

---

# Budgeted stream-based active learning via adaptive submodular maximization (with supplementary materials)

---

**Kaito Fujii**

Kyoto University

JST, ERATO, Kawarabayashi Large Graph Project

fujii@ml.ist.i.kyoto-u.ac.jp

**Hisashi Kashima**

Kyoto University

kashima@i.kyoto-u.ac.jp

## Abstract

Active learning enables us to reduce the annotation cost by adaptively selecting unlabeled instances to be labeled. For pool-based active learning, several effective methods with theoretical guarantees have been developed through maximizing some utility function satisfying adaptive submodularity. In contrast, there have been few methods for stream-based active learning based on adaptive submodularity. In this paper, we propose a new class of utility functions, policy-adaptive submodular functions, which includes many existing adaptive submodular functions appearing in real world problems. We provide a general framework based on policy-adaptive submodularity that makes it possible to convert existing pool-based methods to stream-based methods and give theoretical guarantees on their performance. In addition we empirically demonstrate their effectiveness by comparing with existing heuristics on common benchmark datasets.

## 1 Introduction

Active learning is a problem setting for sequentially selecting unlabeled instances to be labeled, and it has been studied with much practical interest as an efficient way to reduce the annotation cost. One of the most popular settings of active learning is the pool-based one, in which the learner is given the entire set of unlabeled instances in advance, and iteratively selects an instance to be labeled next. The stream-based setting, which we deal with in this paper, is another important setting of active learning, in which the entire set of unlabeled instances are hidden initially, and presented one by one to the learner. This setting also has many real world applications, for example, sentiment analysis of web stream data [26], spam filtering [25], part-of-speech tagging [10], and video surveillance [23].

Adaptive submodularity [19] is an adaptive extension of submodularity, a natural diminishing return condition. It provides a framework for designing effective algorithms for several adaptive problems including pool-based active learning. For instance, the ones for noiseless active learning [19, 21] and the ones for noisy active learning [20, 9, 8] have been developed in recent years. Not only they have strong theoretical guarantees on their performance, but they perform well in practice compared with existing widely-used heuristics.

In spite of its considerable success in the pool-based setting, little is known about benefits of adaptive submodularity in the stream-based setting. This paper answers the question: is it possible to construct algorithms for stream-based active learning based on adaptive submodularity? We propose a general framework for creating stream-based algorithms from existing pool-based algorithms.

In this paper, we tackle the problem of stream-based active learning with a limited budget for making queries. The goal is collecting an informative set of labeled instances from a data stream of a certain length. The stream-based active learning problem has been typically studied in two settings:

the stream setting and the secretary setting, which correspond to memory constraints and timing constraints respectively; we treat both in this paper.

We formalize these problems as the adaptive stochastic maximization problem in the stream or secretary setting. For solving this problem, we propose a new class of stochastic utility functions: *policy-adaptive submodular functions*, which is another adaptive extension of submodularity. We prove this class includes many existing adaptive submodular functions used in various applications. Assuming the objective function satisfies policy-adaptive submodularity, we propose simple methods for each problem, and give theoretical guarantees on their performance in comparison to the optimal pool-based method. Experiments conducted on benchmark datasets show the effectiveness of our methods compared with several heuristics. Due to our framework, many algorithms developed in the pool-based setting can be converted to the stream-based setting.

In summary, our main contributions are the following:

- We provide a general framework that captures budgeted stream-based active learning and other applications.
- We propose a new class of stochastic utility functions, *policy-adaptive submodular functions*, which is a subclass of the adaptive submodular functions, and prove this class includes many existing adaptive submodular functions in real world problems.
- We propose two simple algorithms, AdaptiveStream and AdaptiveSecretary, and give theoretical performance guarantees on them.

## 2 Problem Settings

In this section, we first describe the general framework, then illustrate applications including stream-based active learning.

### 2.1 Adaptive Stochastic Maximization in the Stream and Secretary Settings

Here we specify the problem statement. This problem is a generalization of budgeted stream-based active learning and other applications.


Let  $V = \{v_1, \dots, v_n\}$  denote the entire set of  $n$  items, and each item  $v_i$  is in a particular state out of the set  $\mathcal{Y}$  of possible states. Denote by  $\phi : V \rightarrow \mathcal{Y}$  a realization of the states of the items. Let  $\Phi$  be a random realization, and  $Y_i$  a random variable representing the state of each item  $v_i$  for  $i = 1, \dots, n$ , i.e.,  $Y_i = \Phi(v_i)$ . Assume that  $\phi$  is generated from a known prior distribution  $p(\phi)$ . Suppose the state  $Y_i$  is revealed when  $v_i$  is selected. Let  $\psi_A : A \rightarrow \mathcal{Y}$  denote the partial realization obtained after the states of items  $A \subseteq V$  are observed. Note that a partial realization  $\psi_A$  can be regarded as the set of observations  $\{(s, \psi_A(s)) \mid s \in A\} \subseteq V \times \mathcal{Y}$ .

We are given a set function<sup>1</sup>  $f : 2^{V \times \mathcal{Y}} \rightarrow \mathbb{R}_{\geq 0}$  that defines the utility of observations made when some items are selected. Consider iteratively selecting an item to observe its state and aiming to make observations of high utility value. A policy  $\pi$  is some decision tree that represents a strategy for adaptively selecting items. Formally it is defined to be a partial mapping that determines an item to be selected next from the observations made so far. Given some budget  $k \in \mathbb{Z}_{>0}$ , the goal is constructing a policy  $\pi$  maximizing  $\mathbb{E}_{\Phi}[f(\psi(\pi, \Phi))]$  subject to  $|\psi(\pi, \phi)| \leq k$  for all  $\phi$  where  $\psi(\pi, \phi)$  denotes the observations obtained by executing policy  $\pi$  under realization  $\phi$ .


This problem has been studied mainly in the pool-based setting, where we are given the entire set  $V$  from the beginning and adaptively observe the states of items in any order. In this paper we tackle the stream-based setting, where the items are hidden initially and arrive one by one. The stream-based setting arises in two kinds of scenarios: one is the stream setting<sup>2</sup>, in which we can postpone deciding whether or not to select an item by keeping it in a limited amount of memory, and at any time observe the state of the stored items. The other is the secretary setting, in which we must decide

<sup>1</sup>In the original definition of stochastic utility functions [19], the objective value depends not only on the partial realization  $\psi$ , but also on the realization  $\phi$ . However, given such  $f : 2^V \times \mathcal{Y}^V \rightarrow \mathbb{R}_{\geq 0}$ , we can redefine  $\tilde{f} : 2^{V \times \mathcal{Y}} \rightarrow \mathbb{R}_{\geq 0}$  as  $\tilde{f}(\psi_A) = \mathbb{E}_{\Phi}[f(A, \Phi) \mid \Phi \sim p(\Phi \mid \psi_A)]$ , and it does not critically change the overall discussion in our problem settings. Thus for notational convenience, we use the simpler definition.

<sup>2</sup>In this paper, “stream-based setting” and “stream setting” are distinguished.



(a) A policy tree for the pool-based setting



(b) A policy tree for the stream-based setting

Figure 1: Examples of a pool-based policy and a stream-based policy in the case of  $\mathcal{Y} = \{+1, -1\}$ . (a) A pool-based policy can select items in an arbitrary order. (b) A stream-based policy must select items under memory or timing constraints taking account of only items that arrived so far.

immediately whether or not to select an item at each arrival. In both settings we assume the items arrive in a random order. The comparison of policies for the pool-based and stream-based settings is indicated in Figure 1.

## 2.2 Budgeted Stream-based Active Learning

We consider a problem setting called *Bayesian active learning*. Here  $V$  represents the set of instances,  $Y_1, \dots, Y_n$  the initially unknown labels of the instances, and  $\mathcal{Y}$  the set of possible labels.

Let  $\mathcal{H}$  denote the set of candidates for the randomly generated true hypothesis  $H$ , and  $p_H$  denote a prior probability over  $\mathcal{H}$ . When observations of the labels are noiseless, every hypothesis  $h \in \mathcal{H}$  represents a particular realization, i.e.,  $h$  corresponds to some  $\phi \in \mathcal{Y}^V$ . When observations are noisy, the probability distribution  $\mathbb{P}[Y_1, \dots, Y_n | H = h]$  of the labels is not necessarily deterministic for each  $h \in \mathcal{H}$ . In both cases, we can iteratively select an instance and query its label to the annotation oracle. The objective is to determine the true hypothesis or one whose prediction error is small. Both the pool-based and stream-based settings have been extensively studied. The stream-based setting contains the stream and secretary settings, both of which have a lot of real world applications.

A common approach for devising a pool-based algorithm is designing some utility function that represents the informativeness of a set of labeled instances, and greedily selecting the instance maximizing this utility in terms of the expected value. We introduce the utility into stream-based active learning, and aim to collect  $k$  labeled instances of high utility where  $k \in \mathbb{Z}_{>0}$  is the budget on the number of queries. While most of the theoretical results for stream-based active learning are obtained assuming the data stream is infinite, we assume the length of the total data stream is given in advance.

## 2.3 Other Applications

We give a brief sketch of two examples that can be formalized as the adaptive stochastic maximization problem in the secretary setting. Both are variations for streaming data of the problems first proposed by Golovin and Krause [19].

One is adaptive viral marketing whose aim is spreading information about a new product through social networks. In this problem we adaptively select  $k$  people to whom a free promotional sample of the product is offered so as to let them recommend the product to their friends. We cannot know if he recommends the product before actually offering a sample to each. The objective is maximizing the number of people that information of the product reaches. There arise some situations where people come sequentially, and at each arrival we must decide whether or not to offer a sample to them.

Another is adaptive sensor placement. We want to adaptively place  $k$  unreliable sensors to cover the information obtained by them. The informativeness of each sensor is unknown before its deploy-

ment. We can consider the cases where the timing of placing sensors at each location is restricted for some reasons such as transportation cost.

### 3 Policy-Adaptive Submodularity

In this section, we discuss conditions satisfied by the utility functions of adaptive stochastic maximization problems.

Submodularity [17] is known as a natural diminishing return condition satisfied by various set functions appearing in a lot of applications, and adaptive submodularity was proposed by Golovin and Krause [19] as an adaptive extension of submodularity. Adaptive submodularity is defined as the diminishing return property about the expected marginal gain of a single item, i.e.,  $\Delta(s|\psi_A) \geq \Delta(s|\psi_B)$  for any partial realization  $\psi_A \subseteq \psi_B$  and item  $s \in V \setminus B$ , where

$$\Delta(s|\psi) = \mathbb{E}_\Phi[f(\psi \cup \{(s, \Phi(s))\}) - f(\psi) \mid \Phi \sim p(\Phi|\psi)].$$

Similarly, adaptive monotonicity, an adaptive analog of monotonicity, is defined to be  $\Delta(s|\psi_A) \geq 0$  for any partial realization  $\psi_A$  and item  $s \in V$ . It is known that many utility functions used in the above applications satisfy the adaptive submodularity and the adaptive monotonicity. In the pool-based setting, greedily selecting the item of the maximal expected marginal gain yields  $(1 - 1/e)$ -approximation if the objective function is adaptive submodular and adaptive monotone [19].

Here we propose a new class of stochastic utility functions, *policy-adaptive submodular functions*. Let  $\text{range}(\pi)$  denote the set containing all items that  $\pi$  selects for some  $\phi$ , and we define policy-adaptive submodularity as the diminishing return property about the expected marginal gain of any policy as follows.

**Definition 3.1** (Policy-adaptive submodularity). *A set function  $f : 2^V \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$  is policy-adaptive submodular with respect to a prior distribution  $p(\phi)$ , or  $(f, p)$  is policy-adaptive submodular, if  $\Delta(\pi|\psi_A) \geq \Delta(\pi|\psi_B)$  holds for any partial realization  $\psi_A, \psi_B$  and policy  $\pi$  such that  $\psi_A \subseteq \psi_B$  and  $\text{range}(\pi) \subseteq V \setminus B$ , where*

$$\Delta(\pi|\psi) = \mathbb{E}_\Phi[f(\psi \cup \psi(\pi, \Phi)) - f(\psi) \mid \Phi \sim p(\Phi|\psi)].$$

Since a single item can be regarded as a policy selecting only one item, policy-adaptive submodularity is a stricter condition than adaptive submodularity.

Policy-adaptive submodularity is also a natural extension of submodularity. The submodularity of a set function  $f : 2^V \rightarrow \mathbb{R}_{\geq 0}$  is defined as the condition that  $f(A \cup \{s\}) - f(A) \geq f(B \cup \{s\}) - f(B)$  for any  $A \subseteq B \subseteq V$  and  $s \in V \setminus B$ , which is equivalent to the condition that  $f(A \cup P) - f(A) \geq f(B \cup P) - f(B)$  for any  $A \subseteq B \subseteq V$  and  $P \subseteq V \setminus B$ . Adaptive extensions of these conditions are adaptive submodularity and policy-adaptive submodularity respectively. Nevertheless there is a counterexample to the equivalence of adaptive submodularity and policy-adaptive submodularity, which is given in the supplementary materials.

Surprisingly, many existing adaptive submodular functions in applications also satisfy the policy-adaptive submodularity. In active learning, the objective function of generalized binary search [12, 19], EC<sup>2</sup> [20], ALuMA [21], and the maximum Gibbs error criterion [9, 8] are not only adaptive submodular, but policy-adaptive submodular. In other applications including influence maximization and sensor placements, it is often assumed that the variables  $Y_1, \dots, Y_n$  are independent, and the policy-adaptive submodularity always holds in this case. The proofs of these propositions are given in the supplementary materials.

To give the theoretical guarantees for the algorithms introduced in the next section, we assume not only the adaptive submodularity and the adaptive monotonicity, but also the policy-adaptive submodularity. However, our theoretical analyses can still be applied to many applications.

### 4 Algorithms

In this section we describe our proposed algorithms for each of the stream and secretary settings, and state the theoretical guarantees on their performance. The full versions of pseudocodes are given in the supplementary materials.

---

**Algorithm 1** AdaptiveStream algorithm & AdaptiveSecretary algorithm

---

**Input:** A set function  $f : 2^{V \times \mathcal{Y}} \rightarrow \mathbb{R}_{\geq 0}$  and a prior distribution  $p(\phi)$  such that  $(f, p)$  is policy-adaptive submodular and adaptive monotone. The number of items in the entire stream  $n \in \mathbb{Z}_{>0}$ . A budget  $k \in \mathbb{Z}_{>0}$ . Randomly permuted stream of the items, denoted by  $(s_1, \dots, s_n)$ .  
**Output:** Some observations  $\psi_k \subseteq V \times \mathcal{Y}$  such that  $|\psi_k| \leq k$ .  
1: Let  $\psi_0 := \emptyset$ .  
2: **for** each segment  $S_l = \{s_i \mid (l-1)n/k < i \leq ln/k\}$  **do**  
3:   Select an item  $s$  out of  $S_l$  by  
     $\begin{cases} \text{selecting the item of the largest expected marginal gain} & (\text{AdaptiveStream}) \\ \text{applying the classical secretary algorithm} & (\text{AdaptiveSecretary}) \end{cases}$   
4:   Observe the state  $y$  of item  $s$  and let  $\psi_l := \psi_{l-1} \cup \{(s, y)\}$ .  
5: **return**  $\psi_k$  as the solution

---

#### 4.1 Algorithm for the Stream Setting

The main idea of our proposed method is simple: divide the entire stream into  $k$  segments and select the best item from each one. For simplicity, we consider the case where  $n$  is a multiple integer of  $k$ . If  $n$  is not, we can add  $k\lceil \frac{n}{k} \rceil - n$  dummy items with no benefit and prove the same guarantee. Our algorithm first divides the item sequence  $s_1, \dots, s_n$  into  $S_l = \{s_i \mid (l-1)n/k < i \leq ln/k\}$  for  $l = 1, \dots, k$ . In each segment, the algorithm selects the item of the largest expected marginal gain, that is,  $\arg\max\{\Delta(s|\psi_{l-1}) \mid s \in S_l\}$  where  $\psi_{l-1}$  is the partial realization obtained before the  $l$ th segment. This can be implemented with only  $O(1)$  space by storing only the item of the maximal expected marginal gain so far in the current segment. We provide the theoretical guarantee on the performance of this algorithm by utilizing the policy-adaptive submodularity of the objective function.

**Theorem 4.1.** *Suppose  $f : 2^{V \times \mathcal{Y}} \rightarrow \mathbb{R}_{\geq 0}$  is policy-adaptive submodular and adaptive monotone w.r.t. a prior  $p(\phi)$ . Assume the items come sequentially in a random order. For any policy  $\pi$  such that  $|\psi(\pi, \phi)| \leq k$  holds for all  $\phi$ , AdaptiveStream selects  $k$  items using  $O(1)$  space and achieves at least 0.16 times the expected total gain of  $\pi$  in expectation.*

#### 4.2 Algorithm for the Secretary Setting

Though our proposed algorithm for the secretary setting is similar in its approach to the one for the stream setting, it is impossible to select the item of the maximal expected marginal gain from each segment in the secretary setting. Then we use *classical secretary algorithm* [13] as a subroutine to obtain the maximal item at least with some constant probability. The classical secretary algorithm lets the first  $\lfloor n/(ek) \rfloor$  items pass and then selects the first item whose value is larger than all items so far. The probability that this subroutine selects the item of the largest expected marginal gain is at least  $1/e$  at each segment. This algorithm can be viewed as an adaptive version of the algorithm for the monotone submodular secretary problem [3]. We give the guarantee similar to the one for the stream setting.

**Theorem 4.2.** *Suppose  $f : 2^{V \times \mathcal{Y}} \rightarrow \mathbb{R}_{\geq 0}$  is policy-adaptive submodular and adaptive monotone w.r.t. a prior  $p(\phi)$ . Assume the items come sequentially in a random order. For any policy  $\pi$  such that  $|\psi(\pi, \phi)| \leq k$  holds for all  $\phi$ , AdaptiveSecretary selects at most  $k$  items and achieves at least 0.08 times the expected total gain of  $\pi$  in expectation.*

### 5 Overview of Theoretical Analysis

In this section we briefly describe the proofs of Theorem 4.1 and 4.2, and compare our techniques with the previous work. The full proofs are given in the supplementary materials.

The methods used in the proofs of both theorems are almost the same. They consist of two steps: in the first step, we bound the expected marginal gain of each item and in the second step, we take summation of one step marginal gains and derive the overall bound for the algorithms. Though our techniques used in the second step are taken from the previous work [3], the first step contains several novel techniques.

Let  $\Delta_i$  be the expected marginal gain of an item picked from the  $i$ th segment  $S_i$ . First we bound it from below with the difference between the optimal pool-based policy  $\pi_T^*$  for selecting  $k$  items from  $T$  and the policy  $\pi_{i-1}^\sigma$  that encodes the algorithm until  $i - 1$ th step under a permutation  $\sigma$  in which the items arrive. For the non-adaptive setting, the items in the optimal set are distributed among the segments uniformly at random, then we can evaluate  $\Delta_i$  by considering whether  $S_i$  contains an item included in the optimal set [3]. On the other hand, in the adaptive setting, it is difficult to consider how  $\pi_T^*$  is distributed in the unrarried items because the policy is closely related not only to the contained items but also to the order of items. Then we compare  $\Delta_i$  and the marginal gain of  $\pi_T^*$  directly. With the adaptive monotonicity, we obtain  $\Delta_i \geq (1 - \exp(-\frac{k}{k-i+1}))(f_{\text{avg}}(\pi_T^*) - f_{\text{avg}}(\pi_{i-1}^\sigma))/k$  where  $f_{\text{avg}}(\pi) = \mathbb{E}_\Phi[f(\psi(\pi, \Phi))]$ .

Next we bound  $f_{\text{avg}}(\pi_T^*)$  with the optimal pool-based policy  $\pi_V^*$  that selects  $k$  items from  $V$ . For the non-adaptive setting, we can apply a widely-used lemma proved by Feige, Mirrokni, and Vondrák [15]. This lemma provides a bound for the expected value of a randomly deleted subset. To extend this lemma to the adaptive setting, we define a partially deleted policy tree, *grafted policy*, and prove the adaptive version of the lemma with the policy-adaptive submodularity. From this lemma we can obtain the bound  $\mathbb{E}_\sigma[f_{\text{avg}}(\pi_T^*)] \geq (k - i + 1)f_{\text{avg}}(\pi_V^*)/k$ . We also provide an example that shows adaptive submodularity is not enough to prove this lemma.

Summing the bounds for each one-step expected marginal gain until  $l$ th step ( $l$  is specified in the full proof for optimizing the resulting guarantees), we can conclude that our proposed algorithms achieve some constant factor approximation in comparison to the optimal pool-based policy. Though AdaptiveSecretary is the adaptive version of the existing algorithm, our resulting constant factor is a little worse than the original  $(1 - 1/e)/7$  due to the above new analyses.

## 6 Experiments

### 6.1 Experimental Setting

We conducted experiments on budgeted active learning in the following three settings: the pool-based, stream, and secretary settings. For each setting, we compare two methods: one is based on the policy-adaptive submodularity and the other is based on *uncertainty sampling* as baselines. Uncertainty sampling is the most widely-used approach in applications. Selecting random instances, which we call *random*, is also implemented as another baseline that can be used in every setting.

We select ALuMA [21] out of several pool-based methods based on adaptive submodularity, and convert it to the stream and secretary settings with AdaptiveStream and AdaptiveSecretary, which we call *stream submodular* and *secretary submodular* respectively. For comparison, we also implement the original pool-based method, which we call *pool submodular*. Though ALuMA is designed for the noiseless case, there is a modification method that makes its hypotheses sampling more noise-tolerant [7], which we employ. The number of hypotheses sampled at each time is set  $N = 1000$  in all settings.

For the pool-based setting, uncertainty sampling is widely-known as a generic and easy-to-implement heuristic in many applications. This selects the most uncertain instance, i.e., the instance that is closest to the current linear separator. In contrast, there is no standard heuristic for the stream and secretary settings. We apply the same conversion to the pool-based uncertain sampling method as AdaptiveStream and AdaptiveSecretary, i.e., in the stream setting, selecting the most uncertain instance from the segment at each step, and in the secretary setting, running the classical secretary algorithm to select the most uncertain instance at least with probability  $1/e$ . A similar one to this approach in the stream setting is used in some applications [26]. In every setting, we first randomly select 10 instances for the initial training of a classifier and after that, select  $k - 10$  instances with each method. We use the linear SVM trained with instances labeled so far to judge the uncertainty. We call these methods *pool uncertainty*, *stream uncertainty*, *secretary uncertainty* respectively, and use them as baselines.

We conducted experiments on two benchmark datasets, WDBC<sup>3</sup> and MNIST<sup>4</sup>. The WDBC dataset contains 569 instances, each of which consists of 32-dimensional features of cells and their diagnosis

<sup>3</sup>[https://archive.ics.uci.edu/ml/datasets/Breast+Cancer+Wisconsin+\(Diagnostic\)](https://archive.ics.uci.edu/ml/datasets/Breast+Cancer+Wisconsin+(Diagnostic))

<sup>4</sup><http://yann.lecun.com/exdb/mnist/>

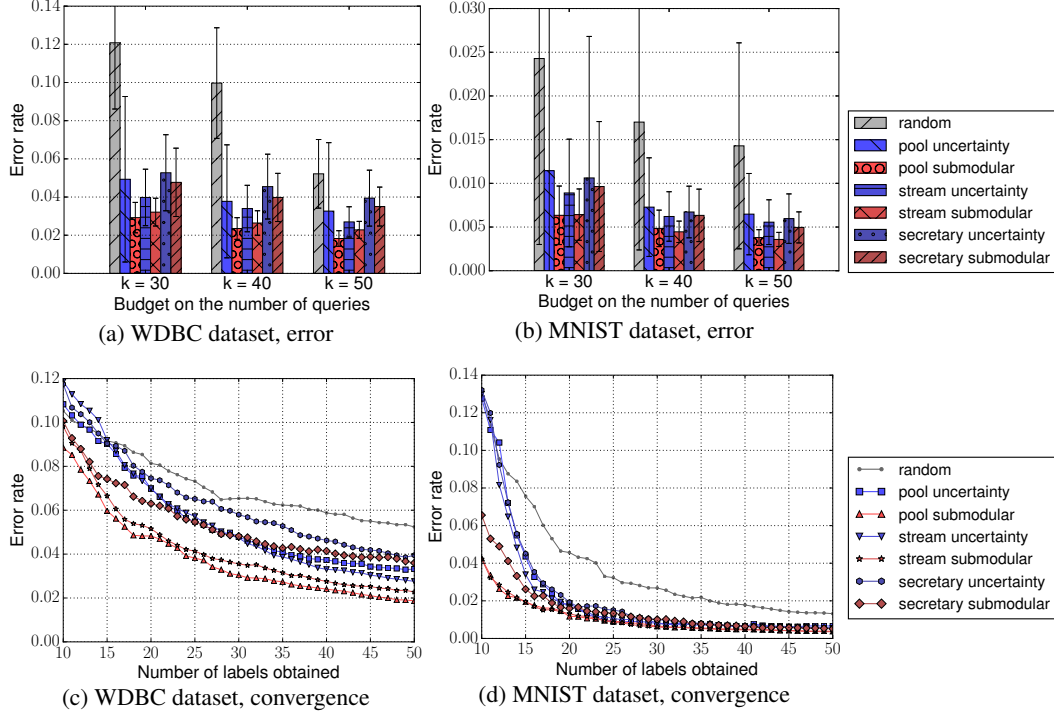


Figure 2: Experimental results

results. From the MNIST dataset, the dataset of handwritten digits, we extract 14780 images of the two classes, 0 and 1, so as to consider the binary classification problem, and apply PCA to reduce its dimensions from 784 to 10. We standardize both datasets so that the values of each feature have zero mean and unit variance.

We evaluate the performance through 100 trials, where at each time an order in which the instances arrive is generated randomly. For all the methods, we calculate the error rate by training linear SVM with the obtained labeled instances and testing with the entire dataset.

## 6.2 Experimental Results

Figure 2(a)(b) illustrate the average error rate achieved by each method with budget  $k = 30, 40, 50$ . Our methods stream submodular and secretary submodular outperform not only random, but also stream uncertainty and secretary uncertainty respectively, i.e., the methods based on policy-adaptive submodularity perform better than the methods based on uncertainty sampling in each of the stream and secretary settings. Moreover, we can observe our methods are stabler than the other methods from the error bars representing the standard deviation.

Figure 2(c)(d) show how the error rate decreases as labels are queried in the case of  $k = 50$ . In both datasets, we can observe the performance of stream submodular is competitive with pool submodular.

## 7 Related Work

**Stream-based active learning.** Much amount of work has been dedicated to devising algorithms for stream-based active learning (also known as *selective sampling*) from both the theoretical and practical aspects. From the theoretical aspects, several bounds on the label complexity have been provided [16, 2, 4], but their interest lies in the guarantees compared to the passive learning, not the optimal algorithm. From the practical aspects, it has been applied to many real world problems such as sentiment analysis of web stream data [26], spam filtering [25], part-of-speech tagging [10], and video surveillance [23], but there is no definitive widely-used heuristic.

Of particular relevance to our work is the one presented by Sabato and Hess [24]. They devised general methods for constructing stream-based algorithms satisfying a budget based on pool-based algorithms, but their theoretical guarantees are bounding the length of the stream needed to emulate the pool-based algorithm, which is a large difference from our work. Das et al. [11] designed the algorithm for adaptively collecting water samples, referring to the submodular secretary problem, but they focused on applications to marine ecosystem monitoring, and did not give any theoretical analysis about its performance.

**Adaptive submodular maximization.** The framework of adaptive submodularity, which is an adaptive counterpart of submodularity, is established by Golovin and Krause [19]. It provides the simple greedy algorithm with the near-optimal guarantees in several adaptive real world problems. Specifically it achieves remarkable success in pool-based active learning. For the noiseless cases, Golovin and Krause [19] described the generalized binary search algorithm [12] as the greedy algorithm for some adaptive submodular function, and improved its approximation factor. Golovin et al. [20] provided an algorithm for Bayesian active learning with noisy observations by reducing it to the equivalence class determination problem. On the other hand, there have been several studies on adaptive submodular maximization in other settings, for example, selecting multiple instances at the same time before observing their states [7], guessing an unknown prior distribution in the bandit setting [18], and maximizing non-monotone adaptive submodular functions [22].

**Submodular maximization in the stream and secretary settings.** Submodular maximization in the stream setting, called *streaming submodular maximization*, has been studied under several constraints. Badanidiyuru et al. [1] provided a  $(1/2 - \epsilon)$ -approximation algorithm that can be executed in  $O(k \log k)$  space for the cardinality constraint. For more general constraints including matching and multiple matroids constraints, Chakrabarti and Kale [5] proposed constant factor approximation algorithms. Chekuri et al. [6] devised algorithms for non-monotone submodular functions.

On the other hand, much effort is also devoted to submodular maximization in the secretary setting, called *submodular secretary problem*, under various constraints. Bateni et al. [3] specified the problem first and provided algorithms for both monotone and non-monotone submodular secretary problems under several constraints, one of which our methods are based on. Feldman et al. [14] improved constant factors of the theoretical guarantees for monotone cases.

## 8 Concluding Remarks

In this paper, we investigated stream-based active learning with a budget constraint in the view of adaptive submodular maximization. To tackle this problem, we introduced the adaptive stochastic maximization problem in the stream and secretary settings, which can formalize stream-based active learning. We provided a new class of objective functions, policy-adaptive submodular functions, and showed this class contains many utility functions that have been used in pool-based active learning and other applications. `AdaptiveStream` and `AdaptiveSecretary`, which we proposed in this paper, are simple algorithms guaranteed to be constant factor competitive with the optimal pool-based policy. We empirically demonstrated their performance by applying our algorithms to the budgeted stream-based active learning problem, and our experimental results indicate their effectiveness compared to the existing methods.

There are two natural directions for future work. One is exploring the possibility of the concept, policy-adaptive submodularity. By studying the nature of this class, we can probably yield theoretical insight for other problems. Another is further developing the practical aspects of our results. In real world problems sometimes it happens that the items arrive not in a random order. For example, in sequential adaptive sensor placement [11], an order of items is restricted to some transportation constraint. In this setting our guarantees do not hold and another algorithm is needed. In contrast to the non-adaptive setting, even in the stream setting, it seems much more difficult to design a constant factor approximation algorithm because the full information of each item is totally revealed when its state is observed and memory is not so powerful as in the non-adaptive setting.

## Acknowledgments

The second author is supported by Grant-in-Aid for Scientific Research on Innovative Areas, Exploration of nanostructure-property relationships for materials innovation.



## References

- [1] A. Badanidiyuru, B. Mirzasoleiman, A. Karbasi, and A. Krause. Streaming submodular maximization: Massive data summarization on the fly. *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pp. 671–680, 2014.
- [2] M.-F. Balcan, A. Beygelzimer, and J. Langford. Agnostic active learning. *Proceedings of the 23rd International Conference on Machine Learning (ICML)*, pp. 65–72, 2006.
- [3] M. Bateni, M. Hajiaghayi, and M. Zadimoghaddam. Submodular secretary problem and extensions. *ACM Transactions on Algorithms (TALG)*, 9(4):32, 2013.
- [4] A. Beygelzimer, S. Dasgupta, and J. Langford. Importance-weighted active learning. *Proceedings of the 26th International Conference on Machine Learning (ICML)*, pp. 49–56, 2009.
- [5] A. Chakrabarti and S. Kale. Submodular maximization meets streaming: Matchings, matroids, and more. *Mathematical Programming Series B*, 154(1), pp. 225–247, 2015.
- [6] C. Chekuri, S. Gupta, and K. Quanrud. Streaming algorithms for submodular function maximization. *Automata, Languages, and Programming (ICALP)*, pp. 318–330, 2015.
- [7] Y. Chen and A. Krause. Near-optimal batch mode active learning and adaptive submodular optimization. *Proceedings of the 30th International Conference on Machine Learning (ICML)*, pp. 160–168, 2013.
- [8] N. V. Cuong, W. S. Lee, and N. Ye. Near-optimal adaptive pool-based active learning with general loss. *Uncertainty in Artificial Intelligence (UAI)*, 2014.
- [9] N. V. Cuong, W. S. Lee, N. Ye, K. M. A. Chai, and H. L. Chieu. Active learning for probabilistic hypotheses using the maximum Gibbs error criterion. *Advances in Neural Information Processing Systems (NIPS)*, pp. 1457–1465, 2013.
- [10] I. Dagan and S. Engelson. Committee-based sampling for training probabilistic classifiers. *Proceedings of the 12th International Conference on Machine Learning (ICML)*, pp. 150–157, 1995.
- [11] J. Das, F. Py, J. B. J. Harvey, J. P. Ryan, A. Gellene, R. Graham, D. A. Caron, K. Rajan, and G. S. Sukhatme. Data-driven robotic sampling for marine ecosystem monitoring. *The International Journal of Robotics Research*, 34(12), pp. 1435–1452, 2015.
- [12] S. Dasgupta. Analysis of a greedy active learning strategy. *Advances in Neural Information Processing Systems (NIPS)*, pp. 337–344, 2004.
- [13] E. B. Dynkin. The optimum choice of the instant for stopping a Markov process. *Soviet Math. Dokl*, 4, pp. 627–629, 1963.
- [14] M. Feldman, J. S. Naor, and R. Schwartz. Improved competitive ratios for submodular secretary problems. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX-RANDOM)*, pp. 218–229, 2011.
- [15] U. Feige, V. Mirrokni, and J. Vondrák. Maximizing non-monotone submodular functions. *SIAM Journal on Computing*, 40(4), pp. 1133–1153, 2011.
- [16] Y. Freund, H. S. Seung, E. Shamir, and N. Tishby. Selective sampling using the query by committee algorithm. *Machine Learning*, 28, pp. 133–168, 1997.
- [17] S. Fujishige. *Submodular Functions and Optimization*, Second Edition. Annals of Discrete Mathematics, Vol. 58, Elsevier, 2005.
- [18] V. Gabillon, B. Kveton, Z. Wen, B. Eriksson, and S. Muthukrishnan. Adaptive submodular maximization in bandit setting. *Advances in Neural Information Processing Systems (NIPS)*, pp. 2697–2705, 2013.
- [19] D. Golovin and A. Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *Journal of Artificial Intelligence Research (JAIR)*, 42, pp. 427–486, 2011.
- [20] D. Golovin, A. Krause, and D. Ray. Near-optimal Bayesian active learning with noisy observations. *Advances in Neural Information Processing Systems (NIPS)*, pp. 766–774, 2010.
- [21] A. Gonen, S. Sabato, and S. Shalev-Shwartz. Efficient active learning of halfspaces: An aggressive approach. *The Journal of Machine Learning Research (JMLR)*, 14(1), pp. 2583–2615, 2013.
- [22] A. Gotovos, A. Karbasi, and A. Krause. Non-monotone adaptive submodular maximization. *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 1996–2003, 2015.
- [23] C. C. Loy, T. M. Hospedales, T. Xiang, and S. Gong. Stream-based joint exploration-exploitation active learning. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [24] S. Sabato and T. Hess. Interactive algorithms: From pool to stream. In *Proceedings of the 29th Annual Conference on Learning Theory (COLT)*, pp. 1419–1439, 2016.
- [25] D. Sculley. Online active learning methods for fast label-efficient spam filtering. *Proceedings of Fourth Conference on Email and Anti-Spam (CEAS)*, 2007.
- [26] J. Smailović, M. Grčar, N. Lavrač, and M. Žnidaršič. Stream-based active learning for sentiment analysis in the financial domain. *Information Sciences*, 285, pp. 181–203, 2014.

## A Proofs of the Policy-Adaptive Submodularity

In this section, we show several adaptive submodular functions used as utility functions in applications are policy-adaptive submodular.

Let  $\mathcal{H}$  denote the set of candidates for the true hypothesis. First we consider the noiseless case. Each hypothesis  $h \in \mathcal{H}$  represents some realization, i.e.,  $h : V \rightarrow \mathcal{Y}$ . Let  $p_H$  be a prior distribution over hypotheses. For convenience, we use  $p_H$  as the sum of the probability  $p_H(\mathcal{H}') = \sum_{h \in \mathcal{H}'} p_H(h)$  for any  $\mathcal{H}' \subseteq \mathcal{H}$ . We can obtain the prior distribution over realizations as  $p(\phi) = p_H(\{h \mid \forall s \in V, h(s) = \phi(s)\})$ . The version space under observations  $\psi$  is defined to be  $\mathcal{H}(\psi) = \{h \in \mathcal{H} \mid \forall (s, y) \in \psi, h(s) = y\}$ .

**The objective function of generalized binary search.** Generalized binary search in the Bayesian setting [12], designed for noiseless active learning, is executed by greedily maximizing

$$f_{\text{GBS}}(\psi) = 1 - p_H(\mathcal{H}(\psi)).$$

Golovin and Krause [19] proved this function is adaptive submodular, and improved the upper bound on the number of labels made by generalized binary search for deciding the true hypothesis.

**Proposition A.1.**  $f_{\text{GBS}}$  is policy-adaptive submodular w.r.t.  $p(\phi)$ .

*Proof.* Let  $\psi_0$  be any partial realization and  $\pi$  any policy such that  $\text{range}(\pi) \subseteq V \setminus \text{dom}(\psi_0)$ . The set of all resulting observations of  $\pi$  is defined to be  $\mathcal{L}(\pi) = \{\psi \mid \exists \phi, \psi(\pi, \phi)\}$ . The expected marginal gain of  $\pi$  under the observations  $\psi_0$  is

$$\begin{aligned} \Delta(\pi|\psi_0) &= \sum_{\psi \in \mathcal{L}(\pi)} p(\psi|\psi_0) \left( f_{\text{GBS}}(\psi \cup \psi_0) - f_{\text{GBS}}(\psi_0) \right) \\ &= \sum_{\psi \in \mathcal{L}(\pi)} \frac{p_H(\mathcal{H}(\psi \cup \psi_0))}{p_H(\mathcal{H}(\psi_0))} \left( p_H(\mathcal{H}(\psi_0)) - p_H(\mathcal{H}(\psi \cup \psi_0)) \right). \end{aligned}$$

Let  $g$  be a function of  $\mathbf{x} = (x_\psi)_\psi = (p_H(\mathcal{H}(\psi \cup \psi_0)))_\psi$  such that  $g(\mathbf{x}) = \Delta(\pi|\psi_0)$ , then  $g(\mathbf{x}) = \sum_\psi x_\psi - \sum_\psi x_\psi^2 / \sum_\psi x_\psi$ . We show this function is monotonically increasing about each  $x_\psi$ .

By differentiating  $g$  with respect to  $x_\psi$ ,

$$\frac{\partial}{\partial x_\psi} g(\mathbf{x}) = 1 + \frac{\sum_{\psi'} x_{\psi'}^2}{(\sum_{\psi'} x_{\psi'})^2} - \frac{2x_\psi}{\sum_{\psi'} x_{\psi'}} = \left( 1 - \frac{x_\psi}{\sum_{\psi'} x_{\psi'}} \right)^2 + \frac{\sum_{\psi' \neq \psi} x_{\psi'}^2}{(\sum_{\psi'} x_{\psi'})^2}$$

which is at least 0. Therefore  $g$  is monotonically increasing w.r.t. each  $x_\psi$ .

If  $\psi_A$  and  $\psi_B$  are partial realizations such that  $\psi_A \subseteq \psi_B$  and  $\text{range}(\pi) \subseteq V \setminus B$ , it holds that  $p_H(\mathcal{H}(\psi \cup \psi_A)) \geq p_H(\mathcal{H}(\psi \cup \psi_B))$  for all  $\psi \in \mathcal{L}(\pi)$ , which implies  $\Delta(\pi|\psi_A) \geq \Delta(\pi|\psi_B)$ .  $\square$

**The objective function of EC<sup>2</sup> algorithm.** Suppose that  $\mathcal{H}$  is partitioned into equivalence classes  $\mathcal{H}_1, \dots, \mathcal{H}_m$ , i.e.,  $\mathcal{H}_1 \cup \dots \cup \mathcal{H}_m = \mathcal{H}$  and  $\mathcal{H}_i \cap \mathcal{H}_j = \emptyset$  for every  $i \neq j$ . The equivalence class determination problem [20] is the problem of deciding which equivalence class the true hypothesis lies in. Golovin et al. [20] proves that

$$f_{\text{EC}}(\psi) = 1 - \sum_{i < j} \left( p_H(\mathcal{H}_i(\psi)) p_H(\mathcal{H}_j(\psi)) \right)$$

is adaptive submodular and adaptive monotone, and this function achieves the value 1 if and only if the equivalence class that the true hypothesis lies in is determined. Now we show  $(f_{\text{EC}}, p)$  is policy-adaptive submodular.

**Proposition A.2.**  $f_{\text{EC}}$  is policy-adaptive submodular w.r.t.  $p(\phi)$ .

*Proof.* Let  $\psi_0$  be any partial realization and  $\pi$  any policy such that  $\text{range}(\pi) \subseteq V \setminus \text{dom}(\psi_0)$ . The set of all resulting observations of  $\pi$  is defined to be  $\mathcal{L}(\pi) = \{\psi \mid \exists \phi, \psi(\pi, \phi)\}$ . The expected

marginal gain of  $\pi$  under the partial realization  $\psi_0$  is

$$\begin{aligned}\Delta(\pi|\psi_0) &= \sum_{\psi \in \mathcal{L}(\pi)} p(\psi|\psi_0) \left( f_{\text{EC}}(\psi \cup \psi_0) - f_{\text{EC}}(\psi_0) \right) \\ &= \sum_{\psi \in \mathcal{L}(\pi)} \frac{p_H(\mathcal{H}(\psi \cup \psi_0))}{p_H(\mathcal{H}(\psi_0))} \left( \sum_{i < j} \left( p_H(\mathcal{H}_i(\psi_0)) p_H(\mathcal{H}_j(\psi_0)) \right) \right. \\ &\quad \left. - \sum_{i < j} \left( p_H(\mathcal{H}_i(\psi \cup \psi_0)) p_H(\mathcal{H}_j(\psi \cup \psi_0)) \right) \right).\end{aligned}$$

Let  $g$  be a function of  $\mathbf{x} = (x_{i,\psi})_{i,\psi} = (p_H(\mathcal{H}_i(\psi \cup \psi_0)))_{i,\psi}$  such that  $g(\mathbf{x}) = \Delta(\pi|\psi_0)$ , then

$$g(\mathbf{x}) = \sum_{\psi} \left[ \frac{\sum_i x_{i,\psi}}{\sum_{i,\psi'} x_{i,\psi'}} \sum_{i < j} \left( \left( \sum_{\psi'} x_{i,\psi'} \right) \left( \sum_{\psi'} x_{j,\psi'} \right) - x_{i,\psi} x_{j,\psi} \right) \right].$$

In the same way as the proof of the adaptive submodularity of  $f_{\text{EC}}$  given by Golovin et al. [20], we can see this function is monotonically increasing about each  $x_{i,\psi}$  by showing the non-negativity of the derivative of  $g$ .

If  $\psi_A$  and  $\psi_B$  are partial realizations such that  $\psi_A \subseteq \psi_B$  and  $\text{range}(\pi) \subseteq V \setminus B$ , it holds that  $p_H(\mathcal{H}_i(\psi \cup \psi_A)) \geq p_H(\mathcal{H}_i(\psi \cup \psi_B))$  for all  $i = 1, \dots, m$  and  $\psi \in \mathcal{L}(\pi)$ , which implies  $\Delta(\pi|\psi_A) \geq \Delta(\pi|\psi_B)$ .  $\square$

**The objective function of ALuMA algorithm.** We consider the binary classification problem, i.e.,  $\mathcal{Y} = \{+1, -1\}$ , for linearly separable instances. Suppose the parameter  $\mathbf{w} \in \mathbb{R}^d$  is uniformly distributed in the  $d$ -dimensional unit ball  $\mathbb{B}_1^d$ , and the classifier of parameter  $\mathbf{w}$  classifies an instance  $\mathbf{x} \in \mathbb{R}^d$  into the class  $\text{sgn}(\langle \mathbf{w}, \mathbf{x} \rangle)$ . Gonen et al. [21] proposed the objective function

$$f_{\text{ALuMA}}(\psi) = 1 - \mathbb{P}[\{\mathbf{w} \in \mathbb{B}_1^d \mid \forall (\mathbf{x}, y) \in \psi, y = \text{sgn}(\langle \mathbf{w}, \mathbf{x} \rangle)\}],$$

and prove it is adaptive submodular. The prior probability over realizations is  $p(\phi) = \mathbb{P}[\{\mathbf{w} \in \mathbb{B}_1^d \mid \forall \mathbf{x} \in V, \Phi(\mathbf{x}) = \text{sgn}(\langle \mathbf{w}, \mathbf{x} \rangle)\}]$ .

**Proposition A.3.** *The objective function  $f_{\text{ALuMA}}$  is policy-adaptive submodular w.r.t.  $p(\phi)$ .*

*Proof.* It is a special case of  $f_{\text{GBS}}$  where the prior probability  $p(\phi)$  can be expressed as above. Hence  $f_{\text{ALuMA}}$  is policy-adaptive submodular.  $\square$

Though running exactly the greedy algorithm for  $f_{\text{ALuMA}}$  is difficult, there is a method to approximately execute the greedy policy within some constant factor [21]. We can apply this approximate method to the stream-based setting and give the performance guarantees by using the policy-adaptive submodularity.

**The objective function of maximum Gibbs error criterion.** Now suppose labels are generated from noisy distribution, i.e.,  $\mathbb{P}[Y_1, \dots, Y_n | H = h]$  is not deterministic. In this case, the prior over realizations is expressed as  $p(\phi) = \sum_{h \in \mathcal{H}} p_H(h) \mathbb{P}[\phi | H = h]$ . By using this distribution, the utility function of the maximum Gibbs error criterion [9, 8] is defined as

$$f_{\text{Gibbs}}(\psi) = 1 - p(\psi).$$

**Proposition A.4.**  *$f_{\text{Gibbs}}$  is policy-adaptive submodular w.r.t.  $p(\phi)$ .*

*Proof.* The same proof as the one of Proposition A.1 can be applied by replacing the probability mass of hypotheses  $p_H(\mathcal{H}(\psi))$  with the probability mass of partial realizations  $p(\psi)$ .  $\square$

**In cases where the states of the items are independent.** In many applications other than active learning, it is often assumed that the states of the items are independent of each other. For instance, this assumption is often used in adaptive variations of sensor placements and influence maximization, both of which are introduced by Golovin and Krause [19]. In this case, the policy-adaptive submodularity follows from the adaptive submodularity.

**Proposition A.5.** *If  $Y_1, \dots, Y_n$  are independent and  $(f, p)$  is adaptive submodular, then  $(f, p)$  is policy-adaptive submodular.*

*Proof.* Let  $\psi_A$  and  $\psi_B$  be any partial realization such that  $\psi_A \subseteq \psi_B$ , and let  $\pi$  denote any policy such that  $\text{range}(\pi) \subseteq V \setminus B$ . We prove  $\Delta(\pi|\psi_A) \geq \Delta(\pi|\psi_B)$  with the induction on the height of policy tree  $\pi$ .

In the case where the height of  $\pi$  is 1, i.e.,  $\pi$  is a policy selecting only a single item, the statement follows from the adaptive submodularity of  $(f, p)$ .

Assume that the statement holds for all policies whose height is at most  $l - 1$ . Suppose the height of  $\pi$  is  $l$ . Let  $s$  denote the first item that  $\pi$  selects and  $\pi_y$  be the subpolicy of  $\pi$  that is executed just after  $Y_s$  turns out to be  $y$  for each  $y \in \mathcal{Y}$ . Since  $Y_1, \dots, Y_n$  are independent and  $s \in V \setminus B$ , it holds that  $p(Y_s = y|\psi_A) = p(Y_s = y|\psi_B)$ . Applying the induction hypothesis to  $\pi_y$ , whose height is at most  $l - 1$ , we obtain  $\Delta(\pi_y|\psi_A \cup \{(s, y)\}) \geq \Delta(\pi_y|\psi_B \cup \{(s, y)\})$ . Therefore, we can see that

$$\begin{aligned} \Delta(\pi|\psi_A) &= \Delta(s|\psi_A) + \sum_{y \in \mathcal{Y}} p(Y_s = y|\psi_A) \Delta(\pi_y|\psi_A \cup \{(s, y)\}) \\ &\geq \Delta(s|\psi_B) + \sum_{y \in \mathcal{Y}} p(Y_s = y|\psi_B) \Delta(\pi_y|\psi_B \cup \{(s, y)\}) \\ &= \Delta(\pi|\psi_B) \end{aligned}$$

□

**A counterexample to the equivalence.** As described above, many adaptive submodular objective functions proposed in previous work are also policy-adaptive submodular. However, there is a simple counterexample to the equivalence of adaptive submodularity and policy-adaptive submodularity.

**Example A.6.** Let  $V = \{a, b, c\}$  and  $\mathcal{Y} = \{+1, -1\}$ . If  $w : V \rightarrow \mathbb{R}_{\geq 0}$  is a weight vector such that  $w(a) = w(b) = w(c) = 1$ , the sum of weights  $f(\psi) = \sum_{s \in \text{dom}(\psi)} w(s)$  is adaptive submodular w.r.t. any prior distribution  $p(\phi)$ .

Let  $\pi$  denote a policy that selects  $b$  first, and proceeds to select  $c$  if  $Y_b = +1$ . Assume the marginal distribution of  $\{Y_a, Y_b\}$  is defined as  $\mathbb{P}[Y_a = +1, Y_b = +1] = 0.09$ ,  $\mathbb{P}[Y_a = +1, Y_b = -1] = \mathbb{P}[Y_a = -1, Y_b = +1] = 0.01$ , and  $\mathbb{P}[Y_a = -1, Y_b = -1] = 0.89$ . Since  $\Delta(\pi|\{(a, +1)\}) = 1.9$  is larger than  $\Delta(\pi|\emptyset) = 1.1$ , we know that  $(f, p)$  is not policy-adaptive submodular.

$(f, p)$  is not only adaptive submodular, but linear and strongly adaptive monotone<sup>5</sup>, hence we can see that policy-adaptive submodularity does not follow from any combination of these conditions.

## B Pseudocodes of the Proposed Algorithms

---

### Algorithm 2 AdaptiveStream algorithm

---

**Input:** A set function  $f : 2^{V \times \mathcal{Y}} \rightarrow \mathbb{R}_{\geq 0}$  and a prior distribution  $p(\phi)$  such that  $(f, p)$  is policy-adaptive submodular and adaptive monotone. The number of items in the entire stream  $n \in \mathbb{Z}_{>0}$ . A budget  $k \in \mathbb{Z}_{>0}$ . Randomly permuted stream of the items, denoted by  $(s_1, \dots, s_n)$ .

**Output:** Some observations  $\psi_k \subseteq V \times \mathcal{Y}$  such that  $|\psi_k| \leq k$ .

- 1: Let  $\psi_0 := \emptyset$ .
  - 2: **for** each segment  $S_l = \{s_i \mid (l-1)n/k < i \leq ln/k\}$  **do**
  - 3:    $\delta_{\max} := 0$
  - 4:   **for** each item  $s \in S_l$  **do**
  - 5:     **if**  $\Delta(s|\psi_{l-1}) \geq \delta_{\max}$  **then**
  - 6:        $s_{\max} := s, \delta_{\max} := \Delta(s|\psi_{l-1})$ .
  - 7:   Observe the state  $y$  of item  $s_{\max}$  and let  $\psi_l := \psi_{l-1} \cup \{(s_{\max}, y)\}$ .
  - 8: **return**  $\psi_k$  as the solution
- 

<sup>5</sup>The strong adaptive monotonicity of  $(f, p)$  is defined to be  $\psi \subseteq \psi' \Rightarrow f(\psi) \leq f(\psi')$ .

---

**Algorithm 3** AdaptiveSecretary algorithm

---

**Input:** A set function  $f : 2^{V \times \mathcal{Y}} \rightarrow \mathbb{R}_{\geq 0}$  and a prior distribution  $p(\phi)$  such that  $(f, p)$  is policy-adaptive submodular and adaptive monotone. The number of items in the entire stream  $n \in \mathbb{Z}_{>0}$ . A budget  $k \in \mathbb{Z}_{>0}$ . Randomly permuted stream of the items, denoted by  $(s_1, \dots, s_n)$ .

**Output:** Some observations  $\psi_k \subseteq V \times \mathcal{Y}$  such that  $|\psi_k| \leq k$ .

```
1: Let  $\psi_0 := \emptyset$ .
2: for each segment  $S_l = \{s_i \mid (l-1)n/k < i \leq ln/k\}$  do
3:   Let the first  $\lfloor n/(ek) \rfloor$  items in  $S_l$  pass,
   and let  $\delta_{\max} := \{\Delta(s_i|\psi_{l-1}) \mid (l-1)n/k < i \leq (l-1)n/k + \lfloor n/(ek) \rfloor\}$ .
4:   select := false
5:   for each  $s_i$  such that  $(l-1)n/k + \lfloor n/(ek) \rfloor < i \leq ln/k$  do
6:     if  $\Delta(s_i|\psi_{l-1}) \geq \delta_{\max}$  then
7:       Observe the state  $y$  of item  $s_i$  and let  $\psi_l := \psi_{l-1} \cup \{(s_i, y)\}$ .
8:       select := true
9:     break
10:  if select = false then
11:     $\psi_l := \psi_{l-1}$ .
12: return  $\psi_k$  as the solution
```

---

## C Proofs of the Theoretical Guarantees of the Algorithms

Define the conditional function as  $f(\psi'|\psi) = f(\psi' \cup \psi) - f(\psi)$  where  $\psi$  and  $\psi'$  are partial realizations. We denote the expected marginal gain about function  $f$  by  $\Delta_f$ . We can show the policy-adaptive submodularity is preserved by restriction as follows.

**Lemma C.1.** Suppose  $f : 2^{V \times \mathcal{Y}} \rightarrow \mathbb{R}$  is policy-adaptive submodular w.r.t. a prior  $p(\phi)$ . For an arbitrary partial realization  $\psi_S$ ,  $f(\cdot|\psi_S)$  is also policy-adaptive submodular w.r.t. the prior distribution  $p(\phi|\psi_S)$ .

*Proof.* Let  $\psi_A$  and  $\psi_B$  be partial realizations such that  $\psi_A \subseteq \psi_B$  and  $B \subseteq V \setminus S$ . Let  $\pi$  be any policy such that  $\text{range}(\pi) \subseteq V \setminus (S \cup B)$ . From the definition,  $\Delta_{f(\cdot|\psi_S)}(\cdot|\psi_A)$  is equal to  $\Delta_f(\cdot|\psi_A \cup \psi_S)$ . This leads to

$$\Delta_{f(\cdot|\psi_S)}(\pi|\psi_A) = \Delta_f(\pi|\psi_A \cup \psi_S) \geq \Delta_f(\pi|\psi_B \cup \psi_S) = \Delta_{f(\cdot|\psi_S)}(\pi|\psi_B).$$

The inequality is due to the policy-adaptive submodularity of  $(f, p)$ . It follows that  $f(\cdot|\psi_S)$  is policy-adaptive submodular w.r.t.  $p(\phi|\psi_S)$ .  $\square$

Define  $E(\pi, \phi)$  to be the set of items obtained when  $\pi$  is executed under realization  $\phi$ , that is,  $E(\pi, \phi) = \text{dom}(\psi(\pi, \phi))$ . To prove the next lemma, we use the following lemma proved by Golovin and Krause [19].

**Lemma C.2** (See [19, Lemme 5.3]). Suppose we have made observations  $\psi$  after selecting  $\text{dom}(\psi)$ . Let  $\pi^*$  be any policy such that  $|E(\pi^*, \phi)| \leq k$  for all  $\phi$ . Then for adaptive monotone submodular  $f$

$$\Delta(\pi^*|\psi) \leq \max_{A \subseteq E, |A| \leq k} \sum_{e \in A} \Delta(e|\psi).$$

Define  $f_{\text{avg}}(\pi) = \Delta(\pi|\emptyset)$  to be the average gain of a policy  $\pi$ . Here we prove the lemma about the maximal expected gain of a random selected subset without the policy-adaptive submodularity.

**Lemma C.3.** Let  $f : 2^{V \times \mathcal{Y}} \rightarrow \mathbb{R}_{\geq 0}$  be a set function and  $p$  any prior over realizations such that  $(f, p)$  is adaptive submodular and adaptive monotone. Let  $\pi$  be any policy such that  $|\psi(\pi, \phi)| \leq k$  for all  $\phi$ . Let  $T \subseteq V$  be a random set that contains each item of  $V$  independently with probability  $r \in [0, 1]$ . Then it holds that

$$\mathbb{E}_T \left[ \max_{s \in T} \Delta(s|\emptyset) \right] \geq \frac{1 - e^{-rk}}{k} f_{\text{avg}}(\pi).$$

*Proof.* From the adaptive data dependent bound (Lemma C.2), we have

$$f_{\text{avg}}(\pi) = \Delta(\pi|\emptyset) \leq \max_{A \subseteq V, |A| \leq k} \sum_{s \in A} \Delta(s|\emptyset).$$

Let  $\{s_1, \dots, s_k\} = \arg\max_{A \subseteq V, |A| \leq k} \sum_{s \in A} \Delta(s|\emptyset)$ . The probability that none of  $s_1, \dots, s_k$  is in  $T$  is upper bounded as

$$\mathbb{P}_T[\forall j = 1, \dots, k, s_j \notin T] = (1 - r)^k \leq e^{-rk}.$$

Under the condition that there is at least one  $s_j$  in  $T$ , the expected value of  $\max\{\Delta(s_j|\emptyset) \mid s_j \in T\}$  is no less than the average  $\sum_{j=1}^k \Delta(s_j|\emptyset)/k$  because of the symmetry of  $s_j$ s. Since we have at least one  $s_j$  with probability no less than  $(1 - e^{-rk})$ , it holds that

$$\begin{aligned} & \mathbb{E}_T[\max\{\Delta(s|\emptyset) \mid s \in T\}] \\ &= \mathbb{P}_T[\exists j = 1, \dots, k, s_j \in T] \mathbb{E}_T[\max\{\Delta(s|\emptyset) \mid s \in T\} \mid \exists j = 1, \dots, k, s_j \in T] \\ &\geq \frac{1 - e^{-rk}}{k} \sum_{j=1}^k \Delta(s_j|\emptyset) \\ &\geq \frac{1 - e^{-rk}}{k} f_{\text{avg}}(\pi). \end{aligned}$$

□

Feige, Mirrokni, and Vondrák [15] gave the useful lemma about the value of a randomly selected subset for non-adaptive submodular functions as follows:

**Lemma C.4** (See [15, Lemma 2.2]). *Let  $g : 2^X \rightarrow \mathbb{R}$  be submodular. Denote by  $A(p)$  a random subset of  $A$  where each element appears with probability  $p$ . Then*

$$\mathbb{E}[g(A(p))] \geq (1 - p)g(\emptyset) + pg(A)$$

This lemma is a crucial part of many randomized algorithms for submodular maximization, but cannot be extended straightforward to adaptive submodular functions. However, assuming the policy-adaptive submodularity, we can prove a similar property for the adaptive setting. Before stating the lemma, we define a partially deleted policy tree.

**Definition C.5** (Grafted policy). *Let  $\pi$  be an arbitrary policy, and let  $T$  be a subset of  $V$ . Let  $s$  denote the item that  $\pi$  selects first, and  $\pi_y$  the subpolicy executed by  $\pi$  when  $Y_s$  turns out to be  $y$  for each  $y \in \mathcal{Y}$ . Formally, the grafted policy  $\pi_T$  of  $\pi$  for  $T$  under partial realization  $\psi$  is a randomized policy<sup>6</sup> defined recursively as follows: If  $s \in T$ ,  $\pi_T$  selects  $s$  first, observes its label  $y$  and proceeds to the grafted policy of  $\pi_y$  for  $T$  under partial realization  $\psi \cup \{(s, y)\}$ . If  $s \notin T$ ,  $\pi_T$  proceeds randomly with probability  $\mathbb{P}[Y_s = y|\psi]$  to the grafted policy of  $\pi_y$  for  $T$  under partial realization  $\psi \cup \{(s, y)\}$  for each  $y \in \mathcal{Y}$ . For simplicity, we omit to mention the partial realization if it is the grafted policy under no observation.*

Intuitively, the grafted policy is obtained by repeating the grafting operation, which is indicated in Figure 3, from the leaves to the root for each removed item.

Now we prove the adaptive version of Lemma C.4.

**Lemma C.6.** *Let  $f : 2^{V \times \mathcal{Y}} \rightarrow \mathbb{R}$  be a set function and  $p$  any prior over realizations such that  $(f, p)$  is policy-adaptive submodular. Let  $\pi$  denote an arbitrary policy. Let  $T$  be a random subset of  $V$  where each element appears with probability  $r \in [0, 1]$ . If  $\pi_T$  is the grafted policy for  $T$ , it holds that*

$$\mathbb{E}_T[f_{\text{avg}}(\pi_T)] \geq (1 - r)f(\emptyset) + rf_{\text{avg}}(\pi).$$

<sup>6</sup>Here we must extend the definition of the policy to randomized policies. A randomized policy is a partial mapping from the current partial realization to a distribution over items that represents the item to be selected next. It is easy to see that the optimal deterministic policy is still optimal among randomized policies under any budget since a randomized policy can be expressed as the weighted sum of deterministic policies.




Figure 3: If item  $s$  is not contained in restricted set  $T$ , remove  $s$  from the original policy  $\pi$  and proceed to the subpolicy, which is executed by  $\pi$  when  $Y_s$  turns out to be  $y$ , with probability  $\mathbb{P}[Y_s = y|\psi]$ . Repeating this operation recursively from the leaves to the root, we obtain the grafted policy of  $\pi$ .

*Proof.* We apply the induction on the height  $k$  of the policy tree  $\pi$  to prove the statement.

In the case of  $k = 1$ , the item  $s \in T$  that  $\pi$  selects is also selected by  $\pi_T$  if  $s \in T$ . If  $s \notin T$ , no item is selected by  $\pi_T$ . Since the probability that  $T$  contains  $s$  is  $r$ ,  $\mathbb{E}_T[f_{\text{avg}}(\pi_T)] = (1-r)f(\emptyset) + rf_{\text{avg}}(\pi)$ .

Assume the statement holds in the case of  $k \leq l-1$ , we show it holds in the case of  $k = l$ . Let  $\pi$  be any policy whose height is  $l$ . Let  $s$  be the item  $\pi$  selects first and  $\pi_y$  the subpolicy executed by  $\pi$  just after the label of  $s$  turns out to be  $y \in \mathcal{Y}$ . Denoted by  $\pi_{T,y}$  the grafted policy of  $\pi_y$  for  $T$  under the observation  $\{(s, y)\}$ .

Now we can decompose the expected total gain of the policy  $\pi$  into the root and subtrees, i.e.,

$$f_{\text{avg}}(\pi) = f(\emptyset) + \Delta(s|\emptyset) + \sum_{y \in \mathcal{Y}} p(\psi_y) \Delta(\pi_y|\psi_y) \quad (1)$$

where  $\psi_y = \{(s, y)\}$  for each  $y \in \mathcal{Y}$ .

From Lemma C.1,  $(f(\cdot|\psi_y), p(\cdot|\psi_y))$  is policy-adaptive submodular. The grafted policy  $\pi_{T,y}$  under the observation  $\psi_y$  is the same as the grafted policy of  $\pi_y$  for  $T$  under no observation with respect to  $p(\cdot|\psi_y)$ . Since  $f(\emptyset|\psi_y) = 0$ , applying the induction hypothesis and we obtain that for all  $y \in \mathcal{Y}$ ,

$$\mathbb{E}_T[\Delta(\pi_{T,y}|\psi_y)] \geq r \Delta(\pi_y|\psi_y). \quad (2)$$

We consider separately the cases where  $s \in T$  and  $s \notin T$ . In the case of  $s \in T$ , it holds that

$$f_{\text{avg}}(\pi_T) = f(\emptyset) + \Delta(s|\emptyset) + \sum_{y \in \mathcal{Y}} p(\psi_y) \Delta(\pi_{T,y}|\psi_y). \quad (3)$$

In the case of  $s \notin T$ , we have

$$\begin{aligned} f_{\text{avg}}(\pi_T) &= f(\emptyset) + \sum_{y \in \mathcal{Y}} p(\psi_y) f_{\text{avg}}(\pi_{T,y}) \\ &\geq f(\emptyset) + \sum_{y \in \mathcal{Y}} p(\psi_y) \Delta(\pi_{T,y}|\psi_y). \end{aligned} \quad (4)$$

The second inequality is due to the policy-adaptive submodularity of  $(f, p)$ .

Hence we obtain

$$\begin{aligned}
\mathbb{E}_T[f_{\text{avg}}(\pi_T)] &= \mathbb{P}[s \in T] \mathbb{E}_T[f_{\text{avg}}(\pi_T) \mid s \in T] + \mathbb{P}[s \notin T] \mathbb{E}_T[f_{\text{avg}}(\pi_T) \mid s \notin T] \\
&\geq r \left( f(\emptyset) + \Delta(s|\emptyset) + \sum_{y \in \mathcal{Y}} \left( p(\psi_y) \mathbb{E}_T[\Delta(\pi_{T,y}|\psi_y) \mid s \in T] \right) \right) \\
&\quad + (1-r) \left( f(\emptyset) + \sum_{y \in \mathcal{Y}} \left( p(\psi_y) \mathbb{E}_T[\Delta(\pi_{T,y}|\psi_y) \mid s \notin T] \right) \right) \quad (\text{Since (3) and (4)}) \\
&= f(\emptyset) + r\Delta(s|\emptyset) + \sum_{y \in \mathcal{Y}} p(\psi_y) \mathbb{E}_T[\Delta(\pi_{T,y}|\psi_y)] \\
&\geq f(\emptyset) + r\Delta(s|\emptyset) + \sum_{y \in \mathcal{Y}} p(\psi_y) r\Delta(\pi_y|\psi_y) \quad (\text{Since (2)}) \\
&= (1-r)f(\emptyset) + r \left( f(\emptyset) + \Delta(s|\emptyset) + \sum_{y \in \mathcal{Y}} p(\psi_y) \Delta(\pi_y|\psi_y) \right) \\
&= (1-r)f(\emptyset) + rf_{\text{avg}}(\pi). \quad (\text{Since (1)})
\end{aligned}$$

□

We illustrate adaptive submodularity is not enough to prove this lemma by providing an example that satisfies adaptive submodularity but not the inequality in Lemma C.6.

**Example C.7.** Let  $f$  and  $p$  be the same as Example A.6. Define a policy  $\pi$  as follows. First it selects  $a$  and if the state of  $a$  is  $-1$ , it ends. If the state of  $a$  is  $+1$ , it next selects  $b$  and proceeds to select  $c$  if the state of  $b$  is  $+1$ . Consider the case of  $r = 0.5$ .

Let  $T$  be a random subset of  $\{a, b, c\}$ . The case of  $T = \{b, c\}$  is most important among eight cases. In this case the grafted policy  $\pi_T$  selects  $b$  with probability 0.1, and if the state of  $b$  is  $+1$ , it next selects  $c$ . The original policy  $\pi$  selects  $c$  with probability 0.09 but the grafted policy  $\pi_T$  selects  $c$  with probability 0.01. This difference results in  $\mathbb{E}_T[f_{\text{avg}}(\pi_T)] = 0.585$  and  $(1-r)f(\emptyset) + rf_{\text{avg}}(\pi) = 0.595$ , which violates the inequality in Lemma C.6.

Lemma C.6 implies the following property about the optimal pool-based policy constructed in a randomly selected subset.

**Lemma C.8.** Let  $f : 2^V \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$  be a set function and  $p$  any prior over realizations such that  $(f, p)$  is policy-adaptive submodular. Let  $T$  be a random subset of  $V$  where each element appears with probability  $r \in [0, 1]$ , and  $\pi_V^*$  and  $\pi_T^*$  the optimal pool-based policy selecting  $k$  items from  $V$  and  $T$  respectively. Then it holds

$$\mathbb{E}_T[f_{\text{avg}}(\pi_T^*)] \geq rf_{\text{avg}}(\pi_V^*).$$

*Proof.* Let  $\pi_T$  be the grafted policy of  $\pi_V^*$  for  $T$ . From Lemma C.6, we obtain

$$\mathbb{E}_T[f_{\text{avg}}(\pi_T)] \geq (1-r)f(\emptyset) + rf_{\text{avg}}(\pi_V^*) \geq rf_{\text{avg}}(\pi_V^*)$$

The second inequality holds because of the non-negativity of  $f$ . Due to the optimality of  $\pi_T^*$ , it holds that  $f_{\text{avg}}(\pi_T^*) \geq f_{\text{avg}}(\pi_T)$  for any  $T \subseteq V$ , which leads to the statement. □

Suppose  $\Phi \sim \psi$  means  $\Phi$  follows the distribution  $p(\Phi|\psi)$ . Let  $\pi' @ \pi$  denote the concatenated policy that run  $\pi$  to completion and run  $\pi'$  ignoring the information so far. To prove the main theorem we refer to the following lemma.

**Lemma C.9** (See [19, Lemme A.7]). Fix a function  $f : 2^E \times \mathcal{O}^E \rightarrow \mathbb{R}_{\geq 0}$ . Then  $\Delta(e|\psi) \geq 0$  for all  $\psi$  with  $\mathbb{P}[\Phi \sim \psi] > 0$  and all  $e \in E$  if and only if for all policies  $\pi$  and  $\pi'$ ,  $f_{\text{avg}}(\pi) \leq f_{\text{avg}}(\pi' @ \pi)$ .

We prove the theorem about the theoretical guarantee for AdaptiveStream algorithm.



*Proof of Theorem 4.1.* Let  $\sigma$  denote a random permutation in which the items arrive. Let  $\pi_i^\sigma$  be the stream-based policy executed by running `AdaptiveStream` until it selects an item from the  $i$ th segment  $S_i$  under the permutation  $\sigma$  and then terminating. Let  $\Psi_i^\sigma$  be a random partial realization obtained by  $\pi_i^\sigma$ .

First we bound from below the expected  $i$ th marginal gain  $\mathbb{E}_\sigma[f_{\text{avg}}(\pi_i^\sigma) - f_{\text{avg}}(\pi_{i-1}^\sigma)]$  of the algorithm. Since at the  $i$ th step, the algorithm selects the item of the largest expected marginal gain in  $S_i$ ,

$$\mathbb{E}_\sigma[f_{\text{avg}}(\pi_i^\sigma) - f_{\text{avg}}(\pi_{i-1}^\sigma)] = \mathbb{E}_{\sigma, \Phi} \left[ \max_{s \in S_i} \Delta(s | \Psi_{i-1}^\sigma) \right]. \quad (5)$$

Fix  $S_1 \cup S_2 \cup \dots \cup S_{i-1}$  (formally, fix the part of a random permutation  $\sigma$ ), we can regard  $S_i$  as a random subset that contains each item  $s \in V \setminus (S_1 \cup S_2 \cup \dots \cup S_{i-1})$  independently with probability  $1/(k-i+1)$ . Let  $T = V \setminus (S_1 \cup S_2 \cup \dots \cup S_{i-1})$ . Since  $f(\cdot | \Psi_{i-1}^\sigma)$  is adaptive submodular w.r.t.  $p(\phi | \Psi_{i-1}^\sigma)$ , from Lemma C.3 we obtain

$$\mathbb{E}_{\sigma, \Phi} \left[ \max_{s \in S_i} \Delta(s | \Psi_{i-1}^\sigma) \right] \geq \frac{1 - \exp(-\frac{k}{k-i+1})}{k} \mathbb{E}_{\sigma, \Phi} [\Delta(\pi_T^* | \Psi_{i-1}^\sigma)] \quad (6)$$

where  $\pi_T^*$  is the optimal pool-based policy selecting  $k$  items from  $T$ .

Under any fixed  $\sigma$ , we can view  $\mathbb{E}_\Phi[\Delta(\pi_T^* | \Psi_{i-1}^\sigma)]$  as the expected gain of  $\pi_T^*$  after  $\pi_i^\sigma$  is executed, then it holds that

$$\begin{aligned} \mathbb{E}_\Phi [\Delta(\pi_T^* | \Psi_{i-1}^\sigma)] &= f_{\text{avg}}(\pi_T^* @ \pi_{i-1}^\sigma) - f_{\text{avg}}(\pi_{i-1}^\sigma) \\ &\geq f_{\text{avg}}(\pi_T^*) - f_{\text{avg}}(\pi_{i-1}^\sigma). \end{aligned} \quad (7)$$

The inequality follows from the adaptive monotonicity of  $(f, p)$  and Lemma C.9.

Let  $\pi_V^*$  be the optimal pool-based policy selecting  $k$  items from  $V$ . Since  $T$  can be regarded as a random subset that contains each item of  $V$  independently with probability  $(k-i+1)/k$ , from Lemma C.8 we have

$$\mathbb{E}_\sigma[f_{\text{avg}}(\pi_T^*)] \geq \frac{k-i+1}{k} f_{\text{avg}}(\pi_V^*). \quad (8)$$

From (5), (6), (7), and (8), we obtain

$$\begin{aligned} \mathbb{E}_\sigma[f_{\text{avg}}(\pi_i^\sigma) - f_{\text{avg}}(\pi_{i-1}^\sigma)] &\geq \frac{1 - \exp(-\frac{k}{k-i+1})}{k} \left\{ \frac{k-i+1}{k} f_{\text{avg}}(\pi_V^*) - \mathbb{E}_\sigma[f_{\text{avg}}(\pi_{i-1}^\sigma)] \right\} \\ &\geq \frac{(1 - e^{-1})(k-i+1)}{k^2} f_{\text{avg}}(\pi_V^*) - \frac{1}{k} \mathbb{E}_\sigma[f_{\text{avg}}(\pi_{i-1}^\sigma)]. \end{aligned}$$

The upper and lower bounds  $0 \leq \exp(-\frac{k}{k-i+1}) \leq e^{-1}$  are used in the second inequality.

We use the proof by contradiction to give the performance guarantee of `AdaptiveStream`. Assume  $\mathbb{E}_\sigma[f_{\text{avg}}(\pi_k^\sigma)] < \alpha f_{\text{avg}}(\pi_V^*)$  (the concrete value of  $\alpha$  is specified later). Due to the adaptive monotonicity of  $(f, p)$ , we know that  $\mathbb{E}_\sigma[f_{\text{avg}}(\pi_i^\sigma)]$  is increasing monotonically as  $i$  increases. Hence for all  $i = 1, \dots, k$ , it holds that  $\mathbb{E}_{\sigma, \Phi}[f_{\text{avg}}(\pi_i^\sigma)] < \alpha f_{\text{avg}}(\pi_V^*)$ .

Let  $l = \lfloor (\sqrt{3} - 1)k \rfloor$ . Taking the summation of the above inequality until the  $l$ th step, we obtain the following:

$$\begin{aligned} \mathbb{E}_\sigma[f_{\text{avg}}(\pi_l^\sigma)] &\geq \frac{1 - e^{-1}}{k^2} \sum_{i=1}^l (k-i+1) f_{\text{avg}}(\pi_V^*) - \frac{1}{k} \sum_{i=1}^l \mathbb{E}_\sigma[f_{\text{avg}}(\pi_{i-1}^\sigma)] \\ &> \frac{1 - e^{-1}}{k^2} \sum_{i=1}^l (k-i+1) f_{\text{avg}}(\pi_V^*) - \frac{1}{k} \sum_{i=1}^l \alpha f_{\text{avg}}(\pi_V^*). \end{aligned}$$

Similarly, taking the summation until the  $(l+1)$ th step, we obtain

$$\mathbb{E}_\sigma[f_{\text{avg}}(\pi_{l+1}^\sigma)] > \frac{1 - e^{-1}}{k^2} \sum_{i=1}^{l+1} (k-i+1) f_{\text{avg}}(\pi_V^*) - \frac{1}{k} \sum_{i=1}^{l+1} \alpha f_{\text{avg}}(\pi_V^*).$$

Since  $\mathbb{E}_\sigma[f_{\text{avg}}(\pi_l^\sigma)]$  and  $\mathbb{E}_\sigma[f_{\text{avg}}(\pi_{l+1}^\sigma)]$  are at most  $\alpha f_{\text{avg}}(\pi_V^*)$  from the assumption, we have

$$\begin{aligned} 2\alpha f_{\text{avg}}(\pi_V^*) &> \frac{1 - e^{-1}}{k^2} \left\{ 2 \sum_{i=1}^l (k - i + 1) + (k - l) \right\} f_{\text{avg}}(\pi_V^*) - 2(\sqrt{3} - 1)\alpha f_{\text{avg}}(\pi_V^*) \\ &\geq \left\{ (4\sqrt{3} - 6)(1 - e^{-1}) - 2(\sqrt{3} - 1)\alpha \right\} f_{\text{avg}}(\pi_V^*). \end{aligned}$$

If  $\alpha \leq (2 - \sqrt{3})(1 - e^{-1})$ , it leads to the contradiction. Therefore we conclude that

$$\mathbb{E}_\sigma[f_{\text{avg}}(\pi_k^\sigma)] \geq (2 - \sqrt{3})(1 - e^{-1})f_{\text{avg}}(\pi_V^*) \geq 0.16f_{\text{avg}}(\pi_V^*).$$

□

Next we prove the theorem about the theoretical guarantee for AdaptiveSecretary algorithm.

*Proof of Theorem 4.2.* Similarly as the stream setting, let  $\pi_i^\sigma$  be the stream-based policy obtained by running AdaptiveSecretary until it selects an item from or passes the  $i$ th segment  $S_i$  under the permutation  $\sigma$  and then terminating. Let  $\Psi_i^\sigma$  be a random partial realization obtained by  $\pi_i^\sigma$ . The outline is completely the same as the proof of Theorem 4.1, but the lower bound of the expected marginal gain at each step is different.

The adaptive secretary algorithm uses the classical secretary algorithm as a subroutine. It selects the item of the maximal expected marginal gain at least with probability  $1/e$  from each segment, i.e.,

$$\mathbb{E}_\sigma[f_{\text{avg}}(\pi_i^\sigma) - f_{\text{avg}}(\pi_{i-1}^\sigma)] \geq \mathbb{E}_{\sigma, \Phi} \left[ \frac{1}{e} \max_{s \in S_i} \Delta(s | \Psi_{i-1}^\sigma) \right].$$

The lower bound of the expected marginal gain at each step is

$$\mathbb{E}_\sigma[f_{\text{avg}}(\pi_i^\sigma) - f_{\text{avg}}(\pi_{i-1}^\sigma)] \geq \frac{(1 - e^{-1})(k - i + 1)}{ek^2} f_{\text{avg}}(\pi_V^*) - \frac{1}{ek} \mathbb{E}_\sigma[f_{\text{avg}}(\pi_{i-1}^\sigma)].$$

Applying the similar evaluation to this inequality with  $l = \lfloor (-e + \sqrt{e^2 + 2e})k \rfloor$ , we conclude that

$$\mathbb{E}_\sigma[f_{\text{avg}}(\pi_k^\sigma)] \geq \frac{1 - e^{-1}}{2e\sqrt{1 + 2/e}} f_{\text{avg}}(\pi_V^*) \geq 0.08f_{\text{avg}}(\pi_V^*).$$

□