

---

# Active Cost-aware Labeling of Streaming Data

---

Ting Cai

University of Wisconsin-Madison

Kirthevasan Kandasamy

University of Wisconsin-Madison

## Abstract

We study actively labeling streaming data, where an active learner is faced with a stream of data points and must carefully choose which of these points to label via an expensive experiment. Such problems frequently arise in applications such as healthcare and astronomy. We first study a setting when the data’s inputs belong to one of  $K$  discrete distributions and formalize this problem via a loss that captures the labeling cost and the prediction error. When the labeling cost is  $B$ , our algorithm, which chooses to label a point if the uncertainty is larger than a time and cost dependent threshold, achieves a worst-case upper bound of  $\tilde{O}(B^{\frac{1}{3}} K^{\frac{1}{3}} T^{\frac{2}{3}})$  on the loss after  $T$  rounds. We also provide a more nuanced upper bound which demonstrates that the algorithm can adapt to the arrival pattern, and achieves better performance when the arrival pattern is more favorable. We complement both upper bounds with matching lower bounds. We next study this problem when the inputs belong to a continuous domain and the output of the experiment is a smooth function with bounded RKHS norm. After  $T$  rounds in  $d$  dimensions, we show that the loss is bounded by  $\tilde{O}(B^{\frac{1}{d+3}} T^{\frac{d+2}{d+3}})$  in an RKHS with a squared exponential kernel and by  $\tilde{O}(B^{\frac{1}{2d+3}} T^{\frac{2d+2}{2d+3}})$  in an RKHS with a Matérn kernel. Our empirical evaluation demonstrates that our method outperforms other baselines in several synthetic experiments and two real experiments in medicine and astronomy.

## 1 Introduction

The success of real-world supervised learning methods, where we wish to learn a mapping from inputs to outputs, often relies on the existence of abundant labeled training

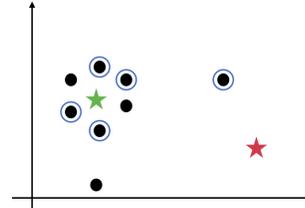


Figure 1: An illustration of the problem set up when the inputs  $\{x_t\}_t$  are in 2D. Each black dot represents a data point that arrived in previous rounds. If it is circled in blue, it indicates that it was chosen for labeling. If the current input  $x_t$  is similar to previously labeled inputs (e.g green star), the algorithm may choose to forego labeling and predict  $f(x_t)$  based on these points to reduce labeling costs. If it is dissimilar to previously labeled points (e.g red star), it may choose to label it to improve the prediction.

data, i.e the outputs are known for the given inputs. However, labeling data can be expensive, as it may require costly experimentation or human effort. The vast literature on active learning has focused on methods for reducing the labeling complexity in such use cases. The majority of such papers study *pool-based active learning*, where there is either a finite or infinite pool of data, and an algorithm should choose which samples from this pool to label in order to learn the mapping well. These methods have been applied successfully in real-world problems, such as text classification, information extraction, materials discovery, astrophysics, and agriculture (McCallum and Nigam, 1998; Wu, 2019; Chandrasekaran et al., 2020; Kusne et al., 2020; Yang et al., 2022; Kandasamy et al., 2017b).

*Streaming* settings for active learning, where an algorithm should choose which data points to label from a possibly infinite stream of data, has received comparatively less attention. Unlike in pool-based settings, where a learner can choose any sample from this pool at any point for labeling, here the learner goes through the data sequentially and has to make decisions to label it on a per-instance basis.

In this work, we consider the following setting where our goal is to learn an underlying function  $f$  from inputs to outputs. On each round  $t$ , the algorithm receives a data’s input  $x_t$ . It may choose to observe a noisy label  $y_t \in \mathbb{R}$  for this point by incurring a *labeling cost*, or not observe this label at no cost. In either case, it must output a prediction  $p_t$  at the end of the round. The algorithm incurs a *faulty prediction cost* proportional to  $|f(x_t) - p_t|$  for choosing

a prediction that is far away from the true function. The goal of the algorithm is to learn the mapping sufficiently well to minimize the sum of both types of costs over a period of time. If we label on each round we incur a large labeling cost; however, not labeling a sufficient number of points incurs a large faulty prediction cost as we will not be able to estimate  $f$  well. To solve this problem effectively, an algorithm should leverage information gained from past experiments to estimate the cost of faulty prediction, and assess this against the cost of labeling to decide if  $x_t$  needs to be labeled. We have illustrated this in Fig. 1.

This problem is motivated by applications in medicine and healthcare. For example, performing an expensive medical exam to assess the risk of a disease for a patient can be costly. In such cases, instead of performing this exam on all patients, a doctor may wish to quickly screen a patient based on their features, past medical history, and cheap medical tests, to determine if an expensive exam is necessary. In addition to saving costs, this will also ensure that valuable medical resources are used prudently, reducing unnecessary waiting times for patients to achieve efficiency. Other applications include reducing the experimentation costs in astronomy, such as when observing streams of astronomical events (e.g supernovae, comets), and reducing investigation costs in finance, such as when streams of customers are applying for large loans.

We study this problem in two different settings: firstly, when each data point belongs to one of  $K$  discrete types; secondly, when  $x_t \in \mathbb{R}^d$  and  $f$  belongs to a reproducing kernel Hilbert space (RKHS) with bounded RKHS norm. In both cases, we formalize this problem via a loss function capturing both costs at each time and describe a simple and efficient algorithm that chooses to label if the uncertainty of  $f(x_t)$  is larger than a threshold determined by the labeling cost and the past arrival pattern. **Our contributions** are:

- First, in §2, we study this problem when the input belongs to one of  $K$  discrete types. Algorithm 1 achieves a  $\tilde{O}(B^{1/3} \sum_{t=1}^T M_{T,k}^{2/3})$  upper bound on the loss where  $B$  is the cost of labeling and  $M_{T,k}$  is the total number of arrivals for type  $k$  after  $T$  rounds. We complement this upper bound with a matching (up to log factors) lower bound. Under the worst possible arrival pattern, the upper and lower bounds become  $\Theta(B^{1/3} K^{1/3} T^{2/3})$ .
- Second, in §3, to model real-world use cases, we study the setting where  $x_t \in \mathbb{R}^d$  and the underlying function  $f$  is in an RKHS. Algorithm 2, which is similar in spirit to Algorithm 1, achieves a  $\tilde{O}(B^{\frac{1}{d+3}} T^{\frac{d+2}{d+3}})$  worst-case upper bound on the loss when the RKHS has a squared exponential kernel and a  $\tilde{O}(B^{\frac{1}{2d+3}} T^{\frac{2d+2}{2d+3}})$  worst-case upper bound when the RKHS has a Matérn kernel.
- Third, in §4, we evaluate both algorithms empirically on several synthetic examples and two real-world use

cases in medicine and astronomy. Our methods outperform baseline methods under different arrival patterns of input data points. When compared to the next best baseline, our method achieves an average loss that is lower by a factor of  $\sim 27\%$  in the medical application and  $\sim 57\%$  in the astronomical application.

## Related work

The majority of the work in active learning studies the *pool-based* setting where a learner can choose to label a data point at any time in the learning process. In addition to the works which focus on applications for the pool-based setting of active learning (McCallum and Nigam, 1998; Settles, 2009; Wu, 2019; Chandrasekaran et al., 2020; Kusner et al., 2020; Yang et al., 2022), there is a rich body of theoretical work (Zhang and Chaudhuri, 2015; Chaudhuri et al., 2015; Balcan and Feldman, 2013; Balcan et al., 2010; Hanneke and Yang, 2019, 2015; Hanneke, 2011; Gentile et al., 2022). However, the *streaming* setting, which we consider in this paper, is more challenging. Since if we choose not to label the data point when it arrives, we may never be able to label that point again in the future.

The majority of the works in the *streaming* setting focus on applications. Zliobaite et al. (2014) explored active learning for streaming data when the distribution of the incoming data shifts over time. They choose to label each data point when the uncertainty is below a threshold, but unlike us, this threshold does not depend on the labeling cost. Moreover, their threshold is chosen in an ad-hoc manner while our threshold is chosen based on theoretical analysis. Lindstrom et al. (2010) considered a batch version of this problem when the incoming data shifts with time. They group a stream of unlabelled data into batches and feed each batch into a classifier to select a subset for labeling. Zhu et al. (2007) studied a similar setting where they build a classifier on a subset of a data batch randomly and use uncertainty sampling, which queries the data that the learner is most uncertain about, to label more instances within this batch. Our setting is different from these works since we don't assume data comes in batches.

Among works with theoretical guarantees in the streaming setting, Wang et al. (2021); Ban et al. (2022); DeSalvo et al. (2021) studied active learning to solve binary or multiple classification tasks. Chu et al. (2011) studied a similar setting while focusing on unbiased online active learning. In addition to the technical and algorithmic differences, our work focuses more on regression problems and accounts for the cost of labeling. Werner et al. (2022) aimed at finding a subset of the streaming data whose value is close to the same-sized subset that has the optimal value. One important difference in our work, when compared to both the applied and theoretical work above, is that we explicitly model the cost of labeling in our formalism and algorithm. Moreover, unlike the above theoretical work, our formal-

ism is based on a loss term that encapsulates both the labeling cost and prediction errors; the learner’s goal is to balance both effects out to minimize the loss.

There are works studying online learning and active learning under budget constraints (Donmez and Carbonell, 2008; Vijayanarasimhan et al., 2010; Badanidiyuru et al., 2012). This is different from our formalism since we model the cost of labeling relative to the prediction error without any budget and aim to minimize a loss composed of both effects. In particular, our setting allows the learner to label more (possibly infinite) points when there is a benefit to doing so over an infinite time horizon.

Our method in the RKHS setting builds on a line of work on Bayesian active learning and adaptive decision-making which use Gaussian processes (GPs). In this flavor, there are several works specifically focused on applications such as large-scale classification, preference learning and posterior approximation (Pinsler et al., 2019; Houlsby et al., 2011; Kandasamy et al., 2017b). Moreover, Garnett et al. (2013) study active learning for discovering low-rank structure in high-dimensional GP. Bui et al. (2017) and Terry and Choe (2020) use GPs for approximation and regression with streaming data. On the theoretical side, Chen et al. (2017) consider learning an unknown variable through a sequence of noisy tests and propose an algorithm to choose the test which maximizes the gain in a surrogate objective with a theoretical guarantee. Javdani et al. (2014) consider a similar problem as Chen et al. (2017) when the hypothesis tests may have correlated regions. The seminal work of Srinivas et al. (2010) provided the first theoretical results for smooth bandits using GPs. Our theoretical analysis in the RKHS setting uses some ideas from this paper.

## 2 $K$ Discrete Types

We will first study this problem in a setting where the input  $x_t$  belongs to one of  $K$  discrete types  $\{1, \dots, K\} \doteq [K]$  (i.e similar to  $K$  arms in  $K$ -armed bandits). This simple setting will help us develop key intuitions and optimality properties which we will then use to design an algorithm under more realistic assumptions in §3. On each round  $t$ , an *input* data point  $x_t \in [K]$  arrives. Each type  $k \in [K]$  is associated with a  $\sigma$ -subGaussian distribution with mean  $\mu_k \in [0, 1]$ . If the learner chooses to label  $x_t$ , they will observe a label  $y_t \in \mathbb{R}$  which is drawn from this distribution. Hence,  $\mathbb{E}[y_t] = \mu_{x_t}$ . We can therefore think of the function  $f$  as a mapping from  $[K]$  to  $[0, 1]$ , where  $f(k) = \mu_k$ , although we prefer the notation  $\mu_k$  in this section.

At the end of round  $t$ , the learner should output a prediction  $p_t$  for  $\mu_{x_t}$ . The learner’s *faulty prediction cost* is simply the difference between the prediction and the true mean of  $x_t$ , i.e  $|p_t - \mu_{x_t}|$ . If the learner chooses to label a point to reduce the prediction cost, they incur a *labeling cost*  $B$

(relative to  $|p_t - \mu_{x_t}|$ ). Therefore, the *total loss*  $L_T$  after  $T$  rounds can be written as,

$$\begin{aligned} L_T &= \sum_{t=1}^T \left( B \cdot \mathbb{1}_{\{x_t \text{ is labeled}\}} + |\mu_{x_t} - p_t| \right) \\ &= BN_T + \sum_{t=1}^T |\mu_{x_t} - p_t|, \end{aligned} \quad (1)$$

where  $N_T$  is the total number of points labeled after  $T$  rounds. To obtain sublinear loss, a learner should balance the cost of labeling against the faulty prediction cost. If the learner chooses to label too frequently, the latter might be reduced, but risks incurring a large labeling cost.

We will make no explicit assumption about the arrival pattern, i.e order of arrival of  $x_t$ ’s on each round. They can even be chosen by an adversary. Our theoretical results will show that our algorithm can naturally adapt to the difficulty of the arrival pattern. In practical applications, the exact value for  $B$  may be chosen based on the monetary cost of a labeling experiment, the opportunity cost for using valuable experimental resources, and the risk of adverse consequences if the true mean is not estimated correctly.

### 2.1 Algorithm

Our algorithm for this problem is outlined in Algorithm 1. On each round  $t$ , We will maintain an estimate  $\hat{\mu}_{t,k}$  for each mean  $\mu_k$  along with confidence intervals. If the uncertainty for  $\mu_k$  as measured by the confidence interval is larger than a certain threshold, we will label  $x_t$ ; otherwise, we will not label it. The exact value of the threshold depends on the number of times type  $k$  has arrived and the cost  $B$  of labeling. As we will show in §2.2, this algorithm, while simple and intuitive, is optimal for this problem.

To describe the algorithm formally, we will let  $\ell_t = 1$  if we chose to label  $x_t$  on round  $t$  and let  $\ell_t = 0$  otherwise. Next,  $M_{t,k}$  and  $N_{t,k}$  respectively will denote the number of times type  $k$  arrived in  $t$  rounds and the number of times type  $k$  arrived and was labeled. Finally, let  $N_t$  denote the total number of rounds we labeled. We have:

$$\begin{aligned} M_{t,k} &= \sum_{s=1}^t \mathbb{1}_{\{x_s=k\}}, \\ N_{t,k} &= \sum_{s=1}^t \mathbb{1}_{\{x_s=k, \ell_s=1\}}, \quad N_t = \sum_{k \in [K]} N_{t,k}. \end{aligned} \quad (2)$$

Next, for each  $k \in [K]$ , let  $\hat{\mu}_{t,k}$  denote the sample mean for type  $k$  using samples obtained via the first  $t$  rounds:

$$\hat{\mu}_{t,k} = \frac{1}{N_{t,k}} \sum_{s=1}^t y_t \cdot \mathbb{1}_{\{x_s=k, \ell_s=1\}} \quad (3)$$

Using sub-Gaussian concentration, we can quantify the uncertainty of this estimate via  $U_{t,k}$  defined below. That is,  $(\hat{\mu}_{t,k} - U_{t,k}(\delta), \hat{\mu}_{t,k} + U_{t,k}(\delta))$  is a confidence interval which traps  $\mu_k$  with probability at least  $1 - \delta$ . We have:

$$U_{t,k}(\delta) = \begin{cases} \sqrt{2\sigma^2 \log(2/\delta)/N_{t,k}} & \text{if } N_{t,k} > 0 \\ \infty & \text{otherwise} \end{cases} \quad (4)$$

We can now describe our algorithm. When  $x_t$  arrives, we choose to label it if the uncertainty  $U_{t,k}(\delta)$  is larger than the threshold value  $B^{1/3}M_{t,x_t}^{-1/3}$  (line 4), which depends on the labeling cost  $B$  and the number of times  $x_t$  has arrived (not observed)  $M_{t,x_t}$ . The threshold increases with  $B$ , discouraging frequent labeling when it is expensive, and decreases with  $M_{t,x_t}$ , encouraging frequent labeling the more times a type is observed to reduce accumulating faulty prediction costs. If we choose to label, we update the mean and confidence intervals for  $x_t$  (line 8). Finally, we output the sample mean  $\hat{\mu}_{t,x_t}$  as our prediction  $p_t$  (line 12).

---

**Algorithm 1** Algorithm for  $K$  discrete types
 

---

**Input:** confidence parameter  $\delta$ , labeling cost  $B$

```

1: for time  $t = 1, 2, \dots$  do
2:    $x_t$  appears,  $M_{t,x_t} \leftarrow M_{t-1,x_t} + 1$ 
3:    $\forall k \neq x_t, M_{t,k} \leftarrow M_{t-1,k}$ 
4:   if  $U_{t-1,x_t}(\delta/N_t^3) > B^{1/3}M_{t,x_t}^{-1/3}$  then
      # See (4) for  $U_{t,k}$ , (2) for  $N_t$ 
5:     Label  $x_t$  and observe  $y_t$ 
6:      $N_{t,x_t} \leftarrow N_{t-1,x_t} + 1$ 
7:     Update  $\hat{\mu}_{t,x_t}, U_{t-1,x_t}$  according to (3), (4).
8:      $\forall k \neq x_t, N_{t,k} \leftarrow N_{t-1,k}, \hat{\mu}_{t,k} \leftarrow \hat{\mu}_{t-1,k},$ 
        $U_{t,k} \leftarrow U_{t-1,k}.$ 
9:   else
10:     $\forall k, N_{t,k} \leftarrow N_{t-1,k}, \hat{\mu}_{t,k} \leftarrow \hat{\mu}_{t-1,k},$ 
       $U_{t,k} \leftarrow U_{t-1,k}.$ 
11:   end if
12:   Output prediction  $p_t \leftarrow \hat{\mu}_{t,x_t}$ 
13: end for
    
```

---

## 2.2 Theoretical Results

We now present our theoretical contributions. We defer all proofs to Appendix §6, but outline the main proof intuitions at the end of this section.

**Upper bounds:** To state our first result, recall from (2) that  $M_{T,k}$  denotes the number of arrivals of type  $k \in [K]$  in  $T$  rounds. Theorem 1 upper bounds the loss for Algorithm 1 in terms of these  $M_{T,k}$  values.

**Theorem 1** (Arrival-dependent Upper Bound). *Assume  $\mu_k \in [0, 1]$  for all  $k \in [K]$  and let  $\delta, B$  be given. Then, for Algorithm 1,  $L_T \in \tilde{O}\left(\sum_{k=1}^K B^{1/3}M_{T,k}^{2/3}\right)$ . Precisely,*

with probability at least  $1 - \delta\pi^2/12, \forall T \geq 1$ ,

$$L_T \leq \sum_{k=1}^K \left( (B^{1/3} + B^{-2/3})2\sigma^2 \log \frac{2T^3}{\delta} M_{T,k}^{2/3} + \frac{3}{2}B^{1/3}M_{T,k}^{2/3} \right).$$

The corollary below, which follows from Theorem 1, states the upper bound under the worst possible arrival pattern for Algorithm 1. It follows via an application of the power mean inequality and noting that  $M_{T,1} + \dots + M_{T,k} = T$ .

$$\left( \frac{M_{T,1}^{2/3} + \dots + M_{T,K}^{2/3}}{K} \right)^{3/2} \leq \frac{M_{T,1} + \dots + M_{T,K}}{K} = \frac{T}{K}.$$

**Corollary 2** (Worst-case Upper Bound). *Under the assumptions of Theorem 1,  $L_T \in \tilde{O}(B^{1/3}T^{2/3}K^{1/3})$  for Algorithm 1.*

To interpret these results, observe that while the worst-case bound scales with  $K$ , our algorithm can naturally adapt to the arrival pattern of the types. In particular, when the arrivals are such that only a few types appear frequently, the bound is significantly better. For instance, in the easiest case, when only one type  $k \in [K]$  appears on all rounds, we obtain  $L_T \in \tilde{O}(B^{1/3}T^{2/3})$ . The worst case bound of Corollary 2 occurs when all types appear an equal number of times, i.e.  $M_{t,k} = T/K$  for all  $k$ . Intuitively, if only a few types appear most of the time, then the learner only needs to learn those types well to obtain a small loss whereas if all types appear a large number of times then all their mean values need to be estimated well to achieve a small loss.

**Lower bounds:** Next, we present arrival-dependent and worst-case hardness results to show that the above observations and bounds are fundamental to this problem. Our first result is a lower bound on the expected loss for any algorithm in terms of the number of arrivals for each type. Let  $\mathcal{A}$  be the class of all algorithms for this problem and  $\mathcal{P}$  denote the class of all problems with  $K$  types with  $\sigma$ -subGaussian observations. We will write  $L_T(A, P)$  to explicitly denote the loss (1) of algorithm  $A$  on problem  $P$ .

**Theorem 3** (Arrival-dependent Lower Bound). *Let  $K > 0$ , and  $T \geq 1$  be given and fix the number of arrivals  $\{M_{T,k}\}_{k=1}^K$  for each  $k \in [K]$ . Then,*

$$\inf_{A \in \mathcal{A}} \sup_{P \in \mathcal{P}} \mathbb{E}L_T(A, P) \geq \frac{1}{4}\sigma^{2/3}B^{1/3} \sum_{k=1}^K M_{T,k}^{2/3}$$

Comparing this with the upper bound in Theorem 1, we see that it matches the lower bound up to constant and logarithmic terms. While the upper bound is in high probability, it is easy to obtain a bound in expectation by setting  $\delta = 1/T$  and using the tower property of expectation. Finally, we can obtain a worst-case lower bound, stated formally below, by setting  $M_{T,k} = T/K$  for all  $k$ .

**Corollary 4** (Worst-case Lower Bound). *Let  $K > 0$ , and  $T \geq 1$  be given. Then,*

$$\inf_{A \in \mathcal{A}} \sup_{P \in \mathcal{P}} \sup_{\{M_{t,k}\}_k} \mathbb{E} L_T(A, P) \geq \frac{1}{4} \sigma^{\frac{2}{3}} B^{\frac{1}{3}} K^{\frac{1}{3}} T^{\frac{2}{3}}$$

**Proof Sketches:** We will conclude this section by outlining our main proof ideas. For the upper bound, we separately bound the prediction errors for rounds where we labeled and for rounds we did not. In the rounds that were not labeled, we used concentration results of sub-Gaussian random variables and the fact that the uncertainty is bounded by the threshold  $B^{1/3} M_{t,x_t}^{-1/3}$  to bound the sum of prediction errors. To bound the total number of rounds that type  $k$  was labeled, we used the fact that if  $N_{t,k}$  was too large, that would reduce the uncertainty  $U_{t,k}$  sufficiently fast, and hence it will fall below the threshold resulting in type  $k$  not being labeled. The careful choice of the threshold gives the optimal trade-off between both terms in this bound. For the lower bound, we use Le Cam’s method (Le Cam, 1960) to derive a lower bound of the expected prediction error for type  $k$  as a function of  $N_{t,k}$ . The lower bound is achieved by the best possible growth rate of  $N_{t,k}$  against  $M_{t,k}$ .

### 3 RKHS Setting

While the setting in §2 was useful for developing intuition and optimality results, it does not reflect practical use cases for this problem. Hence, in this section, we will assume that each input is a  $d$ -dimensional (finite) feature vector and that the output can be modeled using a smooth function.

**Assumptions and problem set up:** Specifically, we will assume  $x_t \in [0, 1]^d$  and that if we label  $x_t$ , we will observe  $y_t = f(x_t) + \epsilon_t$  where the noise  $\epsilon_t$  is bounded, i.e  $|\epsilon_t| \leq \sigma$ , and that  $\mathbb{E}[\epsilon_t] = 0$ . To make the problem tractable, we will assume that  $f$  belongs to a reproducing kernel Hilbert space (RKHS) with a kernel  $\kappa : [0, 1]^d \times [0, 1]^d \rightarrow \mathbb{R}_+$ , and moreover has bounded RKHS norm (Smola and Schölkopf, 1998). Our loss for this problem takes a similar form to (1):

$$L_T = N_T B + \sum_{t=1}^T |f(x_t) - p_t|, \quad (5)$$

where recall that  $p_t$  is the prediction made by the learner for  $f(x_t)$  and  $N_t$  is the total number of times we chose to label in  $t$  rounds. For what follows, we will find it necessary to construct estimates along with confidence intervals for the unknown function  $f$  using past data. We will do so via Gaussian processes (GPs). While GPs are a Bayesian construct, they can be used to obtain frequentist confidence intervals when  $f$  is in an RKHS. Hence, we will begin with a brief review of GPs in the Bayesian setting.

**Gaussian Process:** A GP, written as  $f \sim \mathcal{GP}(\mu, \kappa)$ , is characterized by a prior mean function  $\mu : \mathcal{X} \rightarrow \mathbb{R}$  and

prior covariance kernel  $\kappa : \mathcal{X}^2 \rightarrow \mathbb{R}$ . Some common choices for the kernel are the squared exponential (SE) kernel  $\kappa_{\text{SE}}$  and the Matérn kernel  $\kappa_{\text{M}}$ , defined below:

$$\begin{aligned} \kappa_{\text{SE}}(x, x') &= \exp(- (2l^2)^{-1} \|x - x'\|^2), \\ \kappa_{\text{M}}(x, x') &= \frac{2^{1-\nu}}{\Gamma(\nu)} r^\nu B_\nu(r), \quad r = (\sqrt{2\nu}/l) \|x - x'\|. \end{aligned} \quad (6)$$

Here  $l$  is a lengthscale parameter,  $\nu \in \mathbb{R}_+$  is a smoothness parameter for Matérn kernels and  $B_\nu$  is a modified Bessel function. If  $f \sim \mathcal{GP}(\mu, \kappa)$ , then  $f(x)$  is distributed normally as  $\mathcal{N}(\mu(x), \kappa(x, x))$  for all  $x$ . Given  $n$  observations  $A = \{(x_i, y_i)\}_{i=1}^n$ , the posterior for  $f$  is a GP with covariance  $\kappa_A$ , standard deviation  $\sigma_A$ , and mean  $\mu_A$  given by,

$$\begin{aligned} \kappa_A(x, \tilde{x}) &= \kappa(x, \tilde{x}) - k^\top (K + \sigma^2 I_n)^{-1} \tilde{k}, \\ \sigma_A(x) &= \sqrt{\kappa_A(x, x)}, \quad \mu_A(x) = k^\top (K + \sigma^2 I_n)^{-1} Y, \end{aligned} \quad (7)$$

Here  $Y \in \mathbb{R}^n$  such that  $Y_i = y_i$ ,  $k, \tilde{k} \in \mathbb{R}^n$  are such that  $k_i = \kappa(x, x_i)$ ,  $\tilde{k}_i = \kappa(\tilde{x}, x_i)$ ,  $K \in \mathbb{R}^{n \times n}$  is given by  $K_{i,j} = \kappa(x_i, x_j)$ , and  $I_n \in \mathbb{R}^{n \times n}$  is the identity matrix. While GPs usually assume  $y_i = f(x_i) + \epsilon_i$  and  $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$  in the Bayesian setting, when the noise  $\epsilon_i$  is bounded as outlined in the beginning of this section, we can use (7) to obtain frequentist confidence intervals for  $f$ .

**Maximum Information Gain (MIG):** One quantity of interest going forward will be the maximum information gain  $\gamma_n$  for a GP after  $n$  observations (Srinivas et al., 2010):

$$\gamma_n = \max_{\{x_1, \dots, x_n\}} \frac{1}{2} \log \left( 1 + \sigma^{-2} \sigma_{\{(x_j, y_j)\}_{j=1}^{i-1}}^2(x_i) \right). \quad (8)$$

Here,  $\sigma$  is bound on the noise and  $\sigma_{\{(x_j, y_j)\}_{j=1}^{i-1}}$  is the posterior standard deviation using the first  $i-1$  points (7). The MIG is used to quantify the problem complexity in adaptive decision-making with GPs (Srinivas et al., 2010; Kandasamy et al., 2017a; Chen et al., 2017). In the Bayesian setting, the term inside the max operator is the mutual information between  $f$  and the  $n$  observations, and  $\gamma_n$  can be interpreted as the maximum amount of information  $n$  observations gives about  $f$ . The following bounds are known for GPs with SE and Matérn kernels (Srinivas et al., 2010):

$$\begin{aligned} \text{SE kernel: } \gamma_n &\in \mathcal{O}((\log n)^{d+1}), \\ \text{Matérn kernel: } \gamma_n &\in \mathcal{O}\left(n^{\frac{d(d+1)}{2\nu+d(d+1)}} \log n\right). \end{aligned} \quad (9)$$

#### 3.1 Algorithm

Our algorithm for this formulation, outlined in Algorithm 2, is similar in spirit to Algorithm 1, but instead uses the GP posterior mean and standard deviation for the prediction and the uncertainty. For brevity, let us denote the GP mean using the observations in the first  $t-1$  rounds by  $\mu_{t-1}$ . Precisely,  $\mu_{t-1}(x) = \mu_{A_{t-1}}(x)$ , where

$A_{t-1} = \{(x_s, y_s); 1 \leq s \leq t-1 \text{ and } \ell_s = 1\}$ ; recall  $\mu_A$  from (7) and that  $\ell_s = 1$  if we label on round  $s$ . Similarly, let  $\sigma_{t-1}$  denote the GP standard deviation after  $t-1$  rounds.

On each round, when point  $x_t$  arrives, we will choose to label  $x_t$  if the GP standard deviation  $\sigma_{t-1}(x_t)$ , which quantifies the uncertainty of  $\mu_{t-1}(x_t)$ , is larger than a cost and time-dependent threshold  $\tau(t)$  (line 2). This threshold is set separately for the SE and Matérn kernels as follows:

$$\text{SE kernel: } \tau(t) = \sqrt{2\sigma^2 B^{\frac{1}{d+3}} t^{-\frac{1}{d+3}}}, \quad (10)$$

$$\text{Matérn kernel: } \tau(t) = \sqrt{2\sigma^2 B^{\frac{1}{2d+3}} t^{-\frac{1}{2d+3}}}.$$

While these specific choices for  $\tau$  were chosen to optimize for our final bound, we see that, similar to the discrete case, the threshold increases with cost  $B$  to discourage frequent labeling when it is expensive, and decreases with the total number of arrivals  $t$  to encourage labeling as time goes on to minimize faulty prediction costs. The resulting procedure is summarized in Algorithm 2.

---

**Algorithm 2** Algorithm for  $f$  in an RKHS
 

---

**Input:** GP Prior  $\mu_0, \kappa_0$ , labeling cost  $B$

- 1: **for** time  $t = 1, 2, \dots, T$  **do**
  - 2:   **if**  $\sigma_{t-1}(x_t) > \tau(t)$  **then**
  - 3:     Label  $x_t$ . Observe  $y_t = f(x_t) + \epsilon_t$
  - 4:     Update  $\mu_t$  and  $\sigma_t$  with  $(x_t, y_t)$ .   # See (7)
  - 5:   **else**
  - 6:      $\mu_t \leftarrow \mu_{t-1}, \quad \sigma_t \leftarrow \sigma_{t-1}$
  - 7:   **end if**
  - 8:   Output  $p_t \leftarrow \mu_t(x_t)$
  - 9: **end for**
- 

### 3.2 Theoretical Results

Our main theorem in this section is the following upper bound for the loss (5) for Algorithm 2.

**Theorem 5.** *Let  $\delta \in (0, 1)$  be given. Assume that  $f$  lies in an RKHS  $\mathcal{H}_\kappa(\mathcal{X})$  corresponding to the kernel  $\kappa(x, x')$ , and denote its RKHS norm by  $\|f\|_\kappa$ . Assume that the noise  $\epsilon_t$  has a zero mean conditioned on the history and is bounded by  $\sigma$  almost surely. Recall the MIG  $\gamma_n$  be as defined in (8) and (9), and denote  $\beta_t = 2\|f\|_\kappa^2 + 300 \log^3(t/\delta)$ . Then, with probability at least  $1 - \delta$ , for all  $T \geq 0$ , Algorithm 2 satisfies the following for the SE and Matérn kernels.*

For the SE kernel  $L_T \in \tilde{O}\left(B^{\frac{1}{d+3}} T^{\frac{d+2}{d+3}}\right)$  with,

$$L_T \leq (C_1 + C_2 \beta_{T+1}^{1/2}) B^{1-d_1} T^{d_1} + \sqrt{C_3 \beta_{T+1} B^{-d_1} T^{d_1} \gamma_{C_1 B^{-d_1} T^{d_1}}},$$

where  $d_1 = \frac{d+2}{d+3}$ ,  $C_1 = (\sigma l)^{-d} d^{d/2}$ ,  $C_2 = \frac{\sqrt{2}\sigma(d+3)}{d+2}$ ,  $C_3 = \frac{2}{\log(1+\sigma^{-2})} C_1$ , and  $l$  is from (6).

For a Matérn kernel  $L_T \in \tilde{O}\left(B^{\frac{1}{2d+3}} T^{\frac{2d+2}{2d+3}}\right)$  with,

$$L_T \leq (C_3 + C_4 \beta_{T+1}^{1/2}) B^{1-d_2} T^{d_2} + \sqrt{C_5 \beta_{T+1} B^{-d_2} T^{d_2} \gamma_{C_3 B^{-d_2} T^{d_2}}}$$

where  $d_2 = \frac{2d+2}{2d+3}$ ,  $C_3 = \sigma^{-2d} (2\sqrt{d} L_M)^d$ ,  $C_4 = \frac{\sqrt{2}\sigma(2d+3)}{2d+2}$ ,  $C_5 = \frac{2}{\log(1+\sigma^{-2})} C_3$ , and  $L_M$  is a Lipschitz constant for the Matérn kernel (6).

The theorem establishes that Algorithm 2 is able to achieve sublinear loss for both kernels. While the bound admittedly worsens with dimensionality  $d$ , this is to be expected in high dimensions under nonparametric assumptions. For instance, for bandit problems with a Matérn kernel, the regret bound scales at rate  $\tilde{O}\left(T^{\frac{\nu+d(d+1)}{2\nu+d(d+1)}}\right)$ . One key difference in this result, when compared to Theorem 1, is that it only provides a worst-case upper bound. This is primarily because characterizing the number of arrivals in different regions of  $[0, 1]^d$  is not as straightforward as it was in the discrete setting. Our empirical evaluation in §4 shows that Algorithm 2, much like Algorithm 1, does indeed perform better when the arrivals are skewed than if they are uniform.

Finally, it is worth mentioning that in both bounds, the second term of the RHS is of a lower order than the first term; after accounting for the MIG (9), for the SE kernel, the exponent of  $T$  in the second term is  $\frac{d+2}{2(d+3)}$  and for the Matérn kernel, it is  $\frac{2d+2}{2d+3} \frac{\nu+d(d+1)}{2\nu+d(d+1)}$  which is smaller than  $\frac{2d+2}{2d+3}$ .

**Proof sketch:** The main conceptual difference between the RKHS setting and the  $K$  discrete setting is that here we use the GP standard deviation as our measure of uncertainty, which is not directly connected to the number of labeled points. Thus it is difficult to choose a threshold to balance with the number of labels. We instead use a discretization argument which allows us to bound the uncertainty in each bin by the number of labeled points in that bin, and carefully choose the size of the discretization to optimize for the final bound. For this proof, we also used a prior result to obtain frequentist confidence intervals for  $f$  from the GP updates (Srinivas et al., 2010).

## 4 Experiments

We evaluate our algorithms on synthetic and real experiments, both in the discrete setting and when  $f$  is a smooth function. We compare against the following two baselines:

1. *Random Select:* On each round, this baseline labels with probability 0.5. We experimented with different values in  $[0, 1]$  and found that 0.5 worked best across all our experiments. (Note that this includes, as a special case, the baseline which labels all points when the probability of 0.5)

2. *VAR-UNCERTAINTY* (Zliobaite et al., 2014): This baseline chooses to label if the uncertainty is less than a thresh-

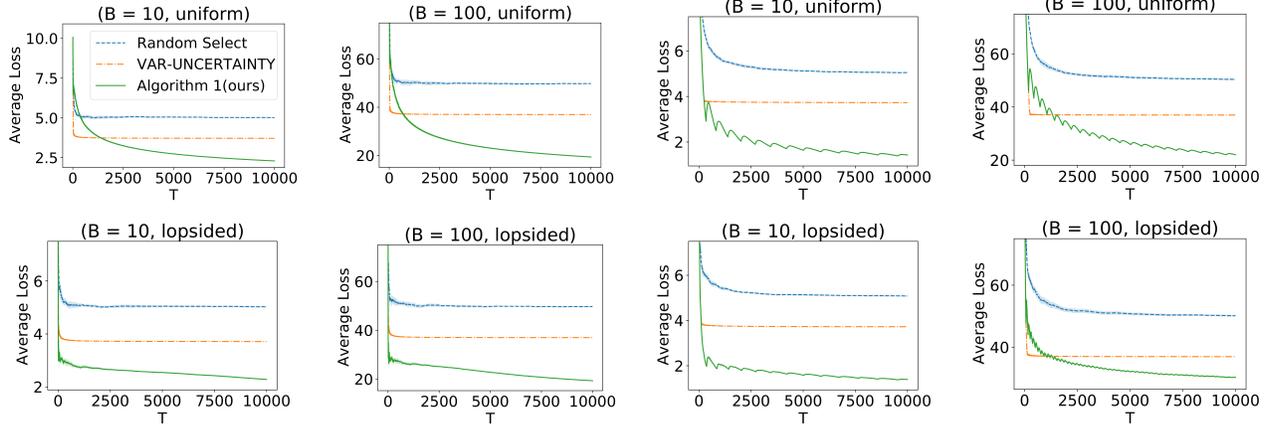


Figure 2: Results comparing the three methods for  $K$  discrete setting of §4.1.  $K = 10$  for column 1 and 2.  $K = 100$  for columns 3 and 4. The results are averaged over 10 trials, and we report the 95% the confidence interval of the standard error.

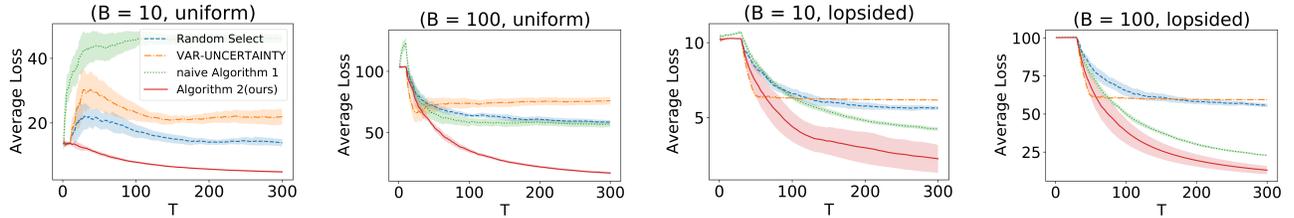


Figure 3: Results comparing the four methods for synthetic functions in §4.2. Columns 1 and 2 show the results when the arrival pattern is uniform for Branin function. Column 3 and 4 show the results when the arrival pattern is lopsided for the Hartmann function. The results are averaged over 10 trials, and we report the 95% confidence interval of the standard error.

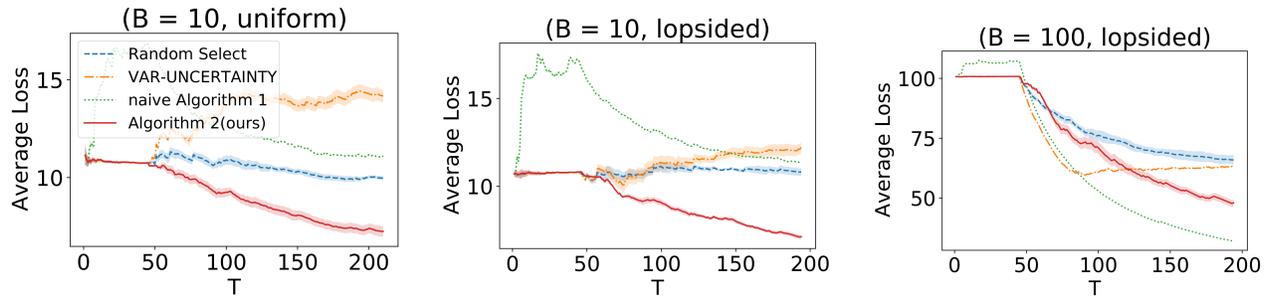


Figure 4: Results comparing the four methods for Parkinsons dataset. The results are averaged over 5 trials and we report the 95% confidence interval of the standard error.

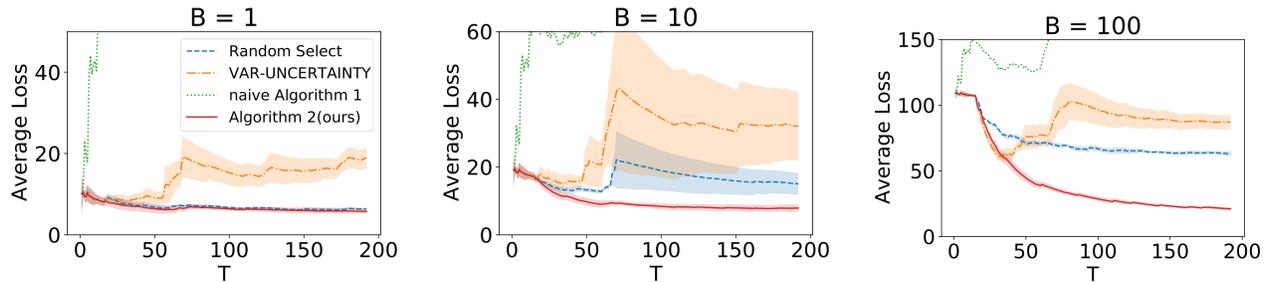


Figure 5: Results comparing the four methods for the supernova dataset. The results are averaged over 5 trials and we report the 95% confidence interval of the standard error. The naive Algorithm 1 performs poorly so we limit the  $y$ -axis to focus on the rest methods.

old  $\theta$ , which is decreased ( $\theta \leftarrow \theta/2$ ) each time if it labels and increased ( $\theta \leftarrow 2\theta$ ) each time if not.  $\theta$  is initially set to be 1 as suggested by the authors. We use the same uncertainty measure as we did for our algorithms in both settings.

For both baselines, we determine  $p_t$  via the sample mean of  $x_t$ 's labels for the  $K$  discrete setting and via the mean of a fitted GP model when  $f$  is a smooth function. We also considered other methods outlined in the related work section but they could not be included for the following reasons. The batch active learning methods (Lindstrom et al., 2010; Zhu et al., 2007) are hard to adapt to our setting since we assume data comes one at a time. Some methods (Werner et al., 2022; DeSalvo et al., 2021) are tailored for classification and have no straightforward generalization to our (regression) problem. The remaining methods are for pool-based active learning and do not apply to our setting.

#### 4.1 Synthetic experiments with $K$ discrete types

First, we compare the performance between Algorithm 1 with *Random Select* and *VAR-UNCERTAINTY* when  $K \in \{10, 100\}$ . For each value of  $K$ , the actual mean values are generated by random sampling in the domain  $[0, 1]$ . The labels  $y_t$  are drawn from a Gaussian distribution with variance 0.01 and the corresponding mean value.

We design two different arrival patterns: (i) *uniform*: when all  $K$  types appear uniformly at random, (ii) *lopsided*: when 20% of the types appear 80% of the time and the remaining 80% of the types appear 20% of the time. The cost of labeling is set to be  $B = \{10, 100\}$  and the total number of rounds is set to be  $T = 10^4$ . We slightly modify our threshold function to  $\lambda B^{1/3} M_{t,k}^{-1/3}$ , where  $\lambda > 0$  is a hyperparameter tuned on a held-out synthetic dataset. (Note that our rates do not change by adding a constant to the threshold function.) The values of  $\lambda$  used in each experiment can be found in Section 7 in the Appendix. Figure 2 shows the experimentation results when the arrival pattern is uniform and lopsided for both  $K = 10$  and  $K = 100$ .

#### 4.2 Synthetic experiments with smooth $f$

In this section, we compare Algorithm 2 with *Random Select*, *VAR-UNCERTAINTY*, and a naive version of Algorithm 1, which splits each domain into  $2^d$  discrete types by halving along each dimension. For Algorithm 2, we use a GP prior with zero mean and an SE kernel. While our theoretical analysis assumes the kernel parameters (e.g length scales) are known, they need to be tuned well for good empirical performance. Hence, we initialize this method by choosing to label the first  $5d$  points and tuning the kernel parameters by maximizing the GP marginal likelihood. We adopt the same procedure for the baselines since they also use a GP to output the prediction. We run each experiment under different costs  $B = \{10, 10^2\}$  and  $T = 300$ . As we did for the  $K$  discrete setting, we modified the threshold function to  $\tau(t) = \lambda \cdot \sqrt{2\sigma^2} B^{\frac{1}{d+3}} C^{-\frac{1}{d+3}}$  where the hyper-

parameter  $\lambda > 0$  was tuned on a held-out set.

**Test functions:** We use the Branin function ( $d = 2$ ) and Hartmann function ( $d = 6$ ) for our synthetic evaluation (Dixon, 1978). We form a uniform arrival pattern for the Branin function whose input domain is  $[-5, 10] \times [0, 15]$ . For the Hartmann function whose input domain is  $[0, 1]^6$ , we form a lopsided arrival pattern where 80% of the  $x_t$ 's coming from  $[0, 0.2] \times [0, 1]^5$  and 20% coming from  $[0.2, 1] \times [0, 1]^5$ . For the Branin function, we set the standard deviation of the noise to be  $\sigma = 5$  and for the Hartmann function, we set it to be  $\sigma = 0.5$ .

Figure 3 shows the results for both of the synthetic functions. Algorithm 2 performs the best among all four methods under different arrival patterns and costs.

#### 4.3 Real experiments with smooth $f$

In this section, we consider the following two real-world datasets in healthcare and astronomy:

**Healthcare:** We use the Parkinsons telemonitoring dataset (Tsanas et al., 2009). They are 5875 biomedical voice recordings from 42 people who have early-stage Parkinson's disease. Each recording contains 20 measurements (from voice and personal statistics), from which we choose 9 different measurements (*age, sex, Jitter(%)*, *Shimmer, NHR, HNR, RPDE, DFA, PPE*) as the features for the input  $x_t$  and use the measurement *total UPDRS* as the label  $f(x_t)$  that needs to be predicted.

As before, we create both uniform and lopsided arrival patterns for this experiment. For uniform arrivals, we randomly select 5 recordings from each patient and shuffle the order resulting in 210 data points. For lopsided arrivals, we randomly select 20 recordings from patients whose *subject#* are from 1 to 8 and randomly select 1 recording from each of the rest patients. There is a total of 194 data points. We manually add noise to each label, the noises are independent and come from the distribution of  $\mathcal{N}(0, 1)$ .

Figure 4 shows the results of the Parkinsons dataset under the uniform and lopsided arrival patterns. Algorithm 2 outperforms almost all the other methods. It is interesting to note that the naive implementation of Algorithm 1 performs better than Algorithm 2 when  $B = 100$ . We found that while the naive algorithm's prediction was worse than other methods, it was able to achieve a lower overall cost by reducing the amount of labeling.

**Astronomy:** We use the supernova dataset from Davis et al. (2007). There are in total 192 data points in  $\mathbb{R}^3$ . Each dimension of the data point represents the Hubble constant, dark matter fraction, and dark energy fraction respectively. We use prior simulators (Kandasamy et al., 2017a; Robertson, 1936; Shchigolev, 2017) to get the label for each data point and manually add independent  $\mathcal{N}(0, 100)$  noise to each label (we choose the noise variance to be large since

the range of the labels is large  $[-980, -129]$ ).

Figure 5 shows the results for all four methods. We observe that *Random Select* and Algorithm 2 perform relatively similarly when  $B = 1$ . However, when  $B$  increases, Algorithm 2 outperforms all the other methods.

This concludes our experiments. In the Appendix, we have provided additional experiments, plots of average prediction errors for the above methods, and an ablation study of the effect of the tuning parameter  $\lambda$  in our algorithms.

## 5 Conclusion

We studied actively labeling streaming data, where we need to balance between the cost of labeling and the cost of making faulty predictions with few labels. We described two algorithms in two different settings based on the same principle: label a point if its predictive uncertainty is larger than a threshold; this threshold increases with the cost of labeling and decreases with the number of arrivals. In the first setting, where the inputs belong to  $K$  discrete distributions, we provided matching lower and upper bounds on the loss in terms of the arrival pattern of queries. In the second setting, where the expected output can be modeled using an RKHS, we showed that the algorithm achieves sublinear loss in the worst case. We corroborated these results with empirical results on real and synthetic experiments.

## References

- Jean-Yves Audibert, Rémi Munos, and Csaba Szepesvári. Exploration-exploitation tradeoff using variance estimates in multi-armed bandits. *Theoretical Computer Science*, 410(19):1876–1902, 2009. URL <http://www.ualberta.ca/~szepesva/papers/ucbtuned-journal.pdf>.
- Ashwinkumar Badanidiyuru, Robert Kleinberg, and Yaron Singer. Learning on a budget: posted price mechanisms for online procurement. In *Proceedings of the 13th ACM conference on electronic commerce*, pages 128–145, 2012.
- Maria-Florina Balcan, Steve Hanneke, and Jennifer Wortman Vaughan. The true sample complexity of active learning. *Machine learning*, 80(2):111–139, 2010.
- Maria-Florina F Balcan and Vitaly Feldman. Statistical active learning algorithms. *Advances in neural information processing systems*, 26, 2013.
- Yikun Ban, Yuheng Zhang, Hanghang Tong, Arindam Banerjee, and Jingrui He. Improved algorithms for neural active learning. In *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=riIaC2ivcYA>.
- Thang D Bui, Cuong Nguyen, and Richard E Turner. Streaming sparse gaussian process approximations. *Advances in Neural Information Processing Systems*, 30, 2017.
- Varun Chandrasekaran, Kamalika Chaudhuri, Irene Giacomelli, Somesh Jha, and Songbai Yan. Exploring connections between active learning and model extraction. In *29th USENIX Security Symposium (USENIX Security 20)*, pages 1309–1326, 2020.
- Kamalika Chaudhuri, Sham M Kakade, Praneeth Netrapalli, and Sujay Sanghavi. Convergence rates of active learning for maximum likelihood estimation. *Advances in Neural Information Processing Systems*, 28, 2015.
- Yuxin Chen, Hamed Hassani, and Andreas Krause. Near-optimal bayesian active learning with correlated and noisy tests. In *Artificial Intelligence and Statistics*, pages 223–231. PMLR, 2017.
- Wei Chu, Martin Zinkevich, Lihong Li, Achint Thomas, and Belle Tseng. Unbiased online active learning in data streams. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 195–203, 2011.
- Tamara M Davis, Edvard Mörtzell, Jesper Sollerman, Andrew C Becker, Stephanie Blondin, P Challis, Alejandro Clocchiatti, AV Filippenko, RJ Foley, Peter M Garnavich, et al. Scrutinizing exotic cosmological models using essence supernova data combined with other cosmological probes. *The Astrophysical Journal*, 666(2): 716, 2007.
- Giulia DeSalvo, Claudio Gentile, and Tobias Sommer Thune. Online active learning with surrogate loss functions. *Advances in Neural Information Processing Systems*, 34:22877–22889, 2021.
- Laurence Charles Ward Dixon. The global optimization problem. an introduction. *Toward global optimization*, 2:1–15, 1978.
- Pinar Donmez and Jaime G Carbonell. Proactive learning: cost-sensitive active learning with multiple imperfect oracles. In *Proceedings of the 17th ACM conference on Information and knowledge management*, pages 619–628, 2008.
- Roman Garnett, Michael A Osborne, and Philipp Hennig. Active learning of linear embeddings for gaussian processes. *arXiv preprint arXiv:1310.6740*, 2013.
- Claudio Gentile, Zhilei Wang, and Tong Zhang. Achieving minimax rates in pool-based batch active learning. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato, editors, *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 7339–7367. PMLR, 2022. URL <https://proceedings.mlr.press/v162/gentile22a.html>.

- Steve Hanneke. Rates of convergence in active learning. *The Annals of Statistics*, pages 333–361, 2011.
- Steve Hanneke and Liu Yang. Minimax analysis of active learning. *J. Mach. Learn. Res.*, 16(1):3487–3602, 2015.
- Steve Hanneke and Liu Yang. Surrogate losses in passive and active learning. *Electronic Journal of Statistics*, 13(2):4646–4708, 2019.
- Neil Houlsby, Ferenc Huszár, Zoubin Ghahramani, and Máté Lengyel. Bayesian active learning for classification and preference learning. *arXiv preprint arXiv:1112.5745*, 2011.
- Shervin Javdani, Yuxin Chen, Amin Karbasi, Andreas Krause, Drew Bagnell, and Siddhartha Srinivasa. Near optimal bayesian active learning for decision making. In *Artificial Intelligence and Statistics*, pages 430–438. PMLR, 2014.
- Kirthevasan Kandasamy, Gautam Dasarathy, Junier B. Oliva, Jeff G. Schneider, and Barnabás Póczos. Multi-fidelity gaussian process bandit optimisation. *CoRR*, abs/1603.06288, 2016. URL <http://arxiv.org/abs/1603.06288>.
- Kirthevasan Kandasamy, Gautam Dasarathy, Jeff Schneider, and Barnabás Póczos. Multi-fidelity bayesian optimisation with continuous approximations. In *International Conference on Machine Learning*, pages 1799–1808. PMLR, 2017a.
- Kirthevasan Kandasamy, Jeff Schneider, and Barnabás Póczos. Query efficient posterior estimation in scientific experiments via bayesian active learning. *Artificial Intelligence*, 243:45–56, 2017b.
- A Gilad Kusne, Heshan Yu, Changming Wu, Huairuo Zhang, Jason Hattrick-Simpers, Brian DeCost, Suchismita Sarker, Corey Oses, Cormac Toher, Stefano Curtarolo, et al. On-the-fly closed-loop materials discovery via bayesian active learning. *Nature communications*, 11(1):1–11, 2020.
- Lucien Le Cam. An approximation theorem for the poisson binomial distribution. *Pacific Journal of Mathematics*, 10(4):1181–1197, 1960.
- Patrick Lindstrom, Sarah Jane Delany, and Brian Mac Namee. Handling concept drift in a text data stream constrained by high labelling cost. In *Twenty-Third International FLAIRS Conference*, 2010.
- Andrew McCallum and Kamal Nigam. Employing EM and pool-based active learning for text classification. In Jude W. Shavlik, editor, *Proceedings of the Fifteenth International Conference on Machine Learning (ICML 1998)*, Madison, Wisconsin, USA, July 24-27, 1998, pages 350–358. Morgan Kaufmann, 1998.
- Robert Pinsler, Jonathan Gordon, Eric Nalisnick, and José Miguel Hernández-Lobato. Bayesian batch active learning as sparse subset approximation. *Advances in neural information processing systems*, 32, 2019.
- HP Robertson. An interpretation of page’s” new relativity”. *Physical Review*, 49(10):755, 1936.
- Burr Settles. Active learning literature survey. 2009.
- VK Shchigolev. Calculating luminosity distance versus redshift in flrw cosmology via homotopy perturbation method. *Gravitation and Cosmology*, 23(2):142–148, 2017.
- Alex J Smola and Bernhard Schölkopf. *Learning with kernels*, volume 4. Citeseer, 1998.
- Niranjan Srinivas, Andreas Krause, Sham M. Kakade, and Matthias W. Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In Johannes Fürnkranz and Thorsten Joachims, editors, *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, June 21-24, 2010, Haifa, Israel, pages 1015–1022. Omnipress, 2010. URL <https://icml.cc/Conferences/2010/papers/422.pdf>.
- Nick Terry and Youngjun Choe. Splitting gaussian process regression for streaming data. *arXiv preprint arXiv:2010.02424*, 2020.
- Athanasios Tsanas, Max Little, Patrick McSharry, and Lorraine Ramig. Accurate telemonitoring of parkinson’s disease progression by non-invasive speech tests. *Nature Precedings*, pages 1–1, 2009.
- Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- Sudheendra Vijayanarasimhan, Prateek Jain, and Kristen Grauman. Far-sighted active learning on a budget for image and video recognition. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 3035–3042. IEEE, 2010.
- Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.
- Zhilei Wang, Pranjal Awasthi, Christoph Dann, Ayush Sekhari, and Claudio Gentile. Neural active learning with performance guarantees. In *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=iPHnzuU6S94>.
- Mariel A Werner, Anastasios Angelopoulos, Stephen Bates, and Michael I Jordan. Online active learning with dynamic marginal gain thresholding. *arXiv preprint arXiv:2201.10547*, 2022.
- Dongrui Wu. Pool-based sequential active learning for regression. *IEEE Trans. Neural Networks Learn. Syst.*, 30(5):1348–1359, 2019. doi: 10.1109/TNNLS.2018.2868649. URL <https://doi.org/10.1109/TNNLS.2018.2868649>.

- Yana Yang, Yang Li, Jiachen Yang, and Jiabao Wen. Dissimilarity-based active learning for embedded weed identification. *Turkish Journal of Agriculture and Forestry*, 46(3):390–401, 2022.
- Chicheng Zhang and Kamalika Chaudhuri. Active learning from weak and strong labelers. *Advances in Neural Information Processing Systems*, 28, 2015.
- Xingquan Zhu, Peng Zhang, Xiaodong Lin, and Yong Shi. Active learning from data streams. In *Seventh IEEE International Conference on Data Mining (ICDM 2007)*, pages 757–762. IEEE, 2007.
- Indre Zliobaite, Albert Bifet, Bernhard Pfahringer, and Geoffrey Holmes. Active learning with drifting streaming data. *IEEE Trans. Neural Networks Learn. Syst.*, 25(1):27–39, 2014. doi: 10.1109/TNNLS.2012.2236570. URL <https://doi.org/10.1109/TNNLS.2012.2236570>.

## Appendix

### 6 PROOFS

In this section, we present the detailed proofs of the Theorems and Corollaries in the main paper.

#### 6.1 Proof of Theorem 1

*Proof.* The proof goes as follows: for each type, we will decompose the prediction errors into rounds that are labeled and rounds that are not labeled and then bound the two parts separately. After this, we will upper-bound  $N_{T,k}$ , the number of times type  $k$  was labeled, to bound the total labeling cost.

The following decomposition of the loss divides the prediction errors into rounds that are labeled ( $\ell_t = 1$ ) and not labeled ( $\ell_t = 0$ ) separately,

$$\begin{aligned} L_T &= \sum_{k=1}^K \left( B \cdot N_{T,k} + \sum_{t=1}^T |\mu_k - \hat{\mu}_{t,k}| \cdot \mathbb{1}_{\{x_t=k\}} \right) \\ &= \sum_{k=1}^K \left( B \cdot N_{T,k} + \underbrace{\sum_{x_t=k, \ell_t=1} |\mu_k - \hat{\mu}_{t,k}|}_{A_1} + \underbrace{\sum_{x_t=k, \ell_t=0} |\mu_k - \hat{\mu}_{t,k}|}_{A_2} \right) \end{aligned} \quad (11)$$

We will bound  $A_2$  using sub-Gaussian concentration and the threshold function. For this, we will consider the following event where  $|\mu_k - \hat{\mu}_{t,k}|$  can be bounded  $\forall t, \forall k$ .

$$\mathcal{E} = \left\{ \hat{\mu}_{t,k} - \sqrt{\frac{2\sigma^2 \log(2N_t^3/\delta)}{N_{t-1,k}}} \leq \mu_k \leq \hat{\mu}_{t,k} + \sqrt{\frac{2\sigma^2 \log(2N_t^3/\delta)}{N_{t-1,k}}} \quad \forall t \geq 1, \forall k \in [K] \right\}$$

Then we consider the probability that event  $\mathcal{E}$  is false,

$$\begin{aligned} \mathbb{P}(\bar{\mathcal{E}}) &= \mathbb{P} \left( \forall t, \forall k \in [K], \left| \frac{1}{N_{t,k}} \sum_{s=1}^t y_s \mathbb{1}_{\{x_s=k, \ell_s=1\}} - \mu_k \right| \geq \sqrt{\frac{2\sigma^2 \log(2N_t^3/\delta)}{N_{t-1,k}}} \right) \\ &= \mathbb{P} \left( \forall t, \left| \frac{1}{N_{t,x_t}} \sum_{s=1}^t y_s \mathbb{1}_{\{x_s=x_t, \ell_s=1\}} - \mu_{x_t} \right| \geq \sqrt{\frac{2\sigma^2 \log(2N_t^3/\delta)}{N_{t-1,x_t}}} \right) \end{aligned}$$

where the second equality is because we can assume the sequence of the incoming data points is fixed. Then we take a union bound on  $t$  and get

$$\mathbb{P}(\bar{\mathcal{E}}) \leq \sum_{t=1}^{\infty} \mathbb{P} \left( \left| \frac{1}{N_{t,x_t}} \sum_{s=1}^t y_s \mathbb{1}_{\{x_s=x_t, \ell_s=1\}} - \mu_{x_t} \right| \geq \sqrt{\frac{2\sigma^2 \log(2N_t^3/\delta)}{N_{t-1,x_t}}} \right)$$

Since  $N_{t-1,x_t}$  is a random variable that can take values in  $\{1, \dots, t-1\}$ , we take a union bound on  $N_{t-1,x_t}$ . Also since  $N_t \geq N_{t-1,x_t}$ , we replace  $N_t^3$  by  $N_{t-1,x_t}^3$  in the probability and get

$$\begin{aligned} \mathbb{P}(\bar{\mathcal{E}}) &\leq \sum_{t=1}^{\infty} \sum_{j=1}^{t-1} \mathbb{P} \left( \left| \frac{1}{N_{t,x_t}} \sum_{s=1}^t y_s \mathbb{1}_{\{x_s=x_t, \ell_s=1\}} - \mu_{x_t} \right| \geq \sqrt{\frac{2\sigma^2 \log(2j^3/\delta)}{j}} \right) \\ &\leq \sum_{t=1}^{\infty} \sum_{j=1}^{t-1} 2 \exp \left( -\frac{j}{2\sigma^2} \cdot \frac{2\sigma^2 \log(2j^3/\delta)}{j} \right) && \text{(by subGaussian concentration)} \\ &\leq \sum_{t=1}^{\infty} \sum_{j=1}^{t-1} \frac{\delta}{j^3} && \text{(by cancelling common terms)} \\ &\leq \sum_{t=1}^{\infty} \int_0^t \frac{\delta}{j^3} dj \end{aligned}$$

$$\begin{aligned}
 &= \sum_{t=1}^{\infty} \frac{\delta}{2t^2} \\
 &= \frac{\delta\pi^2}{12}.
 \end{aligned}$$

Thus  $\mathbb{P}(\mathcal{E}) \geq 1 - \frac{\delta\pi^2}{12}$ . Now, consider  $L_T$  under event  $\mathcal{E}$ . Recall our assumption that  $\forall k, \mu_k \in [0, 1]$  which allows us to bound part  $A_1$  by replacing each term by 1. Thus we get, with probability at least  $1 - \frac{\delta\pi^2}{12}$ ,

$$\begin{aligned}
 L_T &\leq \sum_{k=1}^K \left( B \cdot N_{T,k} + \sum_{x_t=k, \ell_t=1} 1 + \sum_{x_t=k, \ell_t=0} \sqrt{\frac{2\sigma^2 \log(2N_t^3/\delta)}{N_{t-1,k}}} \right) \\
 &\leq \sum_{k=1}^K \left( (B+1) \cdot N_{T,k} + \sum_{x_t=k, \ell_t=0} B^\beta M_{t,k}^\alpha \right)
 \end{aligned}$$

where the second inequality is because:  $\ell_t = 0$  (not labeled) indicates our current confidence radius  $\sqrt{\frac{2\sigma^2 \log(2N_t^3/\delta)}{N_{t-1,k}}}$  is less than the threshold  $B^\beta M_{t,k}^\alpha$ . Next we bound  $\sum_{x_t=k, \ell_t=0} B^\beta M_{t,k}^\alpha$  by the integral from  $t = 0$  to  $M_{T,k}$ ,

$$\begin{aligned}
 L_T &\leq \sum_{k=1}^K \left( (B+1) \cdot N_{T,k} + \int_0^{M_{T,k}} B^\beta t^\alpha dt \right) \\
 &= \sum_{k=1}^K \left( (B+1) \cdot N_{T,k} + \frac{B^\beta}{\alpha+1} M_{T,k}^{\alpha+1} \right)
 \end{aligned}$$

To bound  $N_{T,k}$ , we use the same strategy as Theorem 2 in [Audibert et al. \(2009\)](#):  $\forall u \in \mathbb{N}_+$ , we define  $S_k = \{t : x_t = k, N_{t-1,k} \geq u\}$ , recall the definition of  $N_{T,k}$ ,

$$N_{T,k} = \sum_{t=1}^{\infty} \mathbb{1}_{\{x_t=k, \ell_t=1\}} \leq u + \sum_{t: N_{t-1,k} \geq u} \mathbb{1}_{\{x_t=k, \ell_t=1\}} = u + \sum_{t \in S_k} \mathbb{1}_{\left\{ \sqrt{\frac{2\sigma^2 \log(2N_t^3/\delta)}{N_{t-1,k}}} > B^\beta M_{t,k}^\alpha \right\}} \quad (12)$$

where the last equality is because  $\ell_t = 1$  indicates the confidence radius exceeds the threshold. Then by rearranging the terms in (12) and using the fact that  $N_{t-1,k} \geq u$  in  $S_k$ , we get

$$N_{T,k} \leq u + \sum_{t \in S_k} \mathbb{1}_{\left\{ u < \frac{2\sigma^2 \log(2N_t^3/\delta)}{B^{2\beta} M_{t,k}^{2\alpha}} \right\}} \quad (13)$$

We let  $u = 2\sigma^2 \log(2N_T^3/\delta) B^{-\frac{2}{3}} M_{T,k}^{\frac{2}{3}}$ ,  $\beta = \frac{1}{3}$  and  $\alpha = -\frac{1}{3}$ , since we would like to choose  $u$  such that  $u \geq \frac{2\sigma^2 \log(2N_t^3/\delta)}{B^{2\beta} M_{t,k}^{2\alpha}}$  and then we can bound  $N_{T,k}$  by  $u$ . Recall that  $M_{T,k} \geq M_{t,k}$ ,  $N_T \geq N_t$ ,  $\forall t \leq T$ , we get

$$u B^{2\beta} M_{t,k}^{2\alpha} = 2\sigma^2 \log(2N_T^3/\delta) B^{-\frac{2}{3}} M_{T,k}^{\frac{2}{3}} \cdot B^{\frac{2}{3}} M_{t,k}^{-\frac{2}{3}} \geq 2\sigma^2 \log(2N_t^3/\delta)$$

The inequality indicates that each term inside the summation of (13) becomes 0, thus  $N_{T,k} \leq u$ . Since  $N_{T,k} \leq T$ , finally we get

$$L_T \leq \sum_{k=1}^K \left( (B^{\frac{1}{3}} + B^{-\frac{2}{3}}) 2\sigma^2 \log(2T^3/\delta) M_{T,k}^{\frac{2}{3}} + \frac{3}{2} B^{\frac{1}{3}} M_{T,k}^{\frac{3}{3}} \right)$$

□

## 6.2 Proof of Theorem 3

In this proof, we follow the standard procedure of Le Cam's method (Theorem 7) to derive the minimax lower bound. We include the proof here for completeness. First, recall the formulation of our expected loss:

$$\mathbb{E}L_T = \sum_{k=1}^K \mathbb{E} \left( B \cdot N_{T,k} + \sum_{x_t=k, t=1}^T |\mu_k - p_t| \right).$$

Since we are looking at the infimum over all problems, we focus on using Gaussian distributions for all the types. For each distribution  $k \in [K]$ , suppose there are  $n_k$  i.i.d. samples  $x_1, \dots, x_{n_k}$  from  $\nu_k = \mathcal{N}(\mu_k, \sigma^2)$ . In order to lower bound the minimax risk of estimating  $\mu_k$  by any estimator  $\hat{\mu}_k$  under the metric  $\rho = |\mu_k - \hat{\mu}_k|$ , we consider two possible distributions for  $\nu_k$ : fix  $\delta \in [0, \frac{1}{4}]$ ,  $P_{1,k} = \mathcal{N}(\mu'_k + \delta, \sigma^2)$  and  $P_{2,k} = \mathcal{N}(\mu'_k - \delta, \sigma^2)$  where  $\mu'_k \sim \text{Uniform}(\frac{1}{4}, \frac{3}{4})$ . It's easy to see that  $\mu'_k + \delta$  and  $\mu'_k - \delta$  are  $2\delta$  apart.

Via Le Cam's Method, suppose we have an equal probability to choose between  $P_{1,k}$  and  $P_{2,k}$ , we consider all tests  $\Psi : \mathcal{X} \rightarrow \{1, 2\}$  and get

$$\begin{aligned} \inf_{\hat{\mu}_k} \sup_{\nu_k} \mathbb{E} |\mu_k - \hat{\mu}_k| &\geq \delta \inf_{\Psi} \left\{ \frac{1}{2} P_{1,k}(\Psi(x_1, \dots, x_{n_k}) \neq 1) + \frac{1}{2} P_{2,k}(\Psi(x_1, \dots, x_{n_k}) \neq 2) \right\} \\ &= \frac{\delta}{2} \left[ 1 - \|P_{1,k}^{n_k} - P_{2,k}^{n_k}\|_{TV} \right] \end{aligned}$$

where  $P_{i,k}^{n_k}$  is the product distribution for  $i = 1, 2$  and  $\|P_{1,k}^{n_k} - P_{2,k}^{n_k}\|_{TV}$  is the total variation distance between  $P_{1,k}^{n_k}$  and  $P_{2,k}^{n_k}$ . By Pinsker's inequality and the chain rule of KL-divergence [Vershynin \(2018\)](#),

$$\|P_{1,k}^{n_k} - P_{2,k}^{n_k}\|_{TV}^2 \leq \frac{1}{2} D_{KL}(P_{1,k}^{n_k} \| P_{2,k}^{n_k}) = \frac{n}{2} D_{KL}(P_{1,k}^{n_k} \| P_{2,k}^{n_k}) = \frac{n}{2} \cdot \frac{(2\delta)^2}{2\sigma^2} = \frac{n\delta^2}{\sigma^2}$$

Thus  $\|P_{1,k}^{n_k} - P_{2,k}^{n_k}\|_{TV} \leq \frac{\sqrt{n_k\delta}}{\sigma}$ . Taking  $\delta = \frac{\sigma}{2\sqrt{n_k}}$  guarantees that  $\|P_{1,k}^{n_k} - P_{2,k}^{n_k}\| \leq \frac{1}{2}$ . Thus

$$\inf_{\hat{\mu}_k} \sup_{\nu_k} \mathbb{E} |\mu_k - \hat{\mu}_k| \geq \frac{\delta}{2} \left(1 - \frac{1}{2}\right) = \frac{\delta}{4} = \frac{\sigma}{8\sqrt{n_k}}.$$

Hence when considering all  $K$  types, we get

$$\inf \sup \mathbb{E} L_T \geq \sum_{k=1}^K \mathbb{E} \left( BN_{T,k} + M_{T,k} \cdot \frac{\sigma}{8\sqrt{N_{T,k}}} \right)$$

We tune the values of  $N_{T,k}$  by letting  $BN_{T,k} = \frac{\sigma M_{T,k}}{8\sqrt{N_{T,k}}}$ , so  $N_{T,k} = \left(\frac{\sigma}{8B}\right)^{\frac{2}{3}} M_{T,k}^{\frac{2}{3}}$ . Thus we get

$$\inf \sup \mathbb{E} L_T \geq \sum_{k=1}^K \frac{1}{4} \sigma^{\frac{2}{3}} B^{\frac{1}{3}} M_{T,k}^{\frac{2}{3}}.$$

### 6.3 Proof of Theorem 5

*Proof.* In this proof, similarly as the idea in Section 6.1, we will first decompose the prediction error into rounds that are labeled ( $\ell_t = 1$ ) and rounds that are not labeled ( $\ell_t = 0$ ) and then bound each part separately. Finally, we will bound  $N_T$  for the SE kernel and Matérn kernel separately to bound the labeling cost. Recall the definition of the cumulative loss,  $L_T$  is given by,

$$L_T = BN_T + \sum_{t=1}^T |f(x_t) - \mu_t(x_t)|. \quad (14)$$

where  $p_t = \mu_t(x_t)$  in the GP setting. Theorem 6 in [Srinivas et al. \(2010\)](#) states that: let  $\delta \in (0, 1)$  and  $\beta_t = 2\|f\|_k^2 + 300\gamma_t \ln^3(t/\delta)$ ,

$$\mathbb{P} \left\{ \forall T, \forall x \in \mathcal{X}, |\mu_T(\mathbf{x}) - f(\mathbf{x})| \leq \beta_{T+1}^{1/2} \sigma_T(\mathbf{x}) \right\} \geq 1 - \delta.$$

Thus with probability at least  $1 - \delta$ , (14) is upper bounded by

$$L_T \leq BN_T + \underbrace{\sum_{t=1}^T \beta_{t+1}^{1/2} \sigma_t(x_t)}_{B_1} = \underbrace{BN_T}_{B_1} + \underbrace{\sum_{\ell_t=1} \beta_{t+1}^{1/2} \sigma_t(x_t)}_{B_2} + \underbrace{\sum_{\ell_t=0} \beta_{t+1}^{1/2} \sigma_t(x_t)}_{B_3}, \quad (15)$$

where the equality comes from the decomposition of labeled rounds and not labeled rounds. We consider bounding the prediction errors  $B_2$  and  $B_3$  first. To bound part  $B_2$ , we follow the same idea and use part of the results as Lemma 5.4 in (Srinivas et al., 2010),

$$\begin{aligned}\beta_{t+1}\sigma_t^2(x_t) &\leq \beta_{T+1}\sigma^2(\sigma^{-2}\sigma_t^2(x_t)) \\ &\leq \beta_{T+1}\sigma^2 C_1 \log(1 + \sigma^{-2}\sigma_t^2(x_t))\end{aligned}$$

with  $C_1 = \sigma^{-2}/\log(1 + \sigma^{-2}) \geq 1$ , since  $s^2 \leq C_1 \log(1 + s^2)$  for  $s \in [0, \sigma^{-2}]$ , and  $\sigma^{-2}\sigma_t^2(x_t) \leq \sigma^{-2}k(x_t, x_t) \leq \sigma^{-2}$  where  $k(x_t, x_t)$  is the kernel. Summarized over all the labeled rounds, we get

$$\sum_{\ell_t=1} \beta_{t+1}\sigma_t^2(x_t) \leq \sum_{\ell_t=1} \beta_{T+1}\sigma^2 C_1 \log(1 + \sigma^{-2}\sigma_t^2(x_t))$$

Recall the definition of MIG (8), we get

$$\sum_{\ell_t=1} \beta_{T+1}\sigma^2 C_1 \log(1 + \sigma^{-2}\sigma_t^2(x_t)) \leq \frac{2\beta_{T+1}}{\log(1 + \sigma^{-2})} \gamma_{N_T}.$$

By the Cauchy-Schwarz inequality,

$$\sum_{\ell_t=1} \beta_{t+1}^{1/2} \sigma_t(x_t) \leq \sqrt{N_T \sum_{\ell_t=1} \beta_{t+1} \sigma_t^2(x_t)} \leq \sqrt{\frac{2\beta_{T+1} \gamma_{N_T} N_T}{\log(1 + \sigma^{-2})}}. \quad (16)$$

For parts  $B_1$  and  $B_3$ , we analyze the SE kernel and Matérn kernel separately since different kernels have different threshold functions.

### Squared Exponential (SE) kernel

Since we decide not to label when  $\sigma_{t-1}(x_t) \leq \tau(t)$  and recall that the threshold function for SE kernel is  $\tau(t) = \sqrt{2\sigma^2} B^{\frac{1}{d+3}} t^{-\frac{1}{d+3}}$ , we get to bound part  $B_3$ :

$$\begin{aligned}\sum_{\ell_t=0} \beta_{t+1}^{1/2} \sigma_t(x_t) &\leq \beta_{T+1}^{1/2} \sum_{\ell_t=0} \sigma_{t-1}(x_t) && \text{(by definition of } \beta \text{ and (7))} \\ &\leq \beta_{T+1}^{1/2} \int_{t=0}^T \sqrt{2\sigma^2} B^{\frac{1}{d+3}} t^{-\frac{1}{d+3}} dt \\ &= \frac{\beta_{T+1}^{1/2} \sqrt{2\sigma^2} B^{\frac{1}{d+3}} (d+3)}{d+2} T^{\frac{d+2}{d+3}}.\end{aligned}$$

In order to bound part  $B_1$ , we discretize  $\mathcal{X} = [0, 1]^d$  into  $K^d$  equal-size hypercubes  $A_1, \dots, A_{K^d}$  whose side lengths are  $\frac{1}{K}$ .  $\forall A_i$ , their  $L_2$  diameters are  $\text{diam}(A_i) = \frac{\sqrt{d}}{K}$ . Let  $N_{t,k}$  denote the number of labels after time  $t$  in  $A_k$  and  $M_{t,k}$  denote the number of points after time  $t$  in  $A_k$ .  $\forall k \in \{1, \dots, K^d\}$  and  $\forall u \in \mathbb{R}_+$ ,

$$N_{T,k} \leq u + \sum_{t \in S_k} \mathbb{1}_{\{\sigma_{t-1}(x_t) > \tau(t)\}} \quad (17)$$

where  $S_k = \{t : x_t \in A_k, N_{t-1,k} \geq u\}$ . Let  $u = B^{-\frac{2}{d+3}} T^{\frac{2}{d+3}}$  and  $K = \sqrt{\frac{d/l^2}{\sigma^2}} B^{-\frac{1}{d+3}} T^{\frac{1}{d+3}}$ , where  $l$  is from (6). Using the results from Lemma 8 and we get

$$\begin{aligned}\sigma_{t-1}^2(x_t) &\leq \frac{d/l^2}{K^2} + \frac{\sigma^2}{N_{t-1,k}} \\ &\leq \frac{d/l^2}{K^2} + \frac{\sigma^2}{u} && \text{(since } N_{t-1,k} \geq u\text{)} \\ &\leq \frac{2\sigma^2}{u} && \text{(by definitions of } K \text{ and } u\text{)} \\ &= \frac{2\sigma^2}{B^{-\frac{2}{d+3}} T^{\frac{2}{d+3}}},\end{aligned}$$

which indicates  $\sigma_{t-1}(x_t) \leq \sqrt{2\sigma^2} B^{\frac{1}{d+3}} T^{-\frac{1}{d+3}} \leq \tau(t)$ . Thus the terms inside the summation of (17) becomes 0 and we can bound  $N_T$  by,

$$N_T \leq K^d u = \left( \sqrt{\frac{d/l^2}{\sigma^2}} \right)^d B^{-\frac{d+2}{d+3}} T^{\frac{d+2}{d+3}}.$$

Combining the results for part  $B_1$ ,  $B_2$  and  $B_3$  to (15), we get

$$\begin{aligned} L_T &\leq \left( \sqrt{\frac{d/l^2}{\sigma^2}} \right)^d B^{\frac{1}{d+3}} T^{\frac{d+2}{d+3}} + \sqrt{\frac{2\beta_{T+1}\gamma_{N_T}N_T}{\log(1+\sigma^{-2})}} + \frac{\beta_{T+1}^{1/2}\sqrt{2\sigma}(d+3)}{d+2} B^{\frac{1}{3+d}} T^{\frac{d+2}{d+3}} \\ &= \left( C_1 + \frac{\beta_{T+1}^{1/2}\sqrt{2\sigma}(d+3)}{d+2} \right) B^{\frac{1}{3+d}} T^{\frac{d+2}{d+3}} + \sqrt{\frac{2\beta_{T+1}}{\log(1+\sigma^{-2})}} \gamma_{C_1 B^{-\frac{d+2}{d+3}} T^{\frac{d+2}{d+3}}} C_1 B^{-\frac{d+2}{d+3}} T^{\frac{d+2}{d+3}} \end{aligned}$$

where  $C_1 = \left( \sqrt{\frac{d/l^2}{\sigma^2}} \right)^d$ .

### Matérn Kernel

Recall the threshold function for Matérn kernel,  $\tau(t) = \sqrt{2\sigma^2} B^{\frac{1}{2d+3}} t^{-\frac{1}{2d+3}}$ , we can bound part  $B_3$  by,

$$\begin{aligned} \sum_{\ell_t=0} \beta_{t+1}^{1/2} \sigma_t(x_t) &\leq \beta_{T+1}^{1/2} \sum_{\ell_t=0} \sigma_{t-1}(x_t) && \text{(by definition of } \beta \text{ and (7))} \\ &\leq \beta_{T+1}^{1/2} \int_{t=0}^T \sqrt{2\sigma^2} B^{\frac{1}{2d+3}} t^{-\frac{1}{2d+3}} dt \\ &= \frac{\beta_{T+1}^{1/2} \sqrt{2\sigma^2} B^{\frac{1}{2d+3}} (2d+3)}{2d+2} T^{\frac{2d+2}{2d+3}}. \end{aligned}$$

We follow the same discretization argument in SE kernel but let  $u = B^{-\frac{2}{2d+3}} T^{\frac{2}{2d+3}}$  and  $K = \frac{2L_M\sqrt{d}}{\sigma^2} B^{-\frac{2}{2d+3}} T^{\frac{2}{2d+3}}$ , where  $L_M$  is from (6). By Lemma 8, we get

$$\begin{aligned} \sigma_{t-1}^2(x_t) &\leq \frac{2L_M\sqrt{d}}{K} + \frac{\sigma^2}{N_{t-1,k}} \\ &\leq \frac{2L_M\sqrt{d}}{K} + \frac{\sigma^2}{u} \\ &\leq \frac{2\sigma^2}{u} \\ &= \frac{2\sigma^2}{B^{-\frac{2}{2d+3}} T^{\frac{2}{2d+3}}} \end{aligned}$$

which indicates  $\sigma_{t-1}(x_t) \leq \sqrt{2\sigma^2} B^{\frac{1}{2d+3}} T^{-\frac{1}{2d+3}} \leq \tau(t)$ . Thus the terms inside the summation of (17) becomes 0 and we can bound  $N_T$  by,

$$N_T \leq K^d u = \left( \frac{2L_M\sqrt{d}}{\sigma^2} \right)^d B^{-\frac{2d+2}{2d+3}} T^{\frac{2d+2}{2d+3}}$$

Combining the results for part  $B_1$ ,  $B_2$ , and  $B_3$  to (15), we get

$$\begin{aligned} L_T &\leq \left( \frac{2L_M\sqrt{d}}{\sigma^2} \right)^d B^{\frac{1}{2d+3}} T^{\frac{2d+2}{2d+3}} + \sqrt{\frac{2\beta_{T+1}\gamma_{N_T}N_T}{\log(1+\sigma^{-2})}} + \frac{\beta_{T+1}^{1/2}\sqrt{2\sigma}(2d+3)}{2d+2} B^{\frac{1}{2d+3}} T^{\frac{2d+2}{2d+3}} \\ &= \left( C_2 + \frac{\beta_{T+1}^{1/2}\sqrt{2\sigma}(2d+3)}{2d+2} \right) B^{\frac{1}{3+2d}} T^{\frac{2d+2}{2d+3}} + \sqrt{\frac{2\beta_{T+1}}{\log(1+\sigma^{-2})}} \gamma_{C_2 B^{-\frac{2d+2}{2d+3}} T^{\frac{2d+2}{2d+3}}} C_2 B^{-\frac{2d+2}{2d+3}} T^{\frac{2d+2}{2d+3}} \end{aligned}$$

where  $C_2 = \left( \frac{2L_M\sqrt{d}}{\sigma^2} \right)^d$ .

□

## 7 Values of $\lambda$

Table 1, 2, 3, 4 provide the values of  $\lambda$  used for Algorithm 1 and Algorithm 2 in all the experiments.

B = 10, uniform	B = 100, uniform	B = 10, uniform	B = 100, uniform
0.5	0.25	2	0.75
B = 10, lopsided	B = 100, lopsided	B = 10, lopsided	B = 100, lopsided
0.5	0.25	2	0.5

Table 1:  $\lambda$  used for Fig 2.  $K = 10$  for columns 1 and 2.  $K = 100$  for columns 3 and 4.

	B = 10, uniform	B = 100, uniform	B = 10, lopsided	B = 100, lopsided
Algorithm 2	1	1	0.5	0.5
Algorithm 1	10	10	10	10

Table 2:  $\lambda$  used for Fig 3, Column 1 and 2 are used for the Branin function. Columns 3 and 4 are used for the Hartmann function.

	B = 10, uniform	B = 10, lopsided	B = 100, lopsided
Algorithm 2	5	5	5
Algorithm 1	5	5	5

Table 3:  $\lambda$  used for Fig 4.

	B = 1	B = 10	B = 100
Algorithm 2	1	1	1
Algorithm 1	50	30	10

Table 4:  $\lambda$  used for Fig 5

## 8 Plots of error

In this section, we provide the average prediction error plots for all the experiments.

The prediction error is defined as the absolute difference between the true function values of  $x_t$  and our prediction, which is  $|f(x_t) - p_t|$  at each round  $t$ . Figures 6, 7, 8, 9 provide the average prediction error plots for all the experiments. Across all the experiments, Algorithm 1 maintains a low prediction error compared to the baselines in the  $K$  discrete setting and Algorithm 2 maintains a low prediction error compared to the baselines in the continuous setting.

The prediction errors for the naive implementation of Algorithm 1 are relatively large compared to other methods in the continuous case, which are normal since we expect the naive prediction by sample mean in the continuous case to perform badly. However, when  $B = 100$  for the Parkinsons dataset, naive Algorithm 1 has the lowest average loss compared to the other three methods as shown in Figure 4. This indicates that when the cost of experimentation is relatively large and we can tolerate large prediction error, we can choose naive Algorithm 1 in the continuous case to reduce the total cost.

For the Parkinsons dataset, the prediction errors for Algorithm 2, *Random Select*, and *VAR-UNCERTAINTY* stay low at the beginning and then increase over time, this is because we allow the three methods to label at the first  $5d$  rounds, where  $d$  is the dimension of the input data, so they have a low prediction error. Then after this initialization period, for data points that are not labeled but have labels very distinct from the predicted labels, the prediction errors can be large.

## Active Cost-aware Labeling of Streaming Data

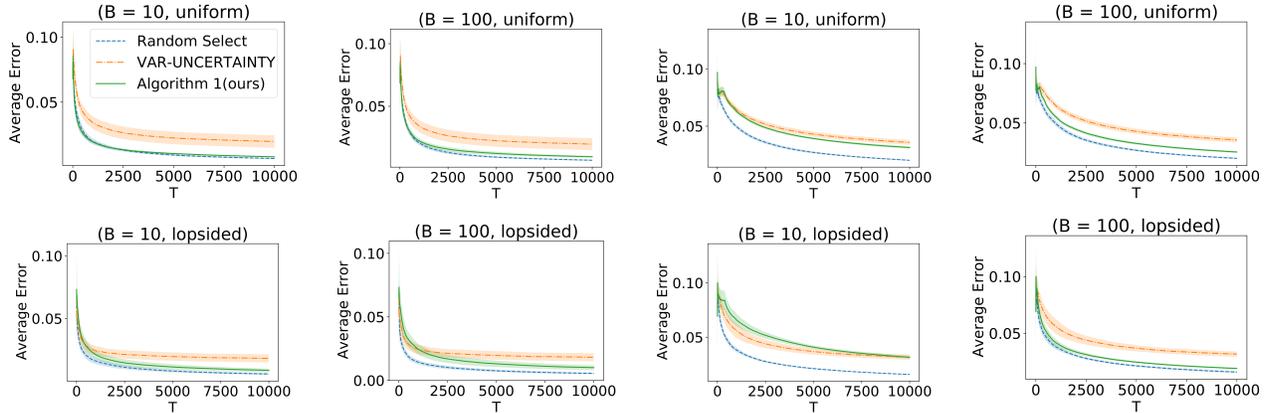


Figure 6: Error plots comparing the three methods for  $K$  discrete setting of §4.1.  $K = 10$  for column 1 and 2.  $K = 100$  for columns 3 and 4. The results are averaged over 10 trials, and we report the 95% the confidence interval of the standard error.

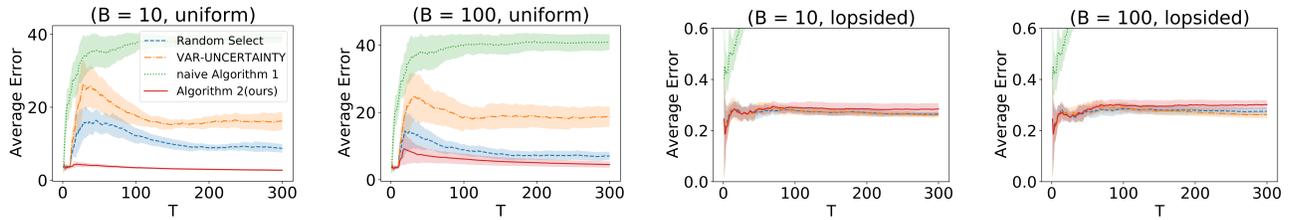


Figure 7: Error plots comparing the four methods for synthetic functions in §4.2. Columns 1 and 2 show the results when the arrival pattern is uniform for the Branin function. Columns 3 and 4 show the results when the arrival pattern is lopsided for the Hartmann function. The results are averaged over 10 trials, and we report the 95% confidence interval of the standard error. The results of Algorithm 1 are poor for columns 3 and 4, so we limit the  $y$ -axis to focus on the rest methods.

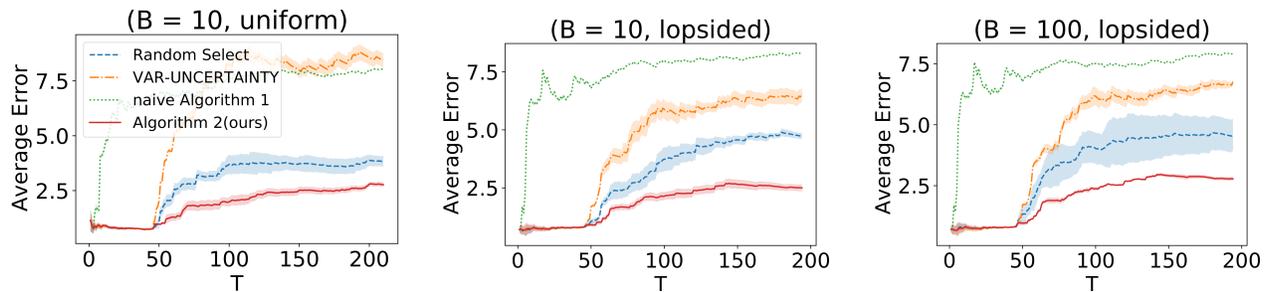


Figure 8: Error plots comparing the four methods for Parkinsons dataset. The results are averaged over 5 trials and we report the 95% confidence interval of the standard error.

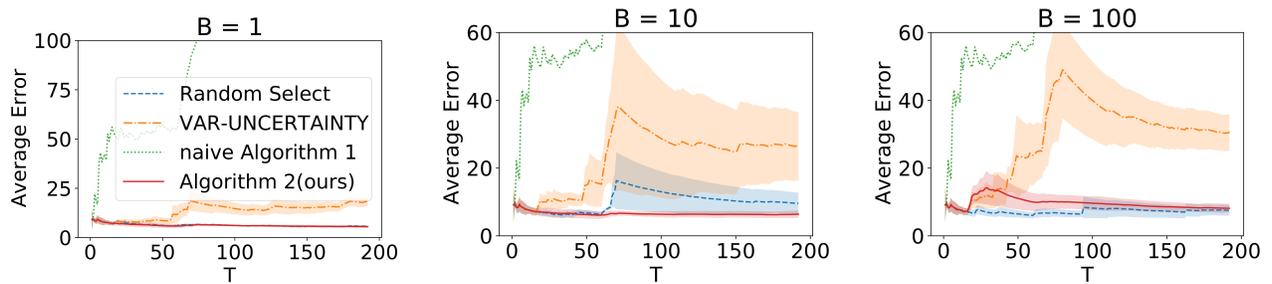


Figure 9: Error plots comparing the four methods for the supernova dataset. The results are averaged over 5 trials and we report the 95% confidence interval of the standard error. The naive Algorithm 1 performs poorly so we limit the  $y$ -axis to focus on the rest methods.

## 9 Ablation study

In this section, we compare the performance of Algorithm 1 and Algorithm 2 under different choices of  $\lambda$ .

Figure 10 shows the results of average loss with different values of  $\lambda$  for the experiment when  $K = 10$ ,  $B = 10$  and the arrival pattern is uniform. Increasing the value of  $\lambda$  can decrease the average loss over time but it may increase the average prediction error since we may not get enough labeling data to do the prediction.

Figure 11 shows the results of average loss with different values of  $\lambda$  when  $B = 10$  and the arrival pattern is uniform for the Branin function.  $\lambda = 1$  achieves the lowest average loss over time since  $\lambda$  being small ( $\lambda = 0.5$ ) limits the number of labeling thus the prediction error is high and  $\lambda$  being large ( $\lambda \geq 1.5$ ) labels a lot thus the total cost of labeling is high.

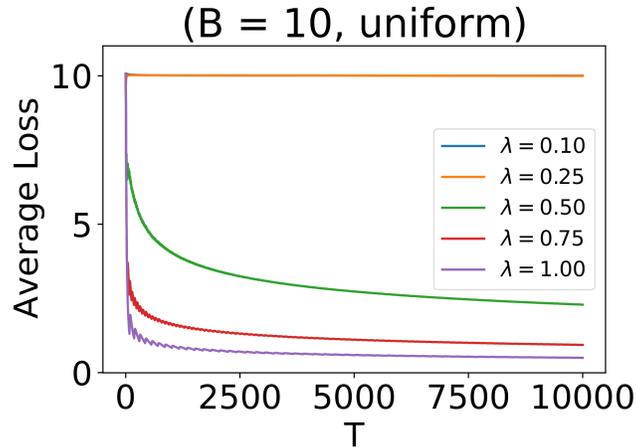


Figure 10: Average loss plots comparing the effects of different  $\lambda$  for Algorithm 1 when  $K = 10$ ,  $B = 10$  for uniform arrival patterns. The result for  $\lambda = 0.1$  is hidden by the result for  $\lambda = 0.25$  since a very small value of  $\lambda$  prohibits labeling at all. The results are averaged over 10 trials and we report the 95% confidence interval.

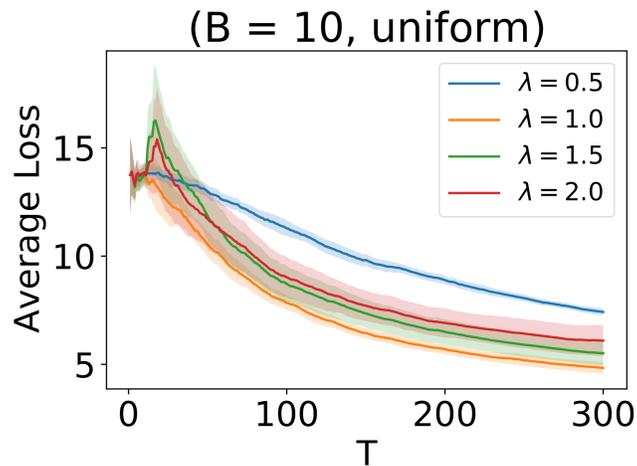


Figure 11: Average loss plots comparing the effects of different  $\lambda$  for Algorithm 2 when  $B = 10$  and the arrival pattern is uniform for Branin function. The results are averaged over 10 trials and we report the 95% confidence interval.

## 10 Supplementary tools

**Theorem 6** (Power Means inequality [Wainwright \(2019\)](#)). *For any  $p < q$ , the following inequality holds*

$$\left( \sum_{i=1}^n w_i x_i^p \right)^{1/p} \leq \left( \sum_{i=1}^n w_i x_i^q \right)^{1/q}$$

where  $w_i \in [0, 1]$  and  $\sum_{i=1}^n w_i = 1$ .

**Theorem 7** (Le Cam's Method [Le Cam \(1960\)](#)). *For any distributions  $P_1$  and  $P_2$  on  $\mathcal{X}$ , we have*

$$\inf_{\Psi} \{P_1(\Psi(X) \neq 1) + P_2(\Psi(X) \neq 2)\} = 1 - \|P_1 - P_2\|_{TV},$$

where the infimum is taken over all tests  $\Psi : \mathcal{X} \rightarrow \{1, 2\}$ .

**Lemma 8** (Lemma 12 in [Kandasamy et al. \(2016\)](#)). *Let  $f \sim \mathcal{GP}(\mathbf{0}, \kappa)$ ,  $f : \mathcal{X} \rightarrow \mathbb{R}$  and we observe  $y = f(x) + \epsilon$  where  $\epsilon \sim \mathcal{N}(0, \eta^2)$ . Let  $A \in \mathcal{X}$  such that its  $L_2$  diameter  $\text{diam}(A) \leq D$ . Say we have  $n$  queries  $(x_t)_{t=1}^n$  of which  $s$  points are in  $A$ . Then the posterior variance of the GP,  $\kappa'(x, x)$  at any  $x \in A$  satisfies*

$$\kappa'(x, x) \leq \begin{cases} D^2/l^2 + \frac{\eta^2}{s} & \text{if } \kappa \text{ is the SE kernel,} \\ 2L_M D + \frac{\eta^2}{s} & \text{if } \kappa \text{ is the Matérn kernel.} \end{cases}$$

Here,  $l$  is the length scale hyperparameter of the SE kernel (6), and  $L_M$  is a Lipschitz constant of the Matérn kernel.