggressive exploration by increasing the value of m , since a larger m leads to less aggressive censoring. Another limitation of the censoring method is that it needs to know the minimum function value. However, when the minimum value is unknown, we can use a suitable lower bound as a proxy for the minimum function value. However, regret analysis of such a method can be challenging and left for future research.

4. Improved Algorithms for Batch BO

Notably, our BO-SDF problem subsumes batch BO as a special case. Specifically, the function queries in batch BO can be viewed as being sequentially selected, in which some queries need to be selected while some other queries of the same batch are still *incomplete* (Desautels et al., 2014). Therefore, in batch BO with a batch size of \mathcal{B}_x , the incomplete queries can be viewed as delayed feedback with fixed delays s.t. $d_s \leq \mathcal{B}_x - 1$. In this case, by choosing $m = \mathcal{B}_x - 1$, we can ensure that $\rho_m = \mathbb{P}(d_s \leq \mathcal{B}_x - 1) = 1$. As a result, our method of censoring (Section 3.1) gives a better treatment to the incomplete queries than the classic technique from batch BO, which we will demonstrate next via both intuitive justification and regret comparison.

The classic technique to handle the incomplete queries in batch BO is *hallucination* which was proposed by the

GP-BUCB algorithm from (Desautels et al., 2014). It has also been adopted by a number of its extensions (Daxberger & Low, 2017; Chowdhury & Gopalan, 2019). In particular, the feedback of incomplete queries is hallucinated to be (i.e., censored with) the GP posterior mean computed using only the completed queries (excluding the incomplete queries). It is equivalent to keeping the GP posterior mean unchanged (i.e., unaffected by the incomplete queries) while updating the GP posterior variance using all selected queries (including the incomplete queries). Note that in contrast to the hallucination, our censoring technique sets the incomplete observations to the minimum function value (i.e., 0) (Section 3.1). As a result, compared with the hallucination where the GP posterior mean is unaffected by the incomplete queries, our censoring can *reduce the value of the GP posterior mean at the incomplete queries* (because we have treated their observations as if they are 0). As a result, our censoring can *discourage these incomplete queries from being unnecessarily queried again*. It encourages the algorithm to explore other unexplored regions, hence leading to better exploration.

Interestingly, the better exploration resulting from our censoring technique also translates to a tighter regret upper bound. Specifically, in batch BO with a batch size of \mathcal{B}_x , after plugging in $m = \mathcal{B}_x - 1$ and $\rho_m = 1$, our next result gives the regret upper bound of **GP-UCB-SDF**.

Proposition 1. *With probability at least $1 - \delta$,*

$$\mathfrak{R}_T = \tilde{O}\left(\gamma_T \sqrt{T} + \mathcal{B}_x \gamma_T\right).$$

As long as $\gamma_T = \Omega(\log T)$ which holds for most commonly used kernels such as the squared exponential (SE) and Matérn kernels, both terms in Proposition 1 grow slower than $\mathfrak{R}_T = \tilde{O}(\gamma_T \sqrt{T} \exp(\gamma_{\mathcal{B}_x-1}))$ which is the regret upper bound of the GP-BUCB algorithm based on hallucination (Desautels et al., 2014). Importantly, to achieve a sub-linear regret upper bound (i.e., to ensure that $\exp(\gamma_{\mathcal{B}_x-1})$ can be upper-bounded by a constant), GP-BUCB (Desautels et al., 2014) and its extensions (Daxberger & Low, 2017; Kandasamy et al., 2018; Chowdhury & Gopalan, 2019) all require a special initialization scheme (i.e., uncertainty sampling). It is often found unnecessary in practice (Kandasamy et al., 2018) and hence represents a gap between theory and practice. Interestingly, our tighter regret upper bound from Proposition 1 has removed the dependence on $\exp(\gamma_{\mathcal{B}_x-1})$ and hence eliminated the need for uncertainty sampling, thereby closing this gap between theory and practice.

Note that our discussions above regarding the advantage of censoring over hallucination also apply to TS-based algorithms. In particular, in batch BO with a batch size of \mathcal{B}_x , the regret upper bound of **GP-TS-SDF** is given by the following result.

Proposition 2. *With probability at least $1 - \delta$,*

$$\mathfrak{R}_T = \tilde{O}\left(\sqrt{T}\gamma_T(\sqrt{\gamma_T} + 1) + \mathcal{B}_x\gamma_T + \mathcal{B}_x\sqrt{T}\right).$$

All three terms in Proposition 2 grow slower than $\mathfrak{R}_T = \tilde{O}(\exp(\gamma_{\mathcal{B}_x-1})\sqrt{T}\gamma_T(\sqrt{\gamma_T} + 1))$ which is the regret upper bound of the GP-BTS algorithm based on hallucination (Chowdhury & Gopalan, 2019) as long as $\gamma_T = \Omega(\log T)$.

To summarize, though batch BO is only a special case of our BO-SDF problem, we have made non-trivial contributions to the algorithm and theory of batch BO. Our GP-UCB-SDF and GP-TS-SDF algorithms and our theoretical analyses may serve as inspiration for future works on batch BO.

5. Contextual Gaussian Process Bandits under Stochastic Delayed Feedback

Before making a decision, sometimes additional information is available to the learner in many real-life scenarios (e.g., users’ profile information for an online platform and patients’ medical history before clinical trials). This additional information is referred to as *context* in the literature, and the value of a function depends on the context (Agrawal & Goyal, 2013b; Chu et al., 2011; Krause & Ong, 2011; Li et al., 2010). The learner can design the algorithms to exploit the contextual information to make a better decision. We can extend our results for the BO-SDF problem to the contextual Gaussian process bandits (Krause & Ong, 2011) with stochastic delayed feedback where a non-linear function maps a context to the feedback.

Let $\mathcal{Q} \subset \mathbb{R}^n$ be a finite set of available actions (or queries) and $\mathcal{Z} \subset \mathbb{R}^{n'}$ be a set of all contexts where $n, n' \geq 1$. In round t , the environment generates a context $z_t \in \mathcal{Z}$. Then, the learner selects an action $x_t \in \mathcal{Q}$ and observes noisy feedback denoted by $y_t = g(z_t, x_t) + \varepsilon_t$. We assume that $g: \mathcal{Z} \times \mathcal{Q} \rightarrow \mathbb{R}$ and ε_t is an R -sub-Gaussian noise. The y_t is only observed after a random delay d_t , which is generated from an unknown distribution \mathcal{D} .

We refer to this new problem as ‘Contextual Gaussian process Bandits under Stochastic Delayed Feedback’ (CGB-SDF). The unknown function g , context space \mathcal{Z} , action space \mathcal{Q} , and unknown delay distribution \mathcal{D} identify an instance of the CGB-SDF problem. The optimal action x_t^* for a given context z_t is the action where the function g has its maximum value, i.e., $x_t^* = \operatorname{argmax}_{x \in \mathcal{Q}} g(z_t, x)$. The learner incurs a penalty of $(g(z_t, x_t^*) - g(z_t, x_t))$ for a context z_t and action x_t . We aim to learn a policy that achieves the minimum cumulative penalties or *regret* which is given by $\mathfrak{R}_T = \sum_{t=1}^T (g(z_t, x_t^*) - g(z_t, x_t))$.

Our goal is to learn a policy that has a small sub-linear regret, i.e., $\lim_{T \rightarrow \infty} \mathfrak{R}_T/T \rightarrow 0$. The sub-linear regret here implies that the policy will eventually select the

optimal action for a given context. We have adapted our algorithms for the BO-SDF problem to the CGB-SDF problem and shown that they also enjoy similar sub-linear regret guarantees; see Appendix B for more details.

6. Experiments

We compare with previous baseline methods that can be applied to BO-SDF problems after modifications. Firstly, we compare with standard GP-UCB (Srinivas et al., 2010) which ignores the delayed observations when selecting the next query. Note that GP-UCB is likely to repeat previously selected queries in some iterations when the GP posterior remains unchanged (i.e., if no observation is newly collected between two iterations). Next, we also compare with the batch BO method of *asy-TS* (Kandasamy et al., 2018) which, similarly to GP-UCB, ignores the delayed observations when using TS to choose the next query. Lastly, we compare with the batch BO methods of *GP-BUCB* (Desautels et al., 2014) and *GP-BTS* (Chowdhury & Gopalan, 2019) which handle the delayed observations via hallucination (Section 4). Following the suggestion from (Vernade et al., 2020), if not specified otherwise, we set m to be 2μ where μ is the mean of the delay distribution. However, this is for convenience only since we will demonstrate (Section 6.1) that our algorithms consistently achieve competitive performances as long as m is large enough. Therefore, μ does not need to be known in practice. We defer some experimental details to Appendix C due to space constraint. Our code is released at <https://github.com/daizhongxiang/BO-SDF>.

6.1. Synthetic Experiments

For this experiment, we use a Poisson distribution (with a mean parameter μ) as the delay distribution and sample the objective function f from a GP with an SE kernel defined on a discrete 1-dimensional domain. Fig. 1a plots the performances of our GP-UCB-SDF for different delay distributions and shows that larger delays (i.e., larger μ ’s) result in worse performances. Fig. 1b shows that for a fixed delay distribution (i.e., $\mu = 10$), an overly small value of m leads to large regrets since it causes our algorithm to ignore a large number of delayed observations (i.e., overly aggressive censoring). On the other hand, a large m is desirable since $m = 20 = 2\mu$ and $m = 40$ produce comparable performances.

However, an overly large value of m may exert an excessive resource requirement since we may need to cater to many pending function evaluations. Therefore, when the delay distribution (hence μ) may be unknown, we recommend using a large value of m that is allowed by the resource constraints. We now compare the performances of our GP-UCB-SDF with the baselines of GP-UCB and




Figure 1: Performances of our **GP-UCB-SDF** for (a) different delay distributions and (b) different values of m (μ is fixed at $\mu = 10$). Performances of **GP-UCB-SDF** and baseline methods for (c) stochastic and (d) deterministic delay distributions.




Figure 2: Performances for hyperparameter tuning of SVM: UBC-based methods with (a) stochastic and (b) deterministic delay distributions, and TS-based methods with (c) stochastic and (d) deterministic delay distributions. Performances for hyperparameter tuning of CNN: UBC-based methods with (e) stochastic and (f) deterministic delay distributions, and TS-based methods with (g) stochastic and (h) deterministic delay distributions.

GP-BUCB for both stochastic ($\mu = 10$, Fig. 1c) and deterministic (delays fixed at 10, Fig. 1d) delay distributions. The figures show that our **GP-UCB-SDF** achieves smaller simple regret than the baselines in both settings.

6.2. Real-world Experiments

In this section, we perform real-world experiments on hyperparameter tuning for ML models under two realistic settings. In the first setting, we consider stochastic delay distributions where the delays are sampled from a Poisson distribution with a mean parameter $\mu = 10$. This setting is used to simulate real-world scenarios where the evaluations of different hyperparameter configurations may be stochastically delayed because they may require a different amount of computations (e.g., training a neural

network with more layers is more time-consuming), or the machines deployed in the experiment may have different computational capabilities. The second setting considers fixed delay distributions where every delay is fixed at $\mathcal{B}_x - 1 = 10$, which is used to emulate real-world asynchronous batch BO problems with a batch size of $\mathcal{B}_x = 11$.

We now tune the hyperparameters of three different ML models to demonstrate the consistency of our algorithms. Specifically, we tune two hyperparameters of SVM (i.e., the penalty parameter and RBF kernel parameter) used for diabetes diagnosis, three hyperparameters of convolution neural networks (CNNs) (i.e., the batch size, learning rate, and learning rate decay) using the MNIST dataset, and the same three hyperparameters of logistic regression (LR) used for breast cancer detection. The experimental results

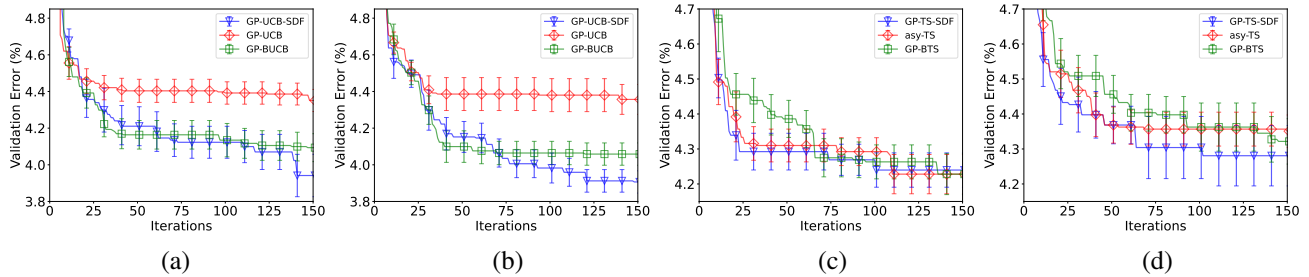


Figure 3: Performances for hyperparameter tuning of LR, UCB-based methods for (a) stochastic and (b) deterministic delay distributions, and TS-based methods for (a) stochastic and (b) deterministic delay distributions.

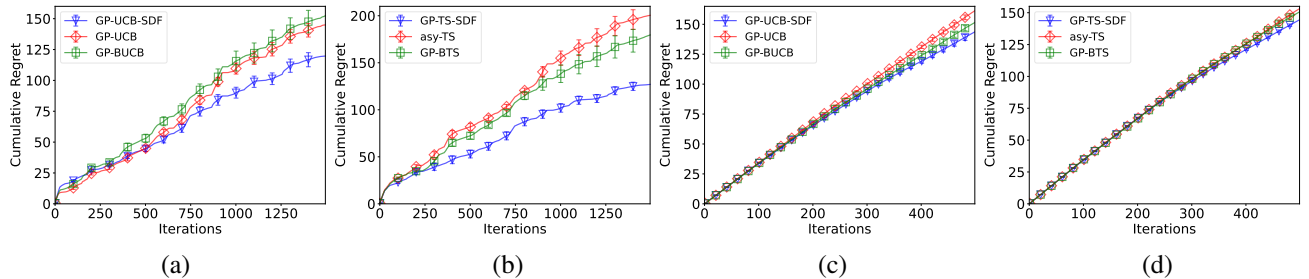


Figure 4: Cumulative regret of contextual GP bandit for (a) UCB-based and (b) TS-based methods for multi-task BO, and (c) UCB-based and (d) TS-based methods for non-stationary BO.

on hyperparameter tuning for SVM and CNN are plotted in Fig. 2, including the results for UCB- and TS-based algorithms for both settings. The results for hyperparameter tuning of LR are shown in Fig. 3. The figures show that our GP-UCB-SDF and GP-TS-SDF consistently outperform the corresponding baselines in both settings with stochastic and deterministic delay distributions. The results here verify the practical effectiveness of our proposed algorithms in real-world problems.

6.3. Real-world Experiments on Contextual Gaussian Process Bandits

In this section, we apply our algorithms to two real-world contextual GP bandit problems (Section 5). Our first experiment performs multi-task BO in the contextual GP bandit setting following the work of Krause & Ong (2011). We use the tabular benchmark dataset for SVM hyperparameter tuning from the work of Wistuba et al. (2015), which consists of the validation accuracy evaluated on a discrete domain of six SVM hyperparameters for 50 different classification tasks. Each of the 50 tasks is associated with a set of meta-features which are used as the contexts z_t . We sequentially tackle the 50 tasks and perform hyperparameter tuning for each task for 30 iterations, i.e., the context z_t for every task is repeated consecutively for 30 times and the contexts for all tasks arrive sequentially.

Our second experiment performs non-stationary BO in

the contextual GP bandit setting. In some real-world BO problems, the objective function is non-stationary (i.e., changing over time). For example, when we tune the hyperparameters of an ML model for diabetes diagnosis (Section 6.2), we may need to perform the hyperparameter tuning task periodically because the size of the dataset is progressively growing due to sequentially arriving patient data. We divide the entire dataset into 20 progressively growing datasets and assign an index $z_t \in \{0, \dots, 19\}$ to each dataset such that a smaller dataset has a smaller index. Then, we use the indices as the contexts because two datasets whose indices are close share many common data points and are hence expected to lead to similar performances. We let the contexts arrive in the order of their indices to simulate sequentially arriving patients and again repeat each context for 30 iterations.

For both experiments, we use a stochastic delay distribution (i.e., Poisson distribution with $\mu = 3$) and set $m = 6$. The cumulative regrets of different algorithms are shown in Fig. 4, in which both GP-UCB-SDF and GP-TS-SDF incur smaller regrets than the other baselines.

7. Related Work

We have developed algorithms for the BO problems with stochastic delayed feedback. In the following, we briefly review the papers relevant to our problem.

Batch BO. Batch BO is the closest BO setting to our BO-SDF problem. Many algorithms have been proposed for different variants of batch BO problems in the literature (Desautels et al., 2014; González et al., 2016; Daxberger & Low, 2017; Kandasamy et al., 2018; Gong et al., 2019; Chowdhury & Gopalan, 2019; Balandat et al., 2020). However, the delay in observing feedback observations in all these batch BO problems is fixed. Some batch BO algorithms are based on the UCB index (Desautels et al., 2014; Daxberger & Low, 2017), while some others are based on TS (Kandasamy et al., 2018; Chowdhury & Gopalan, 2019; Gong et al., 2019). The batch BO methods like ‘local penalization’ (González et al., 2016) and ‘fantasize’ (Balandat et al., 2020) can not work with stochastically delayed feedback and have no theoretical guarantees. All variants of batch BO problems can be treated as special cases of the BO-SDF problem by fixing the delays appropriately. Therefore, our UCB- and TS-based algorithms can be used for any batch BO problem.

Contextual Gaussian Process (GP) Bandits. Several prior works consider the availability of additional information to the learner before making the decision. This line of work is popularly known as contextual bandits (Li et al., 2010). It has many applications in advertising, web search, and e-commerce. In contextual bandits, the mean reward is a function of context and often parameterized (e.g., linear (Li et al., 2010; Chu et al., 2011; Agrawal & Goyal, 2013b), GLM (Li et al., 2017), or non-linear (Valko et al., 2013)). Contextual GP bandits model the posterior belief of a non-linear function using GPs (Krause & Ong, 2011). To the best of our knowledge, our work is the first to introduce stochastic delayed feedback into contextual GP bandits.

Stochastic Delayed Feedback. Due to several practical applications, some works have explored fixed or stochastic delayed feedback in the stochastic multi-armed bandits (MAB) (Slivkins, 2019; Lattimore & Szepesvári, 2020) which is a simpler sequential decision-making framework than BO. In the MAB literature, the problems ranging from known fixed delay (Joulani et al., 2013), unknown delay (Zhong et al., 2017), unknown stochastic delay (Vernade et al., 2017; Grover et al., 2018), and unknown delays in the adversarial setting (Li et al., 2019) are well-studied. The setting of the unknown stochastic delay is further extended to the linear bandits (Zhou et al., 2019; Vernade et al., 2020) which is closest to our BO-SDF problem. However, these works on linear bandits assume that the black-box function is linear and has binary feedback, which may not hold in real-life applications.

8. Conclusion

This paper has studied the Bayesian optimization (BO) problem with stochastic delayed feedback (BO-SDF). We

have used Gaussian processes as a surrogate model for the posterior belief of the unknown black-box function. To exploit the information of queries with delayed feedback (i.e., incomplete queries) for updating the posterior mean function, we have used censored feedback that assigns the minimum function value (or ‘0’) to the delayed feedback. It discourages the learner from selecting the incomplete queries (or queries near them) when choosing the next query and hence leads to more efficient exploration.

We have proposed UCB- and TS-based algorithms for the BO-SDF problem and given their sub-linear regret guarantees. Interestingly, when delays are fixed in a certain way, batch BO becomes a special case of the BO-SDF problem. Hence, we can adapt our algorithms for the batch BO problem. We have shown that our algorithms are theoretically and empirically better and do not need the special initialization scheme required for the existing batch BO algorithms. We also extended our algorithms to contextual Gaussian process bandits with stochastic delayed feedback, which is in itself a non-trivial contribution.

As the learning environment keeps changing in real life, an interesting future direction will be considering the BO-SDF problem with a non-stationary function. We also plan to generalize our algorithms to nonmyopic BO (Kharkovskii et al., 2020b; Ling et al., 2016), high-dimensional BO (Hoang et al., 2018), private outsourced BO (Kharkovskii et al., 2020a), preferential BO (Nguyen et al., 2021d), federated/collaborative BO (Dai et al., 2020b; 2021; Sim et al., 2021), meta-BO (Dai et al., 2022), and multi-fidelity BO (Zhang et al., 2017; 2019) settings, handle risks (Nguyen et al., 2021b;a; Tay et al., 2022) and information-theoretic acquisition functions (Nguyen et al., 2021c;e), incorporate early stopping (Dai et al., 2019), and/or recursive reasoning (Dai et al., 2020a), and consider its application to neural architecture search (Shu et al., 2022a;b) and inverse reinforcement learning (Balakrishnan et al., 2020). For applications with a huge budget of function evaluations, we like to couple our algorithms with the use of distributed/decentralized (Chen et al., 2012; 2013a;b; 2015; Hoang et al., 2016; 2019; Low et al., 2015; Ouyang & Low, 2018), online/stochastic (Hoang et al., 2015; 2017; Low et al., 2014; Xu et al., 2014; Yu et al., 2019b), or deep sparse GP models (Yu et al., 2019a; 2021) to represent the belief of the unknown objective function efficiently.

Acknowledgements

This research/project is supported by A*STAR under its RIE2020 Advanced Manufacturing and Engineering (AME) Industry Alignment Fund – Pre Positioning (IAF-PP) (Award A19E4a0101) and by the Singapore Ministry of Education Academic Research Fund Tier 1.

References

- Agrawal, S. and Goyal, N. Analysis of Thompson sampling for the multi-armed bandit problem. In *Proc. COLT*, pp. 39.1–39.26, 2012.
- Agrawal, S. and Goyal, N. Further optimal regret bounds for Thompson sampling. In *Proc. AISTATS*, pp. 99–107, 2013a.
- Agrawal, S. and Goyal, N. Thompson sampling for contextual bandits with linear payoffs. In *Proc. ICML*, pp. 127–135, 2013b.
- Auer, P., Cesa-Bianchi, N., and Fischer, P. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, pp. 235–256, 2002.
- Balakrishnan, S., Nguyen, Q. P., Low, B. K. H., and Soh, H. Efficient exploration of reward functions in inverse reinforcement learning via Bayesian optimization. In *Proc. NeurIPS*, pp. 4187–4198, 2020.
- Balandat, M., Karrer, B., Jiang, D., Daulton, S., Letham, B., Wilson, A. G., and Bakshy, E. BoTorch: a framework for efficient Monte-Carlo Bayesian optimization. In *Proc. NeurIPS*, pp. 21524–21538, 2020.
- Besson, L. and Kaufmann, E. What doubling tricks can and can't do for multi-armed bandits. arXiv preprint arXiv:1803.06971, 2018.
- Chapelle, O. Modeling delayed feedback in display advertising. In *Proc. KDD*, pp. 1097–1105, 2014.
- Chapelle, O. and Li, L. An empirical evaluation of Thompson sampling. In *Proc. NeurIPS*, pp. 2249–2257, 2011.
- 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.
- Chow, S.-C. and Chang, M. *Adaptive Design Methods in Clinical Trials*. CRC Press, 2006.
- Chowdhury, S. R. and Gopalan, A. On kernelized multi-armed bandits. In *Proc. ICML*, pp. 844–853, 2017.
- Chowdhury, S. R. and Gopalan, A. On batch Bayesian optimization. arXiv preprint arXiv:1911.01032, 2019.
- Chu, W., Li, L., Reyzin, L., and Schapire, R. E. Contextual bandits with linear payoff functions. In *Proc. AISTATS*, pp. 208–214, 2011.
- Contal, E., Buffoni, D., Robicquet, A., and Vayatis, N. Parallel Gaussian process optimization with upper confidence bound and pure exploration. In *Proc. ECML/PKDD*, pp. 225–240, 2013.
- Dai, Z., Yu, H., Low, B. K. H., and Jaillet, P. Bayesian optimization meets Bayesian optimal stopping. In *Proc. ICML*, pp. 1496–1506, 2019.
- Dai, Z., Chen, Y., Low, B. K. H., Jaillet, P., and Ho, T.-H. R2-B2: Recursive reasoning-based Bayesian optimization for no-regret learning in games. In *Proc. ICML*, pp. 2291–2301, 2020a.
- Dai, Z., Low, B. K. H., and Jaillet, P. Federated Bayesian optimization via Thompson sampling. In *Proc. NeurIPS*, pp. 9687–9699, 2020b.
- Dai, Z., Low, B. K. H., and Jaillet, P. Differentially private federated Bayesian optimization with distributed exploration. In *Proc. NeurIPS*, pp. 9125–9139, 2021.
- Dai, Z., Chen, Y., Yu, H., Low, B. K. H., and Jaillet, P. On provably robust meta-Bayesian optimization. In *Proc. UAI*, 2022.
- Daxberger, E. A. and Low, B. K. H. Distributed batch Gaussian process optimization. In *Proc. ICML*, pp. 951–960, 2017.
- Desautels, T., Krause, A., and Burdick, J. W. Parallelizing exploration-exploitation tradeoffs in Gaussian process bandit optimization. *JMLR*, 15:3873–3923, 2014.
- Diemert, E., Meynet, J., Galland, P., and Lefortier, D. Attribution modeling increases efficiency of bidding in display advertising. In *Proc. KDD Workshop on AdKDD and TargetAd*, 2017.
- Frazier, P. I. A tutorial on Bayesian optimization. arXiv preprint arXiv:1807.02811, 2018.
- Garivier, A. and Cappé, O. The KL-UCB algorithm for bounded stochastic bandits and beyond. In *Proc. COLT*, pp. 359–376, 2011.

- Garnett, R. *Bayesian Optimization*. Cambridge Univ. Press, 2022.
- Gong, C., Peng, J., and Liu, Q. Quantile Stein variational gradient descent for batch Bayesian optimization. In *Proc. ICML*, pp. 2347–2356, 2019.
- González, J., Dai, Z., Hennig, P., and Lawrence, N. Batch Bayesian optimization via local penalization. In *Proc. AISTATS*, pp. 648–657, 2016.
- Grover, A., Markov, T., Attia, P., Jin, N., Perkins, N., Cheong, B., Chen, M., Yang, Z., Harris, S., Chueh, W., and Ermon, S. Best arm identification in multi-armed bandits with delayed feedback. In *Proc. AISTATS*, pp. 833–842, 2018.
- 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, 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, B. 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*, 2019.
- Joulani, P., Gyorgy, A., and Szepesvári, C. Online learning under delayed feedback. In *Proc. ICML*, pp. 1453–1461, 2013.
- Kandasamy, K., Krishnamurthy, A., Schneider, J., and Póczos, B. Parallelised Bayesian optimisation via Thompson sampling. In *Proc. AISTATS*, pp. 133–142, 2018.
- Kaufmann, E., Korda, N., and Munos, R. Thompson sampling: An asymptotically optimal finite-time analysis. In *Proc. ALT*, pp. 199–213, 2012.
- Kharkovskii, D., Dai, Z., and Low, B. K. H. Private outsourced Bayesian optimization. In *Proc. ICML*, pp. 5231–5242, 2020a.
- Kharkovskii, D., Ling, C. K., and Low, B. K. H. Nonmyopic Gaussian process optimization with macro-actions. In *Proc. AISTATS*, pp. 4593–4604, 2020b.
- Krause, A. and Ong, C. S. Contextual Gaussian process bandit optimization. In *Proc. NeurIPS*, pp. 2447–2455, 2011.
- Lattimore, T. and Szepesvári, C. *Bandit Algorithms*. Cambridge Univ. Press, 2020.
- Li, B., Chen, T., and Giannakis, G. B. Bandit online learning with unknown delays. In *Proc. AISTATS*, pp. 993–1002, 2019.
- Li, L., Chu, W., Langford, J., and Schapire, R. E. A contextual-bandit approach to personalized news article recommendation. In *Proc. WWW*, pp. 661–670, 2010.
- Li, L., Lu, Y., and Zhou, D. Provably optimal algorithms for generalized linear contextual bandits. In *Proc. ICML*, pp. 2071–2080, 2017.
- Li, Z. and Scarlett, J. Gaussian process bandit optimization with few batches. In *Proc. AISTATS*, pp. 92–107, 2022.
- Ling, C. K., Low, B. 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., 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, 2014.
- 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.
- Nguyen, Q. P., Dai, Z., Low, B. K. H., and Jaillet, P. Optimizing conditional value-at-risk of black-box functions. In *Proc. NeurIPS*, pp. 4170–4180, 2021a.
- Nguyen, Q. P., Dai, Z., Low, B. K. H., and Jaillet, P. Value-at-risk optimization with Gaussian processes. In *Proc. ICML*, pp. 8063–8072, 2021b.
- Nguyen, Q. P., Low, B. K. H., and Jaillet, P. An information-theoretic framework for unifying active learning problems. In *Proc. AAAI*, pp. 9126–9134, 2021c.
- Nguyen, Q. P., Tay, S., Low, B. K. H., and Jaillet, P. Top- k ranking Bayesian optimization. In *Proc. AAAI*, pp. 9135–9143, 2021d.
- Nguyen, Q. P., Wu, Z., Low, B. K. H., and Jaillet, P. Trusted-maximizers entropy search for efficient Bayesian optimization. In *Proc. UAI*, pp. 1486–1495, 2021e.

- 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.
- Rasmussen, C. E. and Williams, C. K. I. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., and De Freitas, N. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2015.
- Shu, Y., Cai, S., Dai, Z., Ooi, B. C., and Low, B. K. H. NASI: Label- and data-agnostic neural architecture search at initialization. In *Proc. ICLR*, 2022a.
- Shu, Y., Chen, Y., Dai, Z., and Low, B. K. H. Neural ensemble search via Bayesian sampling. In *Proc. UAI*, 2022b.
- Sim, R. H. L., Zhang, Y., Low, B. K. H., and Jaillet, P. Collaborative Bayesian optimization with fair regret. In *Proc. ICML*, pp. 9691–9701, 2021.
- Slivkins, A. Introduction to multi-armed bandits. *Foundations and Trends® in Machine Learning*, 12(1–2): 1–286, 2019.
- Snoek, J., Larochelle, H., and Adams, R. P. Practical Bayesian optimization of machine learning algorithms. In *Proc. NeurIPS*, 2012.
- 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.
- Takahashi, A. and Suzuki, T. Bayesian optimization design for dose-finding based on toxicity and efficacy outcomes in phase I/II clinical trials. *Pharmaceutical Statistics*, 20(3):422–439, 2021.
- Tay, S. S., Foo, C. S., Urano, D., Leong, R. C. X., and Low, B. K. H. Efficient distributionally robust Bayesian optimization with worst-case sensitivity. In *Proc. ICML*, 2022.
- Thompson, W. R. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3–4):285–294, 1933.
- Ueno, T., Rhone, T. D., Hou, Z., Mizoguchi, T., and Tsuda, K. Combo: An efficient Bayesian optimization library for materials science. *Materials Discovery*, 4:18–21, 2016.
- Valko, M., Korda, N., Munos, R., Flaounas, I., and Cristianini, N. Finite-time analysis of kernelised contextual bandits. In *Proc. UAI*, pp. 654–663, 2013.
- Verma, A. and Hanawal, M. K. Stochastic network utility maximization with unknown utilities: Multi-armed bandits approach. In *Proc. IEEE INFOCOM*, pp. 189–198, 2020.
- Verma, A., Hanawal, M., Rajkumar, A., and Sankaran, R. Censored semi-bandits: A framework for resource allocation with censored feedback. In *Proc. NeurIPS*, pp. 14499–14509, 2019.
- Verma, A., Hanawal, M. K., Rajkumar, A., and Sankaran, R. Censored semi-bandits for resource allocation. arXiv preprint arXiv:2104.05781, 2021.
- Vernade, C., Cappé, O., and Perchet, V. Stochastic bandit models for delayed conversions. In *Proc. UAI*, 2017.
- Vernade, C., Carpentier, A., Lattimore, T., Zappella, G., Ermis, B., and Brueckner, M. Linear bandits with stochastic delayed feedback. In *Proc. ICML*, pp. 9712–9721, 2020.
- Wistuba, M., Schilling, N., and Schmidt-Thieme, L. Learning hyperparameter optimization initializations. In *Proc. IEEE International Conference on Data Science and Advanced Analytics*, 2015.
- 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, B. 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, H., Liu, D., Low, K. H., and Jaillet, P. Convolutional normalizing flows for deep Gaussian processes. In *Proc. IJCNN*, 2021.
- Zhang, Y., Hoang, T. N., Low, B. 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, B. K. H. Bayesian optimization with binary auxiliary information. In *Proc. UAI*, pp. 1222–1232, 2019.
- Zhong, J., Huang, Y., and Liu, J. Asynchronous parallel empirical variance guided algorithms for the thresholding bandit problem. arXiv preprint arXiv:1704.04567, 2017.
- Zhou, Z., Xu, R., and Blanchet, J. Learning in generalized linear contextual bandits with stochastic delays. In *Proc. NeurIPS*, pp. 5197–5208, 2019.

Appendix

A. Proofs of Theoretical Results

A.1. Proof of Theorem 1

We define $\varphi(x) = k(x, \cdot)$ as the feature map associated with the kernel k , i.e., φ maps any input $x \in \mathcal{Q}$ into a feature in the (potentially infinite-dimensional) feature space of the RKHS associated with k . The reproducing property tells us that $f(x) = \langle \varphi(x), f \rangle_k = \varphi(x)^\top f$. Next, we define $V_t(\lambda) \triangleq \sum_{s=1}^{t-1} \varphi(x_s) \varphi(x_s)^\top + \lambda I$, $B_t \triangleq \sum_{s=1}^{t-1} \varphi(x_s) \tilde{y}_{s,t}$, and $\hat{\theta}_t^W \triangleq V_t(\lambda)^{-1} B_t$. Of note, the GP posterior mean (1) can be equivalently expressed as $\mu_{t-1}(x) = \varphi(x)^\top V_t(\lambda)^{-1} B_t$, and the GP posterior variance (1) can be equivalently written as $\sigma_{t-1}^2(x, x) = \sigma_{t-1}^2(x) = \lambda \varphi(x)^\top V_t(\lambda)^{-1} \varphi(x) = \lambda \|\varphi(x)\|_{V_t(\lambda)^{-1}}^2$. In this section, we prove Theorem 1, which is restated below.

Theorem 1 (Confidence Ellipsoid). *Let $x \in \mathcal{Q}$, $\lambda > 1$, and $t \geq 1$. Then, with probability at least $1 - \delta$,*

$$|\mu_{t-1}(x) - \rho_m f(x)| \leq \nu_t \sigma_{t-1}(x).$$

Proof. Firstly, we can decompose B_t as follows:

$$B_t = \sum_{s=1}^{t-1} \varphi(x_s) \tilde{y}_{s,t} = \sum_{s=1}^{t-1} \varphi(x_s) y_s \mathbb{1}\{d_s \leq m\} + \sum_{s=t-m}^{t-1} \varphi(x_s) y_s (\mathbb{1}\{d_s \leq t-s\} - \mathbb{1}\{d_s \leq m\}).$$

Next, we can plug in this decomposition of B_t to derive:

$$\begin{aligned} |\mu_{t-1}(x) - \rho_m f(x)| &= |\varphi(x)^\top V_t(\lambda)^{-1} B_t - \rho_m \varphi(x)^\top f| \\ &\leq \left| \varphi(x)^\top V_t(\lambda)^{-1} \left(\sum_{s=1}^{t-1} \varphi(x_s) y_s \mathbb{1}\{d_s \leq m\} \right) - \rho_m \varphi(x)^\top f \right| + \\ &\quad \left| \varphi(x)^\top V_t(\lambda)^{-1} \left(\sum_{s=t-m}^{t-1} \varphi(x_s) y_s (\mathbb{1}\{d_s \leq t-s\} - \mathbb{1}\{d_s \leq m\}) \right) \right|. \end{aligned} \quad (4)$$

in which the equality makes use of the expression for the GP posterior mean: $\mu_{t-1}(x) = \varphi(x)^\top V_t(\lambda)^{-1} B_t$, and the reproducing property: $f(x) = \varphi(x)^\top f$. Next, we will separately provide upper bounds on the two terms in Eq. (4).

Firstly, We upper-bound the second term in Eq. (4) as follows:

$$\begin{aligned} &\left| \varphi(x)^\top V_t(\lambda)^{-1} \left(\sum_{s=t-m}^{t-1} \varphi(x_s) y_s (\mathbb{1}\{d_s \leq t-s\} - \mathbb{1}\{d_s \leq m\}) \right) \right| \\ &\leq \|\varphi(x)\|_{V_t(\lambda)^{-1}} \left\| V_t(\lambda)^{-1} \left(\sum_{s=t-m}^{t-1} \varphi(x_s) y_s (\mathbb{1}\{d_s \leq t-s\} - \mathbb{1}\{d_s \leq m\}) \right) \right\|_{V_t(\lambda)} \\ &\leq \sigma_{t-1}(x) \frac{1}{\sqrt{\lambda}} \left\| \sum_{s=t-m}^{t-1} \varphi(x_s) y_s (\mathbb{1}\{d_s \leq t-s\} - \mathbb{1}\{d_s \leq m\}) \right\|_{V_t(\lambda)^{-1}} \\ &\leq \sigma_{t-1}(x) \sum_{s=t-m}^{t-1} \|\varphi(x_s) y_s (\mathbb{1}\{d_s \leq t-s\} - \mathbb{1}\{d_s \leq m\})\|_{V_t(\lambda)^{-1}} \\ &\leq \sigma_{t-1}(x) \mathcal{B}_y \sum_{s=t-m}^{t-1} \|\varphi(x_s)\|_{V_t(\lambda)^{-1}} \\ &= \sigma_{t-1}(x) \frac{\mathcal{B}_y}{\sqrt{\lambda}} \sum_{s=t-m}^{t-1} \sigma_{t-1}(x_s) \\ &\leq \mathcal{B}_y \sigma_{t-1}(x) \sum_{s=t-m}^{t-1} \sigma_{t-1}(x_s). \end{aligned} \quad (5)$$

The second inequality and the equality have made use of the expression for the GP posterior variance: $\sigma_{t-1}^2(x) = \lambda \|\varphi(x)\|_{V_t(\lambda)}^2$, the third and last inequalities has made use of the fact that $\frac{1}{\sqrt{\lambda}} < 1$ because $\lambda > 1$ (note that if $\lambda > 0$, then there will be an additional multiplier factor of $1/\lambda$), the fourth inequality follows since we have assumed that all observations are bounded: $|y_s| \leq \mathcal{B}_y, \forall s \geq 1$.

Next, define $\mu'_{t-1}(x)$ as the standard GP posterior mean *without censoring the observations* (i.e., replace every $\tilde{y}_{s,t}$ by y_s in Eq. (1)). Now the first term in Eq. (4) can be upper-bounded as:

$$\begin{aligned}
 & \left| \varphi(x)^\top V_t(\lambda)^{-1} \left(\sum_{s=1}^{t-1} \varphi(x_s) y_s \mathbb{1}\{d_s \leq m\} \right) - \rho_m \varphi(x)^\top f \right| \\
 &= \left| \varphi(x)^\top V_t(\lambda)^{-1} \left(\sum_{s=1}^{t-1} \varphi(x_s) y_s (\rho_m + \eta_s) \right) - \rho_m \varphi(x)^\top f \right| \\
 &\leq \left| \varphi(x)^\top V_t(\lambda)^{-1} \sum_{s=1}^{t-1} \varphi(x_s) y_s \eta_s \right| + \rho_m \left| \varphi(x)^\top V_t(\lambda)^{-1} \sum_{s=1}^{t-1} \varphi(x_s) y_s - \varphi(x)^\top f \right| \\
 &= \left| \varphi(x)^\top V_t(\lambda)^{-1} \sum_{s=1}^{t-1} \varphi(x_s) y_s \eta_s \right| + \rho_m \left| \mu'_{t-1}(x) - f(x) \right| \\
 &\leq \left| \varphi(x)^\top V_t(\lambda)^{-1} \sum_{s=1}^{t-1} \varphi(x_s) y_s \eta_s \right| + \left(\mathcal{B}_f + R \sqrt{2(\gamma_{t-1} + 1 + \log(2/\delta))} \right) \sigma_{t-1}(x).
 \end{aligned} \tag{6}$$

In the first equality, we have defined $\eta_s \triangleq \mathbb{1}\{d_s \leq m\} - \rho_m$. In the second equality, we have again used the expression for the GP posterior mean and the reproducing property for f . The last inequality follows from Theorem 2 of [Chowdhury & Gopalan \(2017\)](#), which holds with the probability of $\geq 1 - \delta/2$. We have also used the fact that $\rho_m \leq 1$ in the last inequality.

To bound the first term in Eq. (6), we firstly define $\eta'_s \triangleq y_s \eta_s$, which is by definition a \mathcal{B}_y -sub-Gaussian noise. Also define $\boldsymbol{\eta}'_{1:t-1} = [\eta'_s]_{s=1, \dots, t-1}$ which is a $(t-1) \times 1$ -dimensional vector and define $\Phi_{t-1} \triangleq [\varphi(x_1), \dots, \varphi(x_{t-1})]^\top$ which is a $(t-1) \times \infty$ -dimensional matrix. These definitions imply that $\mathbf{K}_{t-1} = \Phi_{t-1} \Phi_{t-1}^\top$ and $V_t(\lambda) = \Phi_{t-1}^\top \Phi_{t-1} + \lambda I$. Also define $\lambda = 1 + \eta^7$ with $\eta > 0$. Based on these definitions, we can upper-bound $\left\| \sum_{s=1}^{t-1} \varphi(x_s) y_s \eta_s \right\|_{V_t(\lambda)^{-1}}$ by

$$\begin{aligned}
 \left\| \sum_{s=1}^{t-1} \varphi(x_s) y_s \eta_s \right\|_{V_t(\lambda)^{-1}} &= \left\| \sum_{s=1}^{t-1} \varphi(x_s) \eta'_s \right\|_{V_t(\lambda)^{-1}} = \sqrt{\left(\sum_{s=1}^{t-1} \eta'_s \varphi(x_s)^\top \right) V_t(\lambda)^{-1} \left(\sum_{s=1}^{t-1} \eta'_s \varphi(x_s) \right)} \\
 &= \sqrt{\boldsymbol{\eta}'_{1:t-1}^\top \Phi_{t-1} (\Phi_{t-1}^\top \Phi_{t-1} + \lambda I)^{-1} \Phi_{t-1}^\top \boldsymbol{\eta}'_{1:t-1}} \\
 &= \sqrt{\boldsymbol{\eta}'_{1:t-1}^\top \Phi_{t-1} \Phi_{t-1}^\top (\Phi_{t-1} \Phi_{t-1}^\top + \lambda I)^{-1} \boldsymbol{\eta}'_{1:t-1}} \\
 &= \sqrt{\boldsymbol{\eta}'_{1:t-1}^\top \mathbf{K}_{t-1} (\mathbf{K}_{t-1} + (1 + \eta)I)^{-1} \boldsymbol{\eta}'_{1:t-1}} \\
 &\leq \sqrt{\boldsymbol{\eta}'_{1:t-1}^\top (\mathbf{K}_{t-1} + \eta I) (\mathbf{K}_{t-1} + (1 + \eta)I)^{-1} \boldsymbol{\eta}'_{1:t-1}} \\
 &= \sqrt{\boldsymbol{\eta}'_{1:t-1}^\top ((\mathbf{K}_{t-1} + \eta I)^{-1} + I)^{-1} \boldsymbol{\eta}'_{1:t-1}} \\
 &= \left\| \boldsymbol{\eta}'_{1:t-1} \right\|_{((\mathbf{K}_{t-1} + \eta I)^{-1} + I)^{-1}}.
 \end{aligned} \tag{7}$$

In the equality in the third line, we have made use of the following matrix equality: $(\Phi_{t-1}^\top \Phi_{t-1} + \lambda I)^{-1} \Phi_{t-1}^\top = \Phi_{t-1}^\top (\Phi_{t-1} \Phi_{t-1}^\top + \lambda I)^{-1}$, and in the second last equality, we have used the matrix equality of $(\mathbf{K}_{t-1} + \eta I)(\mathbf{K}_{t-1} + \eta I + I)^{-1} = ((\mathbf{K}_{t-1} + \eta I)^{-1} + I)^{-1}$. Next, making use of Theorem 1 of [Chowdhury & Gopalan \(2017\)](#) and following

⁷The value of λ is set as $(1 + 2/T)$ ([Chowdhury & Gopalan, 2017](#)). However, one can use a Doubling Trick ([Besson & Kaufmann, 2018](#)) for unknown T .

similar steps as those in the proof there, we have that with probability of $\geq 1 - \delta/2$,

$$\begin{aligned} \|\boldsymbol{\eta}'_{1:t-1}\|_{((\mathbf{K}_{t-1} + \eta I)^{-1} + I)^{-1}} &\leq \mathcal{B}_y \sqrt{2 \log \frac{\sqrt{\det((1 + \eta)I + \mathbf{K}_{t-1})}}{\delta/2}} \\ &\leq \mathcal{B}_y \sqrt{2(\gamma_{t-1} + 1 + \log(2/\delta))}. \end{aligned} \quad (8)$$

Next, combining Eq. (7) and Eq. (8) above allows us to upper-bound the first term in Eq. (6):

$$\begin{aligned} \left| \varphi(x)^\top V_t(\lambda)^{-1} \sum_{s=1}^{t-1} \varphi(x_s) y_s \eta_s \right| &\leq \|\varphi(x)\|_{V_t(\lambda)^{-1}} \left\| V_t(\lambda)^{-1} \sum_{s=1}^{t-1} \varphi(x_s) y_s \eta_s \right\|_{V_t(\lambda)} \\ &\leq \sigma_{t-1}(x) \frac{1}{\sqrt{\lambda}} \left\| \sum_{s=1}^{t-1} \varphi(x_s) y_s \eta_s \right\|_{V_t(\lambda)^{-1}} \\ &\leq \sigma_{t-1}(x) \mathcal{B}_y \sqrt{2(\gamma_{t-1} + 1 + \log(2/\delta))}, \end{aligned} \quad (9)$$

in which the second inequality has made use of the expression for the GP posterior variance: $\sigma_{t-1}^2(x) = \lambda \|\varphi(x)\|_{V_t(\lambda)^{-1}}^2$, and the last inequality follows from Eq. (7) and Eq. (8) (and hence holds with probability of $\geq 1 - \delta/2$), as well as the fact that $\lambda > 1$. Now, we can plug in Eq. (9) as an upper bound on the first term in Eq. (6), and then use the resulting Eq. (6) as an upper bound on the first term in Eq. (4). Combining this with the upper bound on the second term of Eq. (4) given by Eq. (5), we have that

$$\begin{aligned} |\mu_{t-1}(x) - \rho_m f(x)| &\leq \mathcal{B}_y \sigma_{t-1}(x) \sum_{s=t-m}^{t-1} \sigma_{t-1}(x_s) + \left(\mathcal{B}_f + R \sqrt{2(\gamma_{t-1} + 1 + \log(2/\delta))} \right) \sigma_{t-1}(x) + \\ &\quad \sigma_{t-1}(x) \mathcal{B}_y \sqrt{2(\gamma_{t-1} + 1 + \log(2/\delta))} \\ &= \sigma_{t-1}(x) \left(\mathcal{B}_y \sum_{s=t-m}^{t-1} \sigma_{t-1}(x_s) + \mathcal{B}_f + (R + \mathcal{B}_y) \sqrt{2(\gamma_{t-1} + 1 + \log(2/\delta))} \right) \\ &= \nu_t \sigma_{t-1}(x), \end{aligned}$$

in which we have plugged in the definition of ν_t in the last equality. This completes the proof. \square

A.2. Proof of Theorem 2

Theorem 2 (GP-UCB-SDF). Let $C_1 = \sqrt{2/\log(1 + \lambda^{-1})}$. Then, with probability at least $1 - \delta$,

$$\mathfrak{R}_T \leq \frac{2}{\rho_m} \left(C_1 \beta_T \sqrt{T \gamma_T} + m \mathcal{B}_y C_1^2 \gamma_T \right).$$

Proof. To begin with, we can upper-bound the instantaneous regret $r_t = f(x^*) - f(x_t)$ as

$$\begin{aligned} r_t &= f(x^*) - f(x_t) = \frac{1}{\rho_m} [\rho_m f(x^*) - \rho_m f(x_t)] \\ &\leq \frac{1}{\rho_m} [\mu_{t-1}(x^*) + \nu_t \sigma_{t-1}(x^*) - \rho_m f(x_t)] \\ &\leq \frac{1}{\rho_m} [\mu_{t-1}(x_t) + \nu_t \sigma_{t-1}(x_t) - \rho_m f(x_t)] \\ &\leq \frac{1}{\rho_m} 2\nu_t \sigma_{t-1}(x_t), \end{aligned} \quad (10)$$

in which the first and last inequalities have made use of Theorem 1, and the second inequality follows because x_t is selected using the UCB policy: $x_t = \operatorname{argmax}_{x \in \mathcal{Q}} (\mu_{t-1} + \nu_t \sigma_{t-1}(x))$ (2). Now the cumulative regret can be upper-bounded as

$$\mathfrak{R}_T = \sum_{t=1}^T r_t \leq \frac{2}{\rho_m} \sum_{t=1}^T \nu_t \sigma_{t-1}(x_t) = \frac{2}{\rho_m} \sum_{t=1}^T \sigma_{t-1}(x_t) \beta_t + \frac{2}{\rho_m} \sum_{t=1}^T \sigma_{t-1}(x_t) \left(\mathcal{B}_y \sum_{s=t-m}^{t-1} \sigma_{t-1}(x_s) \right), \quad (11)$$

where we have plugged in the expression of ν_t in the equality.

Next, we use $I(\mathbf{y}_{1:T}; \mathbf{f}_{1:T})$ to denote the information gain from the noisy observations in the first T iterations about the objective function f . The first term in Eq. (11) can be upper-bounded as:

$$\begin{aligned}
 \frac{2}{\rho_m} \sum_{t=1}^T \sigma_{t-1}(x_t) \beta_t &\leq \frac{2}{\rho_m} \beta_T \sum_{t=1}^T \sigma_{t-1}(x_t) \\
 &\leq \frac{2}{\rho_m} \beta_T \sqrt{T} \sqrt{\sum_{t=1}^T \sigma_{t-1}^2(x_t)} \\
 &\leq \frac{2}{\rho_m} \beta_T \sqrt{T} \sqrt{\sum_{t=1}^T \frac{1}{\log(1+\lambda^{-1})} \log(1+\lambda^{-1} \sigma_{t-1}^2(x_t))} \\
 &\leq \frac{2}{\rho_m} \sqrt{\frac{1}{\log(1+\lambda^{-1})}} \beta_T \sqrt{T} \sqrt{2 \sum_{t=1}^T \log(1+\lambda^{-1} \sigma_{t-1}^2(x_t))} \\
 &\leq \frac{2}{\rho_m} \sqrt{\frac{1}{\log(1+\lambda^{-1})}} \beta_T \sqrt{T} \sqrt{2I(\mathbf{y}_{1:T}; \mathbf{f})} \\
 &\leq \frac{2}{\rho_m} \sqrt{\frac{2}{\log(1+\lambda^{-1})}} \beta_T \sqrt{T} \sqrt{\gamma_T} \\
 &\leq \frac{2}{\rho_m} C_1 \beta_T \sqrt{T \gamma_T}.
 \end{aligned} \tag{12}$$

The first inequality follows since β_t is monotonically increasing in t , the inequality in the second line follows because $\lambda^{-1}a \leq \frac{\lambda^{-1}}{\log(1+\lambda^{-1})} \log(1+\lambda^{-1}a)$ for $a \in (0, 1)$ (substitute $a = \sigma_{t-1}^2(x_t)$), and the inequality in the fourth line follows from plugging in the expression for the information gain: $I(\mathbf{y}_{1:T}; \mathbf{f}) = \frac{1}{2} \sum_{t=1}^T \log(1+\lambda^{-1} \sigma_{t-1}^2(x_t))$ (Lemma 5.3 of [Srinivas et al. \(2010\)](#)).

Next, we can derive an upper bound on the second term in Eq. (11):

$$\begin{aligned}
 \frac{2}{\rho_m} \sum_{t=1}^T \sigma_{t-1}(x_t) \left(\mathcal{B}_y \sum_{s=t-m}^{t-1} \sigma_{t-1}(x_s) \right) &= \frac{2}{\rho_m} \mathcal{B}_y \sum_{t=1}^T \sum_{s=t-m}^{t-1} \sigma_{t-1}(x_t) \sigma_{t-1}(x_s) \\
 &\leq \frac{1}{\rho_m} \mathcal{B}_y \sum_{t=1}^T \sum_{s=t-m}^{t-1} (\sigma_{t-1}^2(x_t) + \sigma_{t-1}^2(x_s)) \\
 &\leq \frac{1}{\rho_m} \mathcal{B}_y \sum_{t=1}^T \sum_{s=t-m}^{t-1} (\sigma_{t-1}^2(x_t) + \sigma_{s-1}^2(x_s)) \\
 &\leq \frac{2m}{\rho_m} \mathcal{B}_y \sum_{t=1}^T \sigma_{t-1}^2(x_t) \\
 &\leq \frac{2m}{\rho_m} \mathcal{B}_y C_1^2 \gamma_T.
 \end{aligned} \tag{13}$$

The first inequality has made use of the inequality of $ab \leq (a^2 + b^2)/2$, the second inequality follows since $\sigma_{t-1}^2(x_s) \leq \sigma_{s-1}^2(x_s)$ which is because $\sigma_{t-1}^2(x_s)$ is calculated by conditioning on more input locations than $\sigma_{s-1}^2(x_s)$, and the last inequality follows easily from some of the intermediate steps in the derivations of Eq. (12).

Lastly, we can plug in Eq. (12) and Eq. (13) as upper bounds on the two terms in Eq. (11), to obtain:

$$\mathfrak{R}_T \leq \frac{2}{\rho_m} C_1 \beta_T \sqrt{T \gamma_T} + \frac{2m}{\rho_m} \mathcal{B}_y C_1^2 \gamma_T = \frac{2}{\rho_m} \left(C_1 \beta_T \sqrt{T \gamma_T} + m \mathcal{B}_y C_1^2 \gamma_T \right).$$

It completes the proof. \square

A.3. Proof of Theorem 3

First we define $\beta_t \triangleq \mathcal{B}_f + (R + \mathcal{B}_y)\sqrt{2(\gamma_{t-1} + 1 + \log(4/\delta))}$, $\nu_t \triangleq \mathcal{B}_y \sum_{s=t-m}^{t-1} \sigma_{t-1}(x_s) + \beta_t$ and $c_t \triangleq \nu_t(1 + \sqrt{2 \log(|\mathcal{Q}|t^2)})$. Note that we have replaced the δ in the definition of β_t (Theorem 1) by $\delta/2$. We use \mathcal{F}_{t-1} to denote the filtration containing the history of selected inputs and observed outputs up to iteration $t-1$. To begin with, we define two events $E^f(t)$ and $E^{f_t}(t)$ through the following two lemmas.

Lemma 1. *Let $\delta \in (0, 1)$. Define $E^f(t)$ as the event that $|\mu_{t-1}(x) - \rho_m f(x)| \leq \nu_t \sigma_{t-1}(x)$ for all $x \in \mathcal{Q}$. We have that $\mathbb{P}[E^f(t)] \geq 1 - \delta/2$ for all $t \geq 1$.*

Note that Lemma 1 is the same as Theorem 1 except that we have replaced the error probability of δ by $\delta/2$.

Lemma 2. *Define $E^{f_t}(t)$ as the event that $|f_t(x) - \mu_{t-1}(x)| \leq \nu_t \sqrt{2 \log(|\mathcal{Q}|t^2)} \sigma_{t-1}(x)$. We have that $\mathbb{P}[E^{f_t}(t) | \mathcal{F}_{t-1}] \geq 1 - 1/t^2$ for any possible filtration \mathcal{F}_{t-1} .*

Lemma 2 makes use of the concentration of a random variable sampled from a Gaussian distribution, and its proof follows from Lemma 5 of [Chowdhury & Gopalan \(2017\)](#). Next, we define a set of *saturated points* in every iteration, which intuitively represents those inputs that lead to large regrets in an iteration.

Definition 1. *In iteration t , define the set of saturated points as*

$$S_t = \{x \in \mathcal{Q} : \rho_m \Delta(x) > c_t \sigma_{t-1}(x)\},$$

where $\Delta(x) = f(x^*) - f(x)$ and $x^* \in \arg \max_{x \in \mathcal{Q}} f(x)$.

Based on this definition, we will later prove that we can get a lower bound on the probability that the selected input in iteration t is unsaturated (Lemma 4). To do that, we first need the following auxiliary lemma.

Lemma 3. *For any filtration \mathcal{F}_{t-1} , conditioned on the event $E^f(t)$, we have that $\forall x \in \mathcal{Q}$,*

$$\mathbb{P}(f_t(x) > \rho_m f(x) | \mathcal{F}_{t-1}) \geq p, \quad (14)$$

where $p = \frac{1}{4e\sqrt{\pi}}$.

Proof. Adding and subtracting $\frac{\mu_{t-1}(x)}{\nu_t \sigma_{t-1}(x)}$ both sides of $\mathbb{P}(f_t(x) > \rho_m f(x) | \mathcal{F}_{t-1})$, we get

$$\begin{aligned} \mathbb{P}(f_t(x) > \rho_m f(x) | \mathcal{F}_{t-1}) &= \mathbb{P}\left(\frac{f_t(x) - \mu_{t-1}(x)}{\nu_t \sigma_{t-1}(x)} > \frac{\rho_m f(x) - \mu_{t-1}(x)}{\nu_t \sigma_{t-1}(x)} \middle| \mathcal{F}_{t-1}\right) \\ &\geq \mathbb{P}\left(\frac{f_t(x) - \mu_{t-1}(x)}{\nu_t \sigma_{t-1}(x)} > \frac{|\rho_m f(x) - \mu_{t-1}(x)|}{\nu_t \sigma_{t-1}(x)} \middle| \mathcal{F}_{t-1}\right) \\ &\geq \mathbb{P}\left(\frac{f_t(x) - \mu_{t-1}(x)}{\nu_t \sigma_{t-1}(x)} > 1 \middle| \mathcal{F}_{t-1}\right) \\ &\geq \frac{1}{4e\sqrt{\pi}}, \end{aligned} \quad (15)$$

in which the second inequality makes use of Lemma 1 (note that we have conditioned on the event $E^f(t)$ here), and the last inequality follows from the Gaussian anti-concentration inequality: $\mathbb{P}(z > a) \geq \exp(-a^2)/(4\sqrt{\pi}a)$ where $z \sim \mathcal{N}(0, 1)$. \square

The next lemma proves a lower bound on the probability that the selected input x_t is unsaturated.

Lemma 4. *For any filtration \mathcal{F}_{t-1} , conditioned on the event $E^f(t)$, we have that,*

$$\mathbb{P}(x_t \in \mathcal{Q} \setminus S_t | \mathcal{F}_{t-1}) \geq p - 1/t^2.$$

Proof. To begin with, we have that

$$\mathbb{P}(x_t \in \mathcal{Q} \setminus S_t | \mathcal{F}_{t-1}) \geq \mathbb{P}(f_t(x^*) > f_t(x), \forall x \in S_t | \mathcal{F}_{t-1}). \quad (16)$$

This inequality can be justified because the event on the right hand side implies the event on the left hand side. Specifically, according to Definition 1, x^* is always unsaturated because $\Delta(x^*) = 0 < c_t \sigma_{t-1}(x^*)$. Therefore, because x_t is selected by $x_t = \arg \max_{x \in \mathcal{Q}} f_t(x)$, we have that if $f_t(x^*) > f_t(x), \forall x \in S_t$, then the selected x_t is guaranteed to be unsaturated. Now conditioning on both events $E^f(t)$ and $E^{f_t}(t)$, for all $x \in S_t$, we have that

$$f_t(x) \leq \rho_m f(x) + c_t \sigma_{t-1}(x) \leq \rho_m f(x) + \rho_m \Delta(x) = \rho_m f(x) + \rho_m f(x^*) - \rho_m f(x) = \rho_m f(x^*), \quad (17)$$

in which the first inequality follows from Lemma 1 and Lemma 2 and the second inequality makes use of Definition 1. Next, separately considering the cases where the event $E^{f_t}(t)$ holds or not and making use of Eq. (16) and Eq. (17), we have that

$$\begin{aligned} \mathbb{P}(x_t \in \mathcal{Q} \setminus S_t | \mathcal{F}_{t-1}) &\geq \mathbb{P}(f_t(x^*) > f_t(x), \forall x \in S_t | \mathcal{F}_{t-1}) \\ &\geq \mathbb{P}(f_t(x^*) > f(x^*) | \mathcal{F}_{t-1}) - \mathbb{P}(\overline{E^{f_t}(t)} | \mathcal{F}_{t-1}) \\ &\geq p - 1/t^2, \end{aligned} \quad (18)$$

in which the last inequality has made use of Lemma 2 and Lemma 3. \square

Next, we use the following lemma to derive an upper bound on the expected instantaneous regret.

Lemma 5. *For any filtration \mathcal{F}_{t-1} , conditioned on the event $E^f(t)$, we have that,*

$$\mathbb{E}[r_t | \mathcal{F}_{t-1}] \leq \frac{11c_t}{\rho_m p} \mathbb{E}[\sigma_{t-1}(x_t) | \mathcal{F}_{t-1}] + \frac{2\mathcal{B}_f}{t^2}.$$

Proof. To begin with, define \bar{x}_t as the unsaturated input with the smallest GP posterior standard deviation:

$$\bar{x}_t = \arg \min_{x \in \mathcal{Q} \setminus S_t} \sigma_{t-1}(x). \quad (19)$$

This definition gives us:

$$\mathbb{E}[\sigma_{t-1}(x_t) | \mathcal{F}_{t-1}] \geq \mathbb{E}[\sigma_{t-1}(x_t) | \mathcal{F}_{t-1}, x_t \in \mathcal{Q} \setminus S_t] \mathbb{P}(x_t \in \mathcal{Q} \setminus S_t | \mathcal{F}_{t-1}) \geq \sigma_{t-1}(\bar{x}_t)(p - 1/t^2), \quad (20)$$

in which the second inequality makes use of Lemma 4, as well as the definition of \bar{x}_t .

Next, conditioned on both events $E^f(t)$ and $E^{f_t}(t)$, we can upper-bound the instantaneous regret as

$$\begin{aligned} r_t = f(x^*) - f(x_t) &= \frac{1}{\rho_m} [\rho_m f(x^*) - \rho_m f(\bar{x}_t) + \rho_m f(\bar{x}_t) - \rho_m f(x_t)] \\ &\leq \frac{1}{\rho_m} [\rho_m \Delta(\bar{x}_t) + f_t(\bar{x}_t) + c_t \sigma_{t-1}(\bar{x}_t) - f_t(x_t) + c_t \sigma_{t-1}(x_t)] \\ &\leq \frac{1}{\rho_m} [c_t \sigma_{t-1}(\bar{x}_t) + f_t(\bar{x}_t) + c_t \sigma_{t-1}(\bar{x}_t) - f_t(x_t) + c_t \sigma_{t-1}(x_t)] \\ &\leq \frac{1}{\rho_m} [2c_t \sigma_{t-1}(\bar{x}_t) + c_t \sigma_{t-1}(x_t)], \end{aligned}$$

in which the first inequality follows from Lemma 1 and Lemma 2 as well as the definition of $\Delta(\cdot)$ (Definition 1), the second inequality follows because \bar{x}_t is unsaturated, and the last inequality follows because $f_t(\bar{x}_t) - f_t(x_t) \leq 0$ since $x_t = \arg \max_{x \in \mathcal{Q}} f_t(x)$. Next, by separately considering the cases where the event $E^{f_t}(t)$ holds and otherwise, we are ready to upper-bound the expected instantaneous regret:

$$\begin{aligned} \mathbb{E}[r_t | \mathcal{F}_{t-1}] &\leq \frac{1}{\rho_m} \mathbb{E}[2c_t \sigma_{t-1}(\bar{x}_t) + c_t \sigma_{t-1}(x_t) | \mathcal{F}_{t-1}] + \frac{2\mathcal{B}_f}{t^2} \\ &\leq \frac{1}{\rho_m} \mathbb{E} \left[2c_t \frac{\sigma_{t-1}(x_t)}{p - 1/t^2} + c_t \sigma_{t-1}(x_t) | \mathcal{F}_{t-1} \right] + \frac{2\mathcal{B}_f}{t^2} \\ &\leq \frac{c_t}{\rho_m} \left(\frac{2}{p - 1/t^2} + 1 \right) \mathbb{E}[\sigma_{t-1}(x_t) | \mathcal{F}_{t-1}] + \frac{2\mathcal{B}_f}{t^2} \\ &\leq \frac{11c_t}{\rho_m p} \mathbb{E}[\sigma_{t-1}(x_t) | \mathcal{F}_{t-1}] + \frac{2\mathcal{B}_f}{t^2}, \end{aligned} \quad (21)$$

in which the second inequality results from Eq. (20), and the last inequality follows because $\frac{2}{p-1/t^2} \leq 10/p$ and $1 \leq 1/p$. \square

Next, we define the following stochastic process $(Y_t : t = 0, \dots, T)$, which we prove is a super-martingale in the subsequent lemma by making use of Lemma 5.

Definition 2. Define $Y_0 = 0$, and for all $t = 1, \dots, T$,

$$\bar{r}_t = r_t \mathbb{I}\{E^f(t)\}, \quad X_t = \bar{r}_t - \frac{11c_t}{\rho_m p} \sigma_{t-1}(x_t) - \frac{2\mathcal{B}_f}{t^2}, \quad \text{and} \quad Y_t = \sum_{s=1}^t X_s.$$

Lemma 6. $(Y_t : t = 0, \dots, T)$ is a super-martingale with respect to the filtration \mathcal{F}_t .

Proof. As $X_t = Y_t - Y_{t-1}$, we have

$$\begin{aligned} \mathbb{E}[Y_t - Y_{t-1} | \mathcal{F}_{t-1}] &= \mathbb{E}[X_t | \mathcal{F}_{t-1}] \\ &= \mathbb{E}\left[\bar{r}_t - \frac{11c_t}{\rho_m p} \sigma_{t-1}(x_t) - \frac{2\mathcal{B}_f}{t^2} \middle| \mathcal{F}_{t-1}\right] \\ &= \mathbb{E}[\bar{r}_t | \mathcal{F}_{t-1}] - \left[\frac{11c_t}{\rho_m p} \mathbb{E}[\sigma_{t-1}(x_t) | \mathcal{F}_{t-1}] + \frac{2\mathcal{B}_f}{t^2}\right] \\ &\leq 0. \end{aligned}$$

When the event $E^f(t)$ holds, the last inequality follows from Lemma 5; when $E^f(t)$ is false, $\bar{r}_t = 0$ and hence the inequality trivially holds. \square

Lastly, we are ready to prove the upper bound of **GP-TS-SDF** on the cumulative regret \mathfrak{R}_T by applying the Azuma-Hoeffding Inequality to the stochastic process defined above.

Theorem 3 (GP-TS-SDF). With probability at least $1 - \delta$,

$$\mathfrak{R}_T = \tilde{O}\left(\frac{1}{\rho_m} \left(\sqrt{T\gamma_T}(\sqrt{\gamma_T} + 1) + m(\gamma_T + \sqrt{T})\right)\right).$$

Proof. To begin with, we derive an upper bound on $|Y_t - Y_{t-1}|$:

$$\begin{aligned} |Y_t - Y_{t-1}| &= |X_t| \leq |\bar{r}_t| + \frac{11c_t}{\rho_m p} \sigma_{t-1}(x_t) + \frac{2\mathcal{B}_f}{t^2} \\ &\leq 2\mathcal{B}_f + \frac{11c_t}{\rho_m p} + 2\mathcal{B}_f \\ &\leq \frac{1}{\rho_m p} (4\mathcal{B}_f + 11c_t), \end{aligned} \tag{22}$$

where the second inequality follows because $\sigma_{t-1}(x_t) \leq 1$, and the last inequality follows since $\frac{1}{\rho_m p} \geq 1$. Now we are ready to apply the Azuma-Hoeffding Inequality to $(Y_t : t = 0, \dots, T)$ with an error probability of $\delta/2$:

$$\begin{aligned} \sum_{t=1}^T \bar{r}_t &\leq \sum_{t=1}^T \frac{11c_t}{\rho_m p} \sigma_{t-1}(x_t) + \sum_{t=1}^T \frac{2\mathcal{B}_f}{t^2} + \sqrt{2 \log(2/\delta) \sum_{t=1}^T \left(\frac{1}{\rho_m p} (4\mathcal{B}_f + 11c_t)\right)^2} \\ &\leq \frac{11}{\rho_m p} (1 + \sqrt{2 \log(|\mathcal{Q}|T^2)}) \sum_{t=1}^T \nu_t \sigma_{t-1}(x_t) + \frac{B\pi^2}{3} + \frac{4\mathcal{B}_f + 11c_T}{\rho_m p} \sqrt{2T \log(2/\delta)} \\ &\leq \frac{11}{\rho_m p} (1 + \sqrt{2 \log(|\mathcal{Q}|T^2)}) \left(C_1 \beta_T \sqrt{T\gamma_T} + m\mathcal{B}_y C_1^2 \gamma_T\right) + \frac{4\mathcal{B}_f + 11c_T}{\rho_m p} \sqrt{2T \log(2/\delta)} + \frac{B\pi^2}{3}, \end{aligned} \tag{23}$$

in which the second inequality makes use of the fact that c_t is monotonically increasing and that $\sum_{t=1}^T 1/t^2 \leq \pi^2/6$, and the last inequality follows from the proof of Theorem 2 (Appendix A.2) which gives an upper bound on $\sum_{t=1}^T \nu_t \sigma_{t-1}(x_t)$. Note that Eq. (23) holds with probability $\geq 1 - \delta/2$. Also note that $\bar{r}_t = r_t$ with probability of $\geq 1 - \delta/2$ because the event

$E^f(t)$ holds with probability of $\geq 1 - \delta/2$ (Lemma 1), therefore, the upper bound from Eq. (23) is an upper bound on $\mathfrak{R}_T = \sum_{t=1}^T r_t$ with probability of $1 - \delta$.

Lastly, we can simplify the regret upper bound from Eq. (23) into asymptotic notation. Note that $c_t = \nu_t(1 + \sqrt{2 \log(|\mathcal{Q}|t^2)})$, $\nu_t = \mathcal{B}_y \sum_{s=t-m}^{t-1} \sigma_{t-1}(x_s) + \beta_t$ and $\beta_t = \mathcal{B}_f + (R + \mathcal{B}_y) \sqrt{2(\gamma_{t-1} + 1 + \log(4/\delta))}$. To simplify notations, we analyze the asymptotic dependence while ignoring all log factors, and ignore the dependence on constants including B , \mathcal{B}_y and R . This gives us $\beta_T = \tilde{O}(\sqrt{\gamma_T})$ and $c_T = \tilde{O}(m + \sqrt{\gamma_T})$. As a result, the regret upper bound can be simplified as

$$\mathfrak{R}_T = \tilde{O} \left(\frac{1}{\rho_m} (\gamma_T \sqrt{T} + m \gamma_T) + \frac{1}{\rho_m} (m + \sqrt{\gamma_T}) \sqrt{T} \right) = \tilde{O} \left(\frac{1}{\rho_m} (\sqrt{T} \gamma_T (\sqrt{\gamma_T} + 1) + m(\gamma_T + \sqrt{T})) \right).$$

It completes the proof. \square

A.4. Time based Censored Feedback

Let m be the amount of time the learner waits for the delayed feedback, $\bar{\tau}_t$ be the time of initiating the selection process for the t -th query, and τ_t be the time of starting the t -th query. The censored feedback of s -th query before selecting t -th query is denoted by $\tilde{y}_{s,t}$, where $\tilde{y}_{s,t} = y_s \mathbb{1}\{d_s \leq \min(m, \bar{\tau}_t - \tau_s)\}$. Now, the GP posterior mean and covariance functions can be expressed using available information of function queries and their censored feedback as follows:

$$\begin{aligned} \mu_{t-1}(x) &= \mathbf{k}_{t-1}(x)^\top \mathbf{K}_{t,\lambda}^{-1} \tilde{\mathbf{y}}_{t-1}, \\ \sigma_{t-1}^2(x, x') &= k(x, x') - \mathbf{k}_{t-1}(x)^\top \mathbf{K}_{t,\lambda}^{-1} \mathbf{k}_{t-1}(x'), \end{aligned} \quad (24)$$

where $\mathbf{K}_{t,\lambda} = \mathbf{K}_{t-1} + \lambda \mathbf{I}$, $\mathbf{k}_{t-1}(x) = (k(x, x_{t'}))_{\tau_{t'} \in [0, \bar{\tau}_t]}^\top$, $\tilde{\mathbf{y}}_{t-1} = (\tilde{y}_{s,t})_{\tau_s \in [0, \bar{\tau}_t]}^\top$, $\mathbf{K}_{t-1} = (k(x_{t'}, x_{t''}))_{\tau_{t'}, \tau_{t''} \in [0, \bar{\tau}_t]}$ and λ is a regularization parameter to ensures $\mathbf{K}_{t,\lambda}$ is an invertible matrix. Our current regret analysis will work for this setting as well after appropriately changing the used variables.

B. Theoretical Analysis of Contextual Gaussian Process Bandit

Our regret upper bounds in Theorem 2 and Theorem 3 are also applicable in the contextual GP bandit setting after some slight modifications to their proofs. Note that in the contextual GP bandit setting (Section 5), the only major difference with the BO setting is that now the objective function g depends on both a context vector $z_t \in \mathcal{Z}$ and an input vector $x_t \in \mathcal{Q}$, in which every z_t is generated by the environment and x_t is selected by our GP-UCB-SDF or GP-TS-SDF algorithms. Therefore, the most important modifications to the proofs of Theorem 2 and Theorem 3 involve replacing the original input space of \mathcal{Q} by the new input space of $\mathcal{Z} \times \mathcal{Q}$ and accounting for the modified definition of regret in Section 5.

To begin with, with the same definition of ν_t , it is easy to verify that Theorem 1 can be extended to the contextual GP bandit setting:

Theorem 4 (Confidence Ellipsoid for Contextual GP Bandit). *With probability at least $1 - \delta$,*

$$|\mu_{t-1}(z, x) - \rho_m g(z, x)| \leq \nu_t \sigma_{t-1}(z, x), \forall z \in \mathcal{Z}, x \in \mathcal{Q}.$$

The proof of Theorem 4 is the same as that of Theorem 1 except that the input x in the proof of Theorem 1 is replaced by (z, x) .

B.1. GP-UCB-SDF for Contextual Gaussian Process Bandit

For the GP-UCB-SDF algorithm in the contextual GP bandit setting, the instantaneous regret r_t can be analyzed in a similar way to Eq. (10) in the proof of Theorem 2 (Appendix A.2):

$$\begin{aligned} r_t &= g(z_t, x_t^*) - g(z_t, x_t) = \frac{1}{\rho_m} [\rho_m g(z_t, x_t^*) - \rho_m g(z_t, x_t)] \\ &\leq \frac{1}{\rho_m} [\mu_{t-1}(z_t, x_t^*) + \nu_t \sigma_{t-1}(z_t, x_t^*) - \rho_m g(z_t, x_t)] \\ &\leq \frac{1}{\rho_m} [\mu_{t-1}(z_t, x_t) + \nu_t \sigma_{t-1}(z_t, x_t) - \rho_m g(z_t, x_t)] \end{aligned}$$

$$\leq \frac{1}{\rho_m} 2\nu_t \sigma_{t-1}(z_t, x_t),$$

in which the first and last inequalities both make use of Theorem 4, and the second inequality follows from the way in which x_t is selected: $x_t = \arg \max_{x \in \mathcal{Q}} \mu_{t-1}(z_t, x) + \nu_t \sigma_{t-1}(z_t, x)$. Following this, the subsequent steps in the proof in Appendix A.2 immediately follow, hence producing the same regret upper bound as Theorem 2.

B.2. GP-TS-SDF for Contextual Gaussian Process Bandit

To begin with, the events $E^f(t)$ (Lemma 1) and $E^{f_t}(t)$ (Lemma 2) can be similarly defined:

Lemma 7. *Let $\delta \in (0, 1)$. Define $E^f(t)$ as the event that $|\mu_{t-1}(z, x) - \rho_m g(z, x)| \leq \nu_t \sigma_{t-1}(z, x)$ for all $z \in \mathcal{Z}, x \in \mathcal{Q}$. We have that $\mathbb{P}[E^f(t)] \geq 1 - \delta/2$ for all $t \geq 1$.*

Lemma 8. *Define $E^{f_t}(t)$ as the event that $|g_t(z, x) - \mu_{t-1}(z, x)| \leq \nu_t \sqrt{2 \log(|\mathcal{Z}||\mathcal{Q}|t^2)} \sigma_{t-1}(z, x)$. We have that $\mathbb{P}[E^{f_t}(t) | \mathcal{F}_{t-1}] \geq 1 - 1/t^2$ for any possible filtration \mathcal{F}_{t-1} .*

The validity of Lemma 7 follows immediately from Theorem 4 (after replacing the error probability of δ by $\delta/2$ in the definition of ν_t , which we have omitted here for simplicity), and the proof of Lemma 8 is the same as that of Lemma 2. Next, we modify the definition of saturated points from Definition 1:

Definition 3. *In iteration t , given a context vector z_t , define the set of saturated points as*

$$S_t = \{x \in \mathcal{Q} : \rho_m \Delta(z_t, x) > c_t \sigma_{t-1}(z_t, x)\},$$

where $\Delta(z_t, x) = g(z_t, x_t^*) - g(z_t, x)$ and $x_t^* = \arg \max_{x \in \mathcal{Q}} g(z_t, x)$.

Note that according to this definition, x_t^* is always unsaturated. The next lemma is a counterpart of Lemma 3, and the proofs are the same.

Lemma 9. *For any filtration \mathcal{F}_{t-1} , conditioned on the event $E^f(t)$, we have that $\forall z \in \mathcal{Z}, x \in \mathcal{Q}$,*

$$\mathbb{P}(g_t(z, x) > \rho_m g(z, x) | \mathcal{F}_{t-1}) \geq p,$$

where $p = \frac{1}{4e\sqrt{\pi}}$.

The next lemma, similar to Lemma 4, shows that in contextual GP bandit problems, the probability that the selected x_t is unsaturated can also be lower-bounded.

Lemma 10. *For any filtration \mathcal{F}_{t-1} , given a context vector z_t , conditioned on the event $E^f(t)$, we have that,*

$$\mathbb{P}(x_t \in \mathcal{Q} \setminus S_t | \mathcal{F}_{t-1}) \geq p - 1/t^2.$$

Proof. To begin with, we have that

$$\mathbb{P}(x_t \in \mathcal{Q} \setminus S_t | \mathcal{F}_{t-1}) \geq \mathbb{P}(g_t(z_t, x_t^*) > g_t(z_t, x), \forall x \in S_t | \mathcal{F}_{t-1}). \quad (25)$$

The validity of this equation can be seen following the same analysis of Eq. (16), i.e., the event on the right hand side implies the event on the left hand side. Now conditioning on both events $E^f(t)$ and $E^{f_t}(t)$, for all $x \in S_t$, we have that

$$\begin{aligned} g_t(z_t, x) &\leq \rho_m g(z_t, x) + c_t \sigma_{t-1}(z_t, x) \leq \rho_m g(z_t, x) + \rho_m \Delta(z_t, x) \\ &= \rho_m g(z_t, x) + \rho_m g(z_t, x_t^*) - \rho_m g(z_t, x) = \rho_m g(z_t, x_t^*), \end{aligned} \quad (26)$$

in which the first inequality follows from Lemma 7 and Lemma 8 and the second inequality makes use of Definition 3.

Next, separately considering the cases where the event $E^{f_t}(t)$ holds or not and making use of Eq. (25) and Eq. (26), we have that

$$\begin{aligned} \mathbb{P}(x_t \in \mathcal{Q} \setminus S_t | \mathcal{F}_{t-1}) &\geq \mathbb{P}(g_t(z_t, x_t^*) > g_t(z_t, x), \forall x \in S_t | \mathcal{F}_{t-1}) \\ &\geq \mathbb{P}(g_t(z_t, x_t^*) > \rho_m g(z_t, x_t^*) | \mathcal{F}_{t-1}) - \mathbb{P}(\overline{E^{f_t}(t)} | \mathcal{F}_{t-1}) \\ &\geq p - 1/t^2, \end{aligned}$$

where the last inequality has made use of Lemma 8 and Lemma 9. \square

Lemma 11. For any filtration \mathcal{F}_{t-1} , given a context vector z_t , conditioned on the event $E^f(t)$, we have that,

$$\mathbb{E}[r_t | \mathcal{F}_{t-1}] \leq \frac{11c_t}{\rho_m p} \mathbb{E}[\sigma_{t-1}(z_t, x_t) | \mathcal{F}_{t-1}] + \frac{2\mathcal{B}_f}{t^2}.$$

Proof. To begin with, given a context vector z_t , define \bar{x}_t as the unsaturated input with the smallest GP posterior standard deviation:

$$\bar{x}_t = \arg \min_{x \in \mathcal{Q} \setminus S_t} \sigma_{t-1}(z_t, x). \quad (27)$$

This definition gives us:

$$\begin{aligned} \mathbb{E}[\sigma_{t-1}(z_t, x_t) | \mathcal{F}_{t-1}] &\geq \mathbb{E}[\sigma_{t-1}(z_t, x_t) | \mathcal{F}_{t-1}, x_t \in \mathcal{Q} \setminus S_t] \mathbb{P}(x_t \in \mathcal{Q} \setminus S_t | \mathcal{F}_{t-1}) \\ &\geq \sigma_{t-1}(z_t, \bar{x}_t)(p - 1/t^2), \end{aligned} \quad (28)$$

in which the second inequality makes use of Lemma 4, as well as the definition of \bar{x}_t . Next, conditioned on both events $E^f(t)$ and $E^{f^*}(t)$, we can upper-bound the instantaneous regret as

$$\begin{aligned} r_t &= g(z_t, x_t^*) - g(z_t, x_t) \\ &= \frac{1}{\rho_m} [\rho_m g(z_t, x_t^*) - \rho_m g(z_t, \bar{x}_t) + \rho_m g(z_t, \bar{x}_t) - \rho_m g(z_t, x_t)] \\ &\leq \frac{1}{\rho_m} [\rho_m \Delta(z_t, \bar{x}_t) + g_t(z_t, \bar{x}_t) + c_t \sigma_{t-1}(z_t, \bar{x}_t) - g_t(z_t, x_t) + c_t \sigma_{t-1}(z_t, x_t)] \\ &\leq \frac{1}{\rho_m} [c_t \sigma_{t-1}(z_t, \bar{x}_t) + g_t(z_t, \bar{x}_t) + c_t \sigma_{t-1}(z_t, \bar{x}_t) - g_t(z_t, x_t) + c_t \sigma_{t-1}(z_t, x_t)] \\ &\leq \frac{1}{\rho_m} [2c_t \sigma_{t-1}(z_t, \bar{x}_t) + c_t \sigma_{t-1}(z_t, x_t)], \end{aligned}$$

where the first inequality follows from Lemma 7 and Lemma 8 as well as the definition of $\Delta(\cdot)$ (Definition 3), the second inequality follows because \bar{x}_t is unsaturated, and the last inequality follows because $g_t(z_t, \bar{x}_t) - g_t(z_t, x_t) \leq 0$ since $x_t = \arg \max_{x \in \mathcal{Q}} g_t(z_t, x)$. Next, the proof in Eq. (21) can be immediately applied, hence producing the upper bound on the instantaneous regret in this lemma. \square

Now the remaining steps in the proof of Theorem 3 in Appendix A.3 immediately follow. Specifically, we can first define a stochastic process $(Y_t : t = 0, \dots, T)$ in the same way as Definition 2, and then use Lemma 6 to show that $(Y_t : t = 0, \dots, T)$ is super-martingale. Lastly, we can apply the Azuma Hoeffding Inequality to $(Y_t : t = 0, \dots, T)$ in the same way as Theorem 3, which completes the proof.

C. More Experimental Details

In our experiments, since it has been repeatedly observed that the theoretical value of β_t is overly conservative (Srinivas et al., 2010), we set β_t to be a constant ($\beta_t = 1$) for all methods. In all experiments and for all methods, we use the SE kernel for the GP and optimize the GP hyperparameters by maximizing the marginal likelihood after every 10 iterations. In all experiments where we report the simple regret (e.g., Fig. 1, Fig. 2, etc.), we calculate the simple regret in an iteration using only those function evaluations which have converted (i.e., we ignore all pending observations).

C.1. Synthetic Experiment

In the synthetic experiment, the objective function f is sampled from a GP using the SE kernel with a lengthscale of 0.02 defined on a 1-dimensional domain. The domain is an equally spaced grid within the range $[0, 1]$ with a size 1000. The sampled function is normalized into the range of $[0, 1]$.

C.2. Real-world Experiments

In the SVM hyperparameter tuning experiment, the diabetes diagnosis dataset can be found at <https://www.kaggle.com/uciml/pima-indians-diabetes-database>. We use 70% of the dataset as the training set and the remaining 30% as the validation set. For every evaluated hyperparameter configuration, we train the SVM using the




Figure 5: Performances of GP-TS-SDF and other TS-based baseline methods for (a) stochastic and (b) deterministic delay distributions in the synthetic experiment (Section 6.1).

training set with the particular hyperparameter configuration and then evaluate the learned SVM on the validation set, whose validation accuracy is reported as the observations. We tune the penalty parameter within the range of $[10^{-4}, 100]$ and the RBF kernel parameter within $[10^{-4}, 10]$.

In the experiment on hyperparameter tuning for CNN, we use the MNIST dataset and follow the default training-testing sets partitions given by the PyTorch package. The CNN consists of one convolutional layer (with 8 channels and convolutional kernels of size 3), followed by a max-pooling layer (with a pooling size of 3), and then followed by a fully connected layer (with 8 nodes). We use the ReLU activation function and the Adam optimizer. We tune the batch size (within $[128, 512]$), the learning rate (within $[10^{-6}, 1]$) and the learning rate decay (within $[10^{-6}, 1]$).

In the hyperparameter tuning experiment for LR, the breast cancer dataset we have adopted can be found at [https://archive.ics.uci.edu/ml/datasets/Breast+Cancer+Wisconsin+\(Diagnostic\)](https://archive.ics.uci.edu/ml/datasets/Breast+Cancer+Wisconsin+(Diagnostic)). Similar to the SVM experiment, we use 70% of the dataset as the training set and the remaining 30% as the testing set. Here we tune three hyperparameters of the LR model: the batch size (within $[32, 128]$), the learning rate (within $[10^{-6}, 1]$) and learning rate decay (within $[10^{-6}, 1]$).

C.3. Real-world Experiments on Contextual Gaussian Process Bandits

For the multi-task BO experiment, the tabular benchmark on hyperparameter tuning of SVM is introduced by the work of (Wistuba et al., 2015) and can be found at <https://github.com/wistuba/TST>. The dataset consists of 50 tasks, and each task corresponds to a different classification problem with a different dataset. The domain of hyperparameter configurations in this benchmark (which is shared by all 50 tasks) consists of discrete values for six hyperparameters: 3 binary parameters indicating (via one-hot encoding) whether the linear, polynomial, or RBF kernel is used, and the penalty parameter, the degree for the polynomial kernel, and the RBF kernel parameter. The size of the domain is 288. As a result, for each task, the validation accuracy for each one of the 288 hyperparameter configurations is recorded as the observation. As we have mentioned in the main text (Section 6.2), each of the 50 tasks is associated with some meta-features, and here we use the first six meta-features as the contexts. Each of the 50 tasks is associated with a 6-dimensional context vector z_t , which is used to characterize this particular task. For the non-stationary BO task, the dataset and the other experimental settings (e.g., the tuned hyperparameters, the range of the hyperparameters, etc.) are the same as the SVM hyperparameter tuning experiment in Section 6.2. Refer to Appendix C.2 for more details.