Adaptivity in Adaptive Submodularity

Hossein Esfandiari ESFANDIARI@GOOGLE.COM
Google Research, New York City, New York

Amin Karbasi AMIN.KARBASI@YALE.EDU
Yale University, New Haven, Connecticut

Vahab Mirrokni MIRROKNI@GOOGLE.COM
Google Research, New York City, New York

Editors: Mikhail Belkin and Samory Kpotufe

Abstract

Adaptive sequential decision making is one of the central challenges in machine learning and artificial intelligence. In such problems, the goal is to design an interactive policy that plans for an action to take, from a finite set of n actions, given some partial observations. It has been shown that in many applications such as active learning, robotics, sequential experimental design, and active detection, the utility function satisfies adaptive submodularity, a notion that generalizes the notion of diminishing returns to policies. In this paper, we revisit the power of adaptivity in maximizing an adaptive monotone submodular function. We propose an efficient semi adaptive policy that with $O(\log n \times \log k)$ adaptive rounds[1] of observations can achieve an almost tight $1 - 1/e - \varepsilon$ approximation guarantee with respect to an optimal policy that carries out k actions in a fully sequential manner. To complement our results, we also show that it is impossible to achieve a constant factor approximation with $o(\log n)$ adaptive rounds. We also extend our result to the case of adaptive stochastic minimum cost coverage where the goal is to reach a desired utility Q with the cheapest policy. We first prove the long-standing conjecture by Golovin and Krause [24] and show that the greedy policy achieves the asymptotically tight logarithmic approximation guarantee. We then propose a semi adaptive policy that provides the same guarantee in polylogarithmic adaptive rounds through a similar information-parallelism scheme. Our results shrink the adaptivity gap in adaptive submodular maximization by an exponential factor.

1. Introduction

Adaptive stochastic optimization under partial observability is one of the fundamental challenges in artificial intelligence and machine learning with a wide range of applications, including active learning [15], optimal experimental design [40], interactive recommendations [31], viral marketing [41], adaptive influence maximization [44], active detection [12], Wikipedia link prediction [37], and perception in robotics [30], to name a few. In such problems, one needs to adaptively make a sequence of decisions while taking into account the stochastic observations collected in previous rounds. For instance, in active learning, the goal is to learn a classifier by carefully requesting as few labels as possible from a set of unlabeled data points. Similarly, in experimental design, a practitioner may conduct a series of tests in order to reach a conclusion.

Even though it is possible to determine all the selections ahead of time before any observations take place (e.g., select all the data points at once or conduct all the medical tests simultaneously), so called *a priori selection*, it is more efficient to consider a *fully adaptive* procedure that exploits the

1. We also refer to an adaptive round as a batch query and use these two terms interchangeably throughout the paper.

information obtained from past selections in order to make a new selection. Indeed, a priori and fully sequential selections are simply two ends of a spectrum. In this paper, we develop a semi-adaptive policy that enjoys the power of a fully sequential procedure while performing exponentially fewer adaptive rounds compared to previous work. In particular, we only need poly-logarithmic number of rounds for both adaptive stochastic submodular maximization and adaptive stochastic minimum cost coverage problems. In the following, we will state these problems more formally, and then present our results in more details.

1.1. Notations

We mostly follow the notation used by Golovin and Krause [24]. Let the *ground set* $E = \{e_1, \dots, e_n\}$ be a finite set of elements (e.g., tests in medical diagnostics, data points in active learning). Each element $e \in E$ is associated with a random variable $\Phi(e) \in \Omega$ where Ω is the set of all possible outcomes. A realization of the random variable $\Phi(e)$ is denoted by $\phi(e) \in \Omega$. Note that a realization $\phi : E \to \Omega$ is simply a function from the ground set E to the outcomes Ω. For the ease of notation, we can also represent ϕ as a relation $\{(e, \omega) : \phi(e) = \omega, \forall e \in E\}$. For instance, in medical diagnosis, the element e may represent a test, such as the blood pressure, and $\Phi(e)$ its outcome, such as, high or low. Or in active learning, an item e may represent an unlabeled data point and $\Phi(e)$ may represent its label. We assume that there is a prior probability distribution $p(\phi) = p(\Phi = \phi)$ over realizations ϕ. This probability distribution encodes our uncertainty about the outcomes as well as their dependencies. In its simplest form, the outcomes maybe independent and the distribution p completely factorizes. The product distribution may very well be a valid model in the sensor placement scenario where sensors may fail to work independent of one another [2]. However, in many practical settings, such as medical diagnosis and active learning, the underlying distribution may not factorize and the outcomes may depend on each other.

In this paper, we consider adaptive strategies for picking elements where based on our observations so far, we sequentially pick an item e and observe its associated outcome $\Phi(e)$. The set of observations made so far can be represented by a *partial realization* $\psi = \{(e, \omega) : \psi(e) = \omega\}\} \subseteq E \times \Omega$. We use $\mathsf{dom}(\psi) = \{e : \exists \omega \text{ s.t } (e, \omega) \in \psi\}$ to denote the domain of ψ. We say that ψ is a *subrealization* of ψ', and denoted by $\psi \preccurlyeq \psi'$, if $\mathsf{dom}(\psi) \subseteq \mathsf{dom}(\psi')$ and $\forall e \in \mathsf{dom}(\psi)$ we have $\psi(e) = \psi'(e)$. Similarely, a partial realization ψ is *consistent* with a realization ϕ, and denote by $\psi \preccurlyeq \phi$, if they agree everywhere in the domain of ψ. We take a Bayesian approach and assume that after observing ϕ, we can compute the posterior distribution $p(\phi|\psi) = p(\Phi = \phi|\psi \preccurlyeq \Phi)$.

A *policy* $\pi : 2^{E\times\Omega} \to E$ is a partial mapping from partial observations ψ to elements E, stating which element $e \in E$ to select next[2]. Note that any deterministic policy can be visualized by a decision tree. In the proofs we also make use of two notions related to policies, namely, truncation and concatenation [24]. Given a policy π, we define the *level-k-truncation* $\pi_{[k]}$ by running π until it terminates or until it selects k items. Given two policies π_1 and π_2, we define *concatenation* $\pi_1 @ \pi_2$ as the policy obtained by first running π_1 to completion, and then running policy π_2 as if from a fresh start, ignoring the information gathered during the running of π_1. The utility of a set of observations ψ is specified through a utility function $f : 2^{E\times\Omega} \to \mathbb{R}_+$. The expected utility of a

2. Golovin and Krause [24] originally defined a policy as follows $\pi : 2^E \times \Omega^E \to E$. However, in subsequent works [11; 13], the less restrictive form, the one we consider in this paper, is used.

policy is then defined as

$$f_{avg}(\pi) = \mathbb{E}[f(S(\pi, \Phi), \Phi)] = \sum_{\phi} p(\phi) f(S(\pi, \phi), \phi),$$

where the expectation is taken with respect to $p(\phi)$. Throughout the paper $S(\pi, \phi)$ denotes the set of elements selected by policy π under realization ϕ. Without any structural assumptions, it is known that finding an optimal policy, the one that maximizes the expected utility, is notoriously hard as in many cases the utility functions are computationally intractable [39].

Adaptive submodularity [24], a generalization of diminishing returns property from sets to policies, is a sufficient condition under which a partially observable stochastic optimization problem admits (approximate) tractability. This condition ensures that the expected marginal benefit associated with any particular selection never increases as we make more observations. More formally, we define the *conditional expected marginal benefit* $\Delta(e|\psi)$ of an item e conditioned on observing the partial realization ψ as follows:

$$\begin{aligned} \Delta(e|\psi) &\doteq \mathbb{E}[f(\mathsf{dom}(\psi) \cup \{e\}, \Phi) - f(\mathsf{dom}(\psi), \Phi) | \psi \preccurlyeq \Phi] \\ &= \sum_{\omega} p(\Phi(e) = \omega | \psi)[f(\psi \cup \{e, \omega\}) - f(\psi)] \end{aligned}$$

The utility function f is *adaptive submodular* if for all subrealizations $\psi \preccurlyeq \psi'$, and all $e \in E \setminus \mathsf{dom}(\psi')$, we have

$$\Delta(e|\psi) \geq \Delta(e|\psi').$$

Moreover, we say that the utility function f is *adaptive monotone* if for all subrealizations ψ, and all $e \notin \mathsf{dom}(\psi)$ we have $\Delta(e|\psi) \geq 0$.

Whenever we use expectation notation $\mathbb{E}[\chi]$ for a random variable χ, the expectation is over all randomness of χ, unless specified otherwise. Moreover, note that we always use capital letters for random variables, and small letters for realizations. For example Ψ refers to a random variable, and ψ refers to a realization of Ψ, and hence ψ is a deterministic quantity.

1.2. Problem Formulation

The general goal in adaptive stochastic optimization is to develop policies that can maximize the expected utility while minimizing the cost of running the policy. One way to formalize it is through the *adaptive stochastic submodular maximization* problem where we aim to maximize the expected utility subject to a cardinality constraint, i.e.,

$$\pi^* = \arg\max_{\pi} f_{avg}(\pi) \qquad \text{s.t.} \qquad |S(\pi, \phi)| \leq k \qquad \text{whenever } p(\phi) > 0.$$

It is known that when the utility is adaptive submodular and adaptive monotone, the greedy policy, shown in Algorithm (1), achieves the tight $(1 - 1/e)$ approximation ratio with respect to the optimal policy [24]. This result has lead to a surge of applications in decision making problems that are amenable to myopic optimization such as active learning [25], interactive recommender systems [32], value of information [14], and active object detection [12], to name a few.

An alternative formalization is through *adaptive stochastic minimum cost coverage* where we prespecify a quota Q of utility to achieve, and aim to find a policy that achieves it with the cheapest

Algorithm 1 Adaptive Greedy Policy π_{greedy} [24] for Adaptive Stochastic Maximization

Input: Ground set E, size k, distribution $p(\phi)$, function $f(\cdot)$
initialize $A \leftarrow \emptyset$, $\psi \leftarrow \emptyset$
for $i = 1$ **to** k **do**
  $e^* = \arg\max_{e \in E \setminus A} \Delta(e|\psi)$
  $A \leftarrow A \cup \{e^*\}$
  $\psi \leftarrow \psi \cup \{(e^*, \Phi(e^*))\}$
return A

Algorithm 2 Adaptive Greedy Policy π_{greedy} [24] for Adaptive Stochastic Min Cost Coverage

Input: Ground set E, quota Q, cost function $c(\cdot)$, distribution $p(\phi)$, function $f(\cdot)$
initialize $A \leftarrow \emptyset$, $\psi \leftarrow \emptyset$
while $f(A, \psi) < Q$ **do**
  $e^* = \arg\max_{e \in E \setminus A} \Delta(e|\psi)/c(e)$
  $A \leftarrow A \cup \{e^*\}$
  $\psi \leftarrow \psi \cup \{(e^*, \Phi(e^*))\}$
return A

policy, i.e.,

$$\pi^* = \arg\min_{\pi} c_{\text{avg}}(\pi) \qquad \text{s.t.} \qquad f(S(\pi, \phi)) \geq Q \qquad \text{whenever } p(\phi) > 0,$$

where $c_{\text{avg}}(\pi) = \mathbb{E}_p[|S(\pi, \phi)|]$ is the expected number of actions a policy π selects. We can also consider a slightly more general setting where each item e has a non-negative cost $c(e)$ and replace $c_{\text{avg}}(\pi) = \mathbb{E}_p[c(S(\pi, \phi))]$ where $c(S) = \sum_{e \in S} c(e)$. Unlike the adaptive stochastic submodular maximization problem, the performance of the greedy policy, shown in Algorithm (2), is unknown for the above problem unless one makes strong assumptions about the distribution or the utility function. One of the contributions of this paper is to resolve this issue.

1.3. Our Contributions

Fully sequential policies benefit from previous observations in order to make informed decisions. In many scenarios, however, it is more effective (and sometimes the only way) to select multiple elements in parallel and observe their realizations together. Examples include crowdsourcing (where a single task consists of a collection of unlabeled data to be labeled altogether), multi-stage viral marketing (where in each stage a subset of nodes are chosen as seed nodes), batch-mode pool-based active learning (where the label of a set of data points are requested simultaneously), or medical diagnosis (where there is a shared cost among experiments). A batch-mode, semi-adaptive policy is a mix of a priori and fully sequential selections. The focus of this paper is to answer the following question in the context of adaptive stochastic optimization:

> How many adaptive rounds of observations are needed in order to be competitive to an optimal and fully sequential policy?

We answer the above question in the context of adaptive submodularity. In this paper, we consider two adaptive stochastic optimization problems, namely, adaptive stochastic maximization and adaptive stochastic minimum cost cover. We re-examine the required amount of adaptivity in order to be competitive to the optimal and fully sequential policy. In particular, we show the following results in the information-parallel stochastic optimization when the utility function is adaptive submodular and adaptive monotone.

- For the adaptive stochastic submodular maximization problem, we develop a semi adaptive policy that with $O(\log(n)\log(k))$ adaptive rounds (a.k.a., batch queries) achieves the tight

$(1 - 1/e - \varepsilon)$ approximation guarantee with respect to the optimum policy π^* that selects k items fully sequentially, i.e., $f_{\text{avg}}(\pi) \geq (1 - 1/e - \varepsilon) f_{\text{avg}}(\pi^*)$.

- We complement the above result by showing that no policy can achieve a constant factor approximation guarantee with fewer than $o(\log(n))$ adaptive rounds. Moreover, the approximation guarantee of any semi adaptive policy that chooses batches of fixed size r will degrade with a factor of $O(r/\log^2(r))$.

- For the adaptive stochastic minimum cost coverage problem, we show that the greedy policy achieves an asymptoticly tight logarithmic approximation guarantee, effectively proving [24]'s conjecture. More precisely, we show that $c_{\text{avg}}(\pi_{\text{greedy}}) \leq (c_{\text{avg}}(\pi^*) + 1) \log\left(\frac{nQ}{\eta}\right) + 1$ where we make the common assumption that there is a value η such that $f(\psi) > Q - \eta$ implies that $f(\psi) = Q$ for all partial realizations ψ.

- We also develop a semi adaptive policy for the adaptive stochastic minimum cost coverage problem that achieves the same logarithmic approximation guarantee with $O\big(\log n \log(Qn/\eta)\big)$ adaptive rounds.

2. Related Work

Submodularity captures an intuitive diminishing returns property where the gain of adding an element to a set decreases as the set gets larger. More formally, a non-negative set function $f : 2^V \to \mathbb{R}_+$ is **submodular** if for all sets $A \subseteq B \subset V$ and every element $e \in V \setminus B$, we have

$$f(A \cup \{e\}) - f(A) \geq f(B \cup \{e\}) - f(B).$$

Submodular maximization has found numerous applications in machine learning and artificial intelligence [43], including neural network interpretation [17], data summarization [35], crowd teaching [42], privacy [36], fairness [8], and adversarial attacks in deep neural nets [34]. Moreover, in many information gathering and sensing scenarios, the objective functions satisfy submodularity [33; 45; 28]. However, the classic notion of submodularity falls short in interactive information acquisition settings as it requires the decision maker to commit to all of her selections ahead of time, in an open-loop fashion [27].

To circumvent this issue, [24] proposed adaptive submodularity, a generalization of submodularity from sets to policies. Like submodularity, adaptive submodularity is a sufficient condition that ensures tractability in adaptive settings. More precisely, in the adaptive stochastic submodular maximization problem, when the objective function is adaptive monotone and adaptive submodular, the greedy policy achieves the tight $(1 - 1/e)$ approximation guarantee with respect to an optimum policy [24]. More generally, [26] proposed a random greedy policy that not only retains the aforementioned $(1 - 1/e)$ approximation ratio in the monotone setting, but also provides a $(1/e)$ approximation ratio for the non-monotone adaptive submodular functions.

The results for adaptive stochastic minimum cost coverage problem are much weaker. Originally, [24] claimed that the greedy policy also achieves a logarithmic approximation factor but as pointed out by [38] the proof was flawed. Instead, under stronger conditions, namely, strong adaptive submodularity and strong adaptive monotonicity, Golovin and Krause proposed a new proof, with a squared-logarithmic factor approximation, i.e., $c_{\text{avg}}(\pi_{\text{greedy}}) \leq c_{\text{avg}}(\pi^*) \left(\log\left(\frac{nQ}{\eta}\right) + 1\right)^2$.

Note that there are some fundamental technical differences between the notions of submodularity and adaptive submodularity. For example, submodularity preserves under truncation, that is, if a function $f(\cdot)$ is submodular, for any constant c, $\min\big(f(\cdot), c\big)$ is submodular as well. This comes very handy in designing algorithms for submodular functions and often used as a simple way to reduce minimum cost coverage to submodular maximization. However, unfortunately, truncation does not preserve adaptive submodularity (See Appendix B) and thus all the previous attempts to use this reduction are futile. In this paper, we prove the original conjecture of Golovin and Krause [24] and show that under adaptive submodularity (without resorting to stronger conditions), the greedy policy achieves a logarithmic approximation factor, namely, $c_{\text{avg}}(\pi_{\text{greedy}}) \leq (c_{\text{avg}}(\pi^*) + 1) \log\left(\frac{nQ}{\eta}\right) + 1$.

The main focus of this paper is to explore the information parallelism, a.k.a., batch-mode, stochastic optimization [29; 23]. Many active learning problems naturally fall into this setting when it is more cost-effective to request labels in large batches, rather than one-at-a-time (for detailed discussions, we refer the interested reader to [10]). Note that the two extremes of batch-mode stochastic optimization are full batch setting (i.e., all selections are done in a single batch, and hence the batch-mode setting reduces to the non-adaptive, open-loop optimization problem) and full sequential setting (i.e., elements are selected one-by-one in a closed-loop manner where each selection is based on the results of all previous selections). In this paper, we lay out a rigorous foundation for the semi-adaptive setting where elements are selected in a sequential and closed-loop way but with multiple selections at each round.

There are a few partial results regarding the semi-adaptive policy for the adaptive stochastic minimum cost coverage problem. In particular, [11] proposed a policy that selects batches of fixed size r and proved that under strong adaptive submodularity and strong adaptive monotonicity, this policy achieves a poly-logarithmic approximation to an optimal policy that is also constrained to picking up batches of size r. Note that this result does not provide any guarantees with respect to the actual baseline, namely, the optimal and fully sequential policy. Moreover, [10] showed that this policy has a sublinear-approximation[3] guarantee against the fully sequential policy. In fact, we show that for the adaptive stochastic submodular maximization problem, the approximation factor of a fixed batch-policy suffers by at least a factor of $\log^2(r)/r$ in the worst case, so unless r is a fixed constant, no constant factor approximation guarantee is possible.

Back to the adaptive stochastic minimum cost coverage, when the distribution p is fully factorized (i.e., the outcomes are independent), [1] very recently showed that there exists a policy that, using $O(\log(Q)/\log\log(Q))$ rounds of adaptivity, achieves a poly-logarithmic approximation to the optimal sequential policy. In this paper, we propose a (batch-mode) semi adaptive policy that, using only polylogarithmic adaptive rounds, achieves an asymptotically tight logarithmic approximation to the fully sequential policy for general adaptive monotone submodular functions (we do not need to resort to stronger notions of adaptivity and monotonicity).

To the best of our knowledge, no results are known for semi-adaptive policies for the adaptive stochastic submodular maximization problem. We develop a (batch-mode) semi adaptive policy that achieves an almost tight $1 - 1/e - \varepsilon$ approximation guarantee with only polylogarithmic adaptive rounds. We also complement our result by showing that no semi-adaptive policy can achieve a constant factor approximation to the optimal policy by fewer than $o(\log(n))$ adaptive rounds.

3. Unfortunately, the approximation factor grows polynomially in r. Moreover, note that this result assumes strong adaptive submodularity and strong adaptive monotonicity.

Our work is also related to the adaptivity complexity of submodular maximization, which refers to the number of parallel rounds required to achieve a constant factor approximation guarantee in the offline, open-loop setting. [3] developed a parallel algorithm that $O(\log n)$ rounds finds a solution with an approximation arbitrarily close to $\frac{1}{3}$ which was soon improved to $(1-\frac{1}{e}-\epsilon)$-approximation [21; 6; 18]. The adaptivity complexity was also studied in the non-monotone submodular maximization [5; 20; 9], convex minimization [4; 16; 7] and multi-armed bandit [22; 19]. We lift the notion of adaptivity complexity from the offline optimization to the interactive setting where instead of parallelizing the optimization steps we parallelize the information acquisition.

3. Greedy Versus Optimum

Throughout the paper, we assume that the utility function f is an adaptive monotone and adaptive submodular function with respect to the distribution $p(\phi)$. In this section, we show how the expected utility obtained by the greedy policy π_{greedy} is related to the expected utility obtained by the optimum policy π^*. Let us define π^τ to be the policy that runs the greedy policy π_{greedy} and stops when the expected marginal gain of every single remaining element is less than or equal to τ. We define τ_i to be a threshold such that the expected number of elements selected by π^{τ_i} is i, in other words $\sum_\phi p(\phi)|S(\pi^{\tau_i},\phi)| = i$.[4] Also remember that for two policies π and π' we define $\pi@\pi'$ to be a policy that first runs π and then runs π' from a fresh start (i.e., ignoring the information gathered by π). This definition implies $S(\pi@\pi',\phi) = S(\pi,\phi)\cup S(\pi',\phi)$. The following is the key lemma of this paper.

Lemma 1 *For any policy π^* and any positive integer ℓ we have*

$$f_{avg}(\pi^{\tau_\ell}) > \left(1 - e^{-\frac{\ell}{\mathbb{E}[K]+1}}\right) f_{avg}(\pi^*),$$

where K is a random variable that indicates the number of items picked by π^*, i.e. $K = |S(\pi^*,\Phi)|$.

Before proving Lemma 1, we provide some primitives that we use in the proof of this lemma as well as Lemmas 6 and 14. For a randomized policy π, we use Θ_π to indicate the random bits of the policy π. We use θ_π to indicate a realization of Θ_π and use $p(\theta_\pi)$ to indicate the probability that θ_π is realized. We drop π from the notation when it is clear from the context.

For a deterministic policy π, a (potentially) randomized policy π^*, an element e, and two subrealization ψ and $\psi \preccurlyeq \psi'$, we define the event $<\psi,\psi',e,\pi,\pi^*>$ to be the event that

- $\mathsf{dom}(\psi) = \mathsf{dom}(\pi)$, meaning that the policy π selects exactly all of the elements of ψ,
- and, the set of elements selected by the policy $\pi@\pi^*$, at some point during its run, coincides exactly with the domain of ψ',
- and, right after the policy $\pi@\pi^*$ selects all the elements of ψ', it chooses e.

4. Use an arbitrary tie breaking rule to make it exactly equal. For example, if by accepting elements whose expected marginal benefit is strictly larger than τ the greedy policy selects $\alpha < i$ elements in expectation, and by accepting elements with expected marginal benefit larger than or equal to τ the greedy policy selects $\beta > i$ elements in expectation, then we let the policy π^{τ_i} accepts the elements with a marginal gain of τ with probability $\frac{i-\alpha}{\beta-\alpha}$. Hence, the expected number of items accepted by π^{τ_i} is $\alpha + \frac{i-\alpha}{\beta-\alpha}(\beta-\alpha) = i$.

Note that for a fixed θ_{π^*}, and conditioned on $\psi' \preccurlyeq \Phi$, we can simulate $\pi@\pi^*$ and deterministically indicate whether the event $< \psi, \psi', e, \pi, \pi^* >$ happens or not.[5] We define $\mathbb{1}_{\psi',\psi,e,\pi,\pi^*,\theta_{\pi^*}}$ to be a *deterministic* binary variable that is 1 if and only if, conditioned on $\Theta_{\pi^*} = \theta_{\pi^*}$ and $\psi' \preccurlyeq \Phi$, the event $< \psi, \psi', e, \pi, \pi^* >$ happens. We use the shorthand $\mathbb{1}_{\psi',\psi,e,\theta}$ (and drop the notations of policies) when it is clear from the context.

Proof of Lemma 1: First we provide an upper bound on $f_{avg}(\pi^*)$. Pick an arbitrary number $i \in \{0, \ldots, \ell\}$. Note that by adaptive monotonicity, we have $f_{avg}(\pi^*) \leq f_{avg}(\pi^{\tau_i}@\pi^*)$. Next we show that $f_{avg}(\pi^{\tau_i}@\pi^*) \leq f_{avg}(\pi^{\tau_i}) + \mathbb{E}[K]\,\Delta^{\tau_i}(\pi)$, where $\Delta^{\tau_i}(\pi) = f_{avg}(\pi^{\tau_i}) - f_{avg}(\pi^{\tau_{i-1}})$. This implies that

$$f_{avg}(\pi^*) \leq f_{avg}(\pi^{\tau_i}@\pi^*) \leq f_{avg}(\pi^{\tau_i}) + \mathbb{E}[K]\,\Delta^{\tau_i}(\pi). \tag{1}$$

Note that once we run π^{τ_i} the set of selected elements and their realizations depends on Φ. To capture this randomness, we let Ψ_i be a random variable that indicates the partial realization observed by running policy π^{τ_i}. We use ψ_i to indicate a realization of the random variable Ψ_i. Note that, by definition of π^{τ_i} we have $\tau_i \geq \max_{e \in E \setminus \mathsf{dom}(\psi_i)} \Delta(e|\psi_i)$. Moreover, by adaptive submodularity, for all sub-realizations ψ' such that $\psi_i \preccurlyeq \psi'$, and for all $e \in E \setminus \mathsf{dom}(\psi')$ we have $\Delta(e|\psi_i) \geq \Delta(e|\psi')$. Therefore, we have

$$\tau_i \geq \Delta(e|\psi'), \qquad \forall \psi' \text{ s.t. } \psi_i \preccurlyeq \psi', \text{ and } \forall e \in E \setminus \mathsf{dom}(\psi'). \tag{2}$$

Now, we can bound the different between the expected utility obtained by $\pi^{\tau_i}@\pi^*$ and π^{τ_i} as follows:

$$\begin{aligned}
&f_{avg}(\pi^{\tau_i}@\pi^*) - f_{avg}(\pi^{\tau_i})\\
&= \sum_{\psi_i \in \Psi_i} \sum_{\psi_i \preccurlyeq \psi'} \sum_{e \notin \mathsf{dom}(\psi')} \sum_{\theta \in \Theta_{\pi^*}} p(\psi')p(\theta)\mathbb{1}_{\psi_i,\psi',e,\theta}\mathbb{E}\left[f(\psi' \cup \Phi(e)) - f(\psi')|\psi' \preccurlyeq \Phi \text{ and } \Theta_{\pi^*} = \theta\right]\\
&= \sum_{\psi_i \in \Psi_i} \sum_{\psi_i \preccurlyeq \psi'} \sum_{e \notin \mathsf{dom}(\psi')} \sum_{\theta \in \Theta_{\pi^*}} p(\psi')p(\theta)\mathbb{1}_{\psi_i,\psi',e,\theta}\mathbb{E}\left[f(\psi' \cup \Phi(e)) - f(\psi')|\psi' \preccurlyeq \Phi\right]\\
&= \sum_{\psi_i \in \Psi_i} \sum_{\psi_i \preccurlyeq \psi'} \sum_{e \notin \mathsf{dom}(\psi')} \sum_{\theta \in \Theta_{\pi^*}} p(\psi')p(\theta)\mathbb{1}_{\psi_i,\psi',e,\theta}\Delta(e|\psi')\\
&\leq \sum_{\psi_i \in \Psi_i} \sum_{\psi_i \preccurlyeq \psi'} \sum_{e \notin \mathsf{dom}(\psi')} \sum_{\theta \in \Theta_{\pi^*}} p(\psi')p(\theta)\mathbb{1}_{\psi_i,\psi',e,\theta}\tau_i\\
&\leq \sum_{\psi_i \in \Psi_i} \sum_{\psi_i \preccurlyeq \psi'} \sum_{e \notin \mathsf{dom}(\psi')} \sum_{\theta \in \Theta_{\pi^*}} p(\psi')p(\theta)\mathbb{1}_{\psi_i,\psi',e,\theta}\Delta^{\tau_i}(\pi)\\
&= \Big(\sum_{\psi_i \in \Psi_i} \sum_{\psi_i \preccurlyeq \psi'} \sum_{e \notin \mathsf{dom}(\psi')} \sum_{\theta \in \Theta_{\pi^*}} p(\psi')p(\theta)\mathbb{1}_{\psi_i,\psi',e,\theta}\Big)\Delta^{\tau_i}(\pi)\\
&= \mathbb{E}[K]\,\Delta^{\tau_i}(\pi),
\end{aligned}$$

where the first equality is by definition, the second equality is by independency[6], the third equality is by definition, forth inequality is by Inequality 2 and the fifth inequality is by $\Delta^{\tau_i}(\pi) \geq \tau_i$. This

5. Note that if policy $\pi@\pi^*$ attempts to query an element that does not exist in ψ', prior to querying e, we know that the event $< \psi, \psi', e, \pi, \pi^* >$ does not happen and we do not need to simulate the policy any further.

6. Note that ψ' is fixed, hence $f(\psi' \cup \Phi(e)) - f(\psi')$ is independent of Θ.

proves inequality (1) as promised. Let us define $\Delta_i^* = f_{avg}(\pi^*) - f_{avg}(\pi^{\tau_i})$. Inequality 1 implies that $\Delta_i^* \leq \mathbb{E}[K](\Delta_{i-1}^* - \Delta_i^*)$. By a simple rearrangement we have $\Delta_i^* \leq \big(1 - \frac{1}{\mathbb{E}[K]+1}\big)\Delta_{i-1}^*$. By iteratively applying this inequality we have

$$\Delta_\ell^* \leq \big(1 - \frac{1}{\mathbb{E}[K]+1}\big)^\ell \Delta_0^* \leq e^{-\frac{\ell}{\mathbb{E}[K]+1}} \Delta_0^*.$$

By applying the definition of Δ_i^* and further rearrangements we have $f_{avg}(\pi^{\tau_\ell}) > \big(1 - e^{-\frac{\ell}{\mathbb{E}[K]+1}}\big) f_{avg}(\pi^*)$, as desired. □

Next, we will use Lemma 1 to prove the conjecture by [24] that the greedy policy achieves the asymptotically tight logarithmic approximation guarantee.

4. Adaptive Stochastic Minimum Cost Coverage

In this section, we show that the greedy policy, outlined in Algorithm 2, achieves a logarithmic approximation guarantee for adaptive stochastic minimum cost coverage. For the ease of presentation, we focus on the unit cost case, i.e., $c(e) = 1$ for all $e \in E$. The generalization to the non-uniform cost is immediate. To prove this theorem we use Lemma 1 and follow the usual proof for set cover.

Theorem 2 *Assume that there is a value $\eta \in (0, Q]$ such that $f(\psi) > Q - \eta$ implies $f(\psi) = Q$ for all ψ. Let π^* be an arbitrary policy (including the optimum policy[7]) that covers everything i.e., $f(\pi^*) = Q$ for all ϕ. Let π_{greedy} be the greedy policy, and let $n = |E|$. We have*

$$c_{avg}(\pi_{greedy}) \leq \big(c_{avg}(\pi^*) + 1\big) \log\left(\frac{nQ}{\eta}\right) + 1.$$

Proof Let K be a random variable that indicates the number of items picked by π^*, i.e., $K = S(\pi^*, \Phi)$. Set $\ell = (\mathbb{E}[K] + 1)\log(nQ/\eta)$. Note that by definition of π^* we have $f(\pi^*) = Q$ for all ϕ, hence we have $f_{avg}(\pi^*) = Q$. By Lemma 1 we have

$$\begin{aligned}
f_{avg}(\pi^{\tau_\ell}) &> \left(1 - e^{-\frac{\ell}{\mathbb{E}[K]+1}}\right) f_{avg}(\pi^*) && \text{By Lemma 1} \\
&= \left(1 - e^{-\log(nQ/\eta)}\right) f_{avg}(\pi^*) && \text{Since } \ell = \big(\mathbb{E}[K] + 1\big)\log(nQ/\eta) \\
&= \left(1 - \frac{\eta}{nQ}\right) f_{avg}(\pi^*) \\
&= Q - \frac{\eta}{n}. && \text{Since } f_{avg}(\pi^*) = Q
\end{aligned}$$

Recall that by definition we have $\mathbb{E}[f(\pi^{\tau_\ell})] = f_{avg}(\pi^{\tau_\ell}) = Q - \frac{\eta}{n}$. Moreover, by adaptive monotonicity we have $f(\pi^{\tau_\ell}) \leq f(\phi) = Q$. Hence by Markov inequality with probability $1 - 1/n$ we have $f(\pi^{\tau_\ell}) > Q - \eta$. By definition of η this implies that with probability $1 - \frac{1}{n}$ we have $f(\pi^{\tau_\ell}) = Q$. Therefore, with probability $1 - 1/n$, the policy π^{τ_ℓ} reaches the utility Q after selecting

7. One can think of this as an optimal policy that minimizes the expected number of selected items and guarantees that every realization is covered.

$\ell = \big(\mathbb{E}[K]+1\big)\log(nQ/\eta)$ items in expectation. Otherwise, π_{greedy} picks at most all the n items. Hence the expected number of items that π_{greedy} picks is upper bounded by

$$\left(1-\frac{1}{n}\right)\times\big(\mathbb{E}[K]+1\big)\log\left(\frac{nQ}{\eta}\right)+\frac{1}{n}\times n \le \big(c_{avg}(\pi^*)+1\big)\log\left(\frac{nQ}{\eta}\right)+1.$$

■

With slight modifications to the proof of Theorem 2 we achieve the following corollary. We provide the proof of this result in Appendix A.1.

Corollary 3 *Assume that there is a value $\eta \in (0, Q]$ such that $f(\psi) > Q - \eta$ implies $f(\psi) = Q$ for all ψ. Let π^* be an arbitrary policy that covers everything i.e. $f(\pi^*) = Q$ for all ϕ. Let π_{greedy} be the greedy policy, and let $\delta = \min_\phi p(\phi)$. We have*

$$c_{avg}(\pi_{greedy}) \le \big(c_{avg}(\pi^*)+1\big)\log\left(\frac{Q}{\delta\eta}\right)+1.$$

5. Semi Adaptive Stochastic Submodular Maximization

In this section, we provide a policy for adaptive stochastic submodular maximization that makes only $O(\log n \log k)$ batch queries (a.k.a. adaptive rounds). We show that our policy provides a $(1-\frac{1}{e}-\varepsilon)$ approximate solution compare to that of the best fully sequential policy. Our policy is based on two notions, *semi-adaptive values* and *information gap*. We first provide some intuition and notations, and then explicitly define these two notions. At any stage of the algorithm, semi-adaptive value of an item e is our estimate of the expected value of selecting item e. We provide these estimations based on the information of the last batch query that we carried out and the set of items that we are deciding to select but not queried yet. Information gap is our estimate of the accuracy of the maximum semi-adaptive value. We use the information gap to balance between the loss on the performance and the number of batch queries that we make. We iteratively and greedily

Algorithm 3 Semi Adaptive Greedy Policy for Adaptive Submodular Maximization

Input: Ground set E, size k, distribution $p(\phi)$, function $f(\cdot)$, value $\epsilon > 0$
initialize $A \leftarrow \emptyset$, $\psi' \leftarrow \emptyset$, $i \leftarrow 1$
while $|A| < k$ **do**
  while $\text{IG}(i, \psi') \ge (1-\epsilon)$ and $|A| < k$ **do**
    $e^* = \arg\max_{e \in E \setminus A} \text{SAV}(e, i, \psi')$
    $A \leftarrow A \cup \{e^*\}$
    $i \leftarrow i + 1$
  $\psi'' \leftarrow$ query all elements in $A \setminus \mathsf{dom}(\psi')$
  $\psi' \leftarrow \psi' \cup \psi''$
return A

Algorithm 4 Semi Adaptive Greedy Policy for Minimum Cost Coverage

Input: Ground set E, quota Q, distribution $p(\phi)$, function $f(\cdot)$, value $\epsilon > 0$
initialize $A \leftarrow \emptyset$, $\psi' \leftarrow \emptyset$, $i \leftarrow 1$
while $f(A, \psi') < Q$ **do**
  while $\text{RIG}(i, \psi') \ge (1-\epsilon)$ **do**
    $e^* = \arg\max_{e \in E \setminus A} \text{SAV}(e, i, \psi')$
    $A \leftarrow A \cup \{e^*\}$
    $i \leftarrow i + 1$
  $\psi'' \leftarrow$ query all elements in $A \setminus \mathsf{dom}(\psi')$
  $\psi' \leftarrow \psi' \cup \psi''$
return A

select elements based on their semi-adaptive values. We continue this selection non-adaptively until the information gap decreases to $(1-\varepsilon)$. When the information gap drops below $(1-\varepsilon)$ we *query* all the selected elements. We call this algorithm *semi-adaptive greedy* and is outlined in Algorithm 3. In this section, we use π to refer to this policy. We refer to the step i of a policy as the time it selects the i-th element. We use Ψ_i^π to refer to the (random) partial realization up to and including step i of policy π. Again, note that the partial realization we observe by running π depends on Φ. We also use ψ' to refer to the observed partial realization of items that have been queried so far. Since policy π is deterministic, given ψ', we can deterministically indicate the domain of Ψ_i^π. Therefore, $\mathsf{dom}(\Psi_i^\pi)$ is deterministic and well specified (while the state of items $e \in \mathsf{dom}(\Psi_i^\pi) \setminus \mathsf{dom}(\psi')$ is random). We are ready to define the semi-adaptive values and the information gap.

Definition 4 (Semi-Adaptive Value) *At any step i of the policy π, and given the partial realization $\psi' \preccurlyeq \Psi_i^\pi$, the semi-adaptive value of an item $e \in E \setminus \mathsf{dom}(\Psi_i^\pi)$ is defined as follows:*

$$\mathrm{SAV}(e, i, \psi') \doteq \mathbb{E}_{\psi' \preccurlyeq \Psi_i^\pi}\big[\Delta(e|\Psi_i^\pi)\big].$$

Note that the semi-adaptive value of an item e is equal to the expected marginal gain of e over all the unknown random realizations (i.e., not in ψ').

Definition 5 (Information Gap) *At any step i of the policy π, and given the partial realization $\psi' \preccurlyeq \Psi_i^\pi$, the information gap is defined as follows.*

$$\mathrm{IG}(i, \psi') \doteq \frac{\max_{e \notin \mathsf{dom}(\Psi_i^\pi)} \mathbb{E}_{\psi' \preccurlyeq \Psi_i^\pi}\big[\Delta(e|\Psi_i^\pi)\big]}{\mathbb{E}_{\psi' \preccurlyeq \Psi_i^\pi}\big[\max_{e \notin \mathsf{dom}(\Psi_i^\pi)} \Delta(e|\Psi_i^\pi)\big]}.$$

Equipped with these definitions, we show next the utility obtained by the semi adaptive greedy policy, shown in Algorithm 3, along with the total number of batch queries.

5.1. Performance

Golovin and Krause [24] showed that a fully sequential greedy policy π_{greedy} achieves $f_{avg}(\pi_{\text{greedy}[\ell]}) > \big(1-e^{-\frac{\ell}{k}}\big) f_{avg}(\pi^*)$. The next lemma bounds the performance of our semi adaptive greedy policy in a similar fashion. We use the notions of semi-adaptive values and the information gap to prove this lemma. In the following lemma, π is the semi-adaptive greedy policy and $\pi_{[\ell]}$ is a policy that runs π and stops if it selects ℓ items.

Lemma 6 *Let π be the semi-adaptive greedy policy. For any policy π^* and positive integer ℓ we have*

$$f_{avg}(\pi_{[\ell]}) > \big(1 - e^{-\frac{\ell}{k}} - \varepsilon\big) f_{avg}(\pi^*).$$

This is one of the main technical contributions of the paper, but due to the space constraint we provide the proof of this lemma in Appendix A.2. The proof of this lemma relies on Lemma 1 and uses a similar machinery.

5.2. Query Complexity

In this subsection, we bound the number of batch queries of the semi-adaptive greedy policy. We define random variable Ψ'_t to be the partial realization obtained by the t-th batch query (do not confuse it with Ψ_i^π). We use ψ'_t to indicate the realization of random variable Ψ'_t. The next lemma shows that after any $\log_{\frac{1}{1-\varepsilon/2}}\left(\frac{n}{\delta}\right) = O_\varepsilon\big(\log\left(\frac{n}{\delta}\right)\big)$ batch queries, the maximum expected marginal benefit drops by a factor $(1-\frac{\varepsilon}{2})$, with high probability. We later apply this lemma iteratively for $O(\log k)$ times to show that after $O_\varepsilon(\log n \log k)$ batch queries, the maximum expected marginal gain is vanishingly small.

Lemma 7 *Pick an arbitrary t and fix partial realization ψ'_t. Let $\Delta'_t = \max_{e\notin \mathsf{dom}(\psi'_t)} \Delta(e|\psi'_t)$, and let $t^+ = t + \log_{\frac{1}{1-\varepsilon/2}}\left(\frac{n}{\delta}\right)$. With probability at least $1-\delta$ we have*

$$\max_{e\notin \mathsf{dom}(\Psi'_{t^+})} \Delta(e|\Psi'_{t^+}) \le \left(1-\frac{\varepsilon}{2}\right)\Delta'_t.$$

Proof For any $t' \ge t$ we use the random variable $S_{t'}$ to indicate the set of elements such that $\Delta(e|\Psi'_{t'}) \ge (1-\frac{\varepsilon}{2})\Delta'_t$. To prove the above lemma, we show that $\mathbb{E}\left[|S_{t'}|\right] \le (1-\frac{\varepsilon}{2})\mathbb{E}\left[|S_{t'+1}|\right]$. This together with $|S_t| \le n$ implies that $\mathbb{E}\left[|S_{t^+}|\right] \le \delta$ for $t^+ = t + \log_{\frac{1}{1-\varepsilon/2}}\left(\frac{n}{\delta}\right)$. Note that $|S_{t^+}|$ is a non-negative integer, and hence we have $S_{t^+} = \emptyset$ with probability at least $1-\delta$.

Next, we show that $\mathbb{E}\left[|S_{t'}|\right] \le (1-\frac{\varepsilon}{2})\mathbb{E}\left[|S_{t'+1}|\right]$. First note that by adaptive monotonicity $e \in S_{t'+1}$ implies $e \in S_{t'}$, and hence we have $S_{t'+1} \subseteq S_{t'}$. In the following, we use the notion of information gap and show that for any element $e \in S_{t'}$, we have $e \notin S_{t'+1}$ with probability at least $\frac{\varepsilon}{2}$. This directly implies $\mathbb{E}\left[|S_{t'}|\right] \le (1-\frac{\varepsilon}{2})\mathbb{E}\left[|S_{t'+1}|\right]$ as desired. Note that when we query $\Psi'_{t'+1}$ the information gap is at most $(1-\varepsilon)$. Hence, for some Ψ_i^π (which corresponds to $\Psi'_{t'+1}$) we have

$$\begin{aligned}\max_{e\notin \mathsf{dom}(\Psi_i^\pi)} \mathbb{E}_{\psi'_{t'}\preccurlyeq\Psi_i^\pi}\big[\Delta(e|\Psi_i^\pi)\big] &\le (1-\varepsilon)\mathbb{E}_{\psi'_{t'}\preccurlyeq\Psi_i^\pi}\big[\max_{e\notin \mathsf{dom}(\Psi_i^\pi)} \Delta(e|\Psi_i^\pi)\big] && \text{information gap}\\ &\le (1-\varepsilon)\Delta'_t. && \text{by adaptive monotonicity}\end{aligned}$$

This implies that for all $e \notin \mathsf{dom}(\Psi_i^\pi)$, with probability at least $\frac{\varepsilon}{2}$, we have $\Delta(e|\Psi_i^\pi) \le (1-\frac{\varepsilon}{2})\Delta'_t$. Therefore, for any element $e \in S_{t'}$, we have $e \notin S_{t'+1}$ with probability at least $\frac{\varepsilon}{2}$, as promised. ∎

Now, we are ready to prove the main theorem of this section. In the following theorem, π is the semi-adaptive greedy policy, $\pi_{[\ell]}$ is a policy that runs π and stops if it selects ℓ items and $\pi_{[\ell]}^T$ is a policy that runs $\pi_{[\ell]}$ and stops if it makes T batch queries.

Theorem 8 *Let π be the semi-adaptive greedy policy and let $\pi_{[\ell]}^T$ be a policy that runs $\pi_{[\ell]}$ and stops if it makes T batch queries. For any policy π^* (including the optimum policy) and any positive integer ℓ and for some $T \in O_\varepsilon(\log n \log \ell)$ we have*

$$f_{avg}(\pi_{[\ell]}^T) > \big(1 - e^{-\frac{\ell}{k}} - 3\varepsilon\big) f_{avg}(\pi^*).$$

Proof Let us set $\delta = \frac{\varepsilon}{\log_{\frac{1}{1-\varepsilon/2}}(\frac{\ell}{\varepsilon})}$ and let $\Delta_1(\pi)$ be the expected marginal benefit of the first selected item. By applying Lemma 7 iteratively $\log_{\frac{1}{1-\varepsilon/2}}(\frac{\ell}{\varepsilon})$ times we have

$$\max_{e\notin \mathsf{dom}(\Psi'_T)} \Delta(e|\Psi'_T) \le \big(1-\frac{\varepsilon}{2}\big)^{\log_{\frac{1}{1-\varepsilon/2}}(\frac{\ell}{\varepsilon})} \Delta_1(\pi) = \frac{\varepsilon}{\ell}\Delta_1(\pi),$$

with probability $1 - \delta \times \log_{\frac{1}{1-\varepsilon/2}}(\frac{\ell}{\varepsilon}) = 1 - \varepsilon$. This means that with probability $(1-\varepsilon)$ the total expected marginal benefit of the elements added after the T-th batch query is at most

$$\ell \times \frac{\varepsilon}{\ell}\Delta_1(\pi) = \varepsilon\Delta_1(\pi) \le \varepsilon f_{avg}(\pi_{[\ell]}),$$

where $T = \log_{\frac{1}{1-\varepsilon/2}}\left(\frac{n}{\delta}\right) \times \log_{\frac{1}{1-\varepsilon/2}}(\frac{\ell}{\varepsilon}) \in O_\varepsilon(\log n \log \ell)$. This, together with lemma 6, implies that if we stop policy π after $T \in O_\varepsilon(\log n \log \ell)$ batch queries for any policy π^* we have

$$f_{avg}(\pi_{[\ell]}) > \left(1 - e^{-\frac{\ell}{\mathbb{E}[k]}} - 3\varepsilon\right) f_{avg}(\pi^*),$$

as desired. ■

6. Semi Adaptive Stochastic Minimum Cost Coverage

In this section, we bound the efficiency and round complexity of the semi-adaptive greedy policy, outlined in Algorithm 4. To simplify the proofs, we use a more restricted notion of information gap. It is easy to observe that the same proofs in the previous section hold using this version of information gap as well.[8]

Definition 9 (Restricted Information Gap) *At any step i of the policy π, and given the partial realization $\psi' \preccurlyeq \Psi_i^\pi$, the restricted information gap is defined as follows.*

$$\text{RIG}(i, \psi') = \frac{\max_{e \notin \mathsf{dom}(\psi')} \Delta(e|\psi')}{\mathbb{E}_{\psi' \preccurlyeq \Psi_i^\pi}\left[\max_{e \notin \mathsf{dom}(\Psi_i^\pi)} \Delta(e|\Psi_i^\pi)\right]}.$$

6.1. Performance

Next theorem bounds the performance of our policy. The proof of this theorem is a combination of the ideas in Lemma 1, Lemma 6 and Theorem 2 and is presented in Appendix A.4.

Theorem 10 *Assume that there is a value $\eta \in (0, Q]$ such that $f(\psi) > Q - \eta$ implies $f(\psi) = Q$ for all ψ. Let π^* be an arbitrary policy that covers everything, i.e., $f(\pi^*) = Q$ for all ϕ. Let π be the semi-adaptive greedy policy, outlined in Algorithm 4. We have*

$$c_{avg}(\pi) \le \left(\frac{c_{avg}(\pi^*) + 1}{1-\varepsilon}\right) \log\left(\frac{nQ}{\eta}\right) + 1.$$

6.2. Query Complexity

Next, we bound the number of batch queries of the semi-adaptive greedy policy. We use Lemma 7 presented in the previous section together with Theorem 10 to prove the above theorem. The proof of this theorem follows the proof of Theorem 8 and is presented in Appendix A.5.

Theorem 11 *Let $\eta \in (0, Q]$ be a value such that $f(\psi) > Q - \eta$ implies $f(\psi) = Q$ for all ψ. Let π^* be any policy that covers everything, i.e., $f(\pi^*) = Q$ for all ϕ. Let π be the semi-adaptive greedy policy (Algorithm 4) and let π^T be a policy that runs π and stops if it makes T batch queries. For some $T \in O\big(\log n \log(Qn/\eta)\big)$ we have $f(\pi^T) = Q$ with probability at least $1 - 1/n$.*

8. We use this notion in this section for simplicity. However, since the previous notion of information gap is more intuitive, we keep the previous notion as well.

7. Hardness

Next theorem states our hardness result.

Theorem 12 *Any policy with a constant approximation guarantee for adaptive stochastic submodular maximization requires $\Omega(\log n)$ batch queries.*

Proof Consider the following example. We have $n = 2^{k-1} - 1$ elements, and we want to select k elements. The elements are decomposed into k bags of sizes $1, 2, \ldots, 2^k$, where the decomposition is chosen uniformly at random. The objective function for a set S is the number of distinct bags that elements in S belong to. Whenever we select an element e we see all of the elements that are in the same bag as e. It is easy to see that this function is adaptive monotone and adaptive submodular.

Note that one can iteratively select k elements each with a marginal benefit of 1 and hence the value of the optimum solution of this instance is k. Next we upper-bound the value of the solution of a policy with $t \in o(\log n)$ batch-queries.

Let B_i be the i-th batch query and let $b_i = |B_i|$. Note that the marginal gain of each element is either 0 or 1. Moreover, all of the elements with the marginal gain of 1 are symmetric. Hence, without loss of generality, we assume that B_i is random subset of elements with the marginal gain of 1. Hence, with probability at least $(1 - \frac{1}{b_i})$ all of the elements in B_i belong to the $\log^2 b_i$ largest bags with the marginal gain of 1. Hence, the expected marginal benefit of batch B_i is at most $(1 - \frac{1}{b_i}) \log^2 b_i + \frac{1}{b_i} b_i \leq \log^2 b_i + 1$. Hence the expected value of the solution of this policy is at most

$$\sum_{i=1}^{t} (\log^2 b_i + 1) \leq \sum_{i=1}^{t} (\log^2 \frac{k}{t} + 1) = t \log^2 \frac{k}{t} + t \in o(k),$$

where the last inequality is due to $t \in o(\log n) = o(k)$. ■

Notice that in the hard example provided in the above theorem, we upper bound the marginal gain of each batch of size r by $\log^2 r + 1$. Hence if we force each batch to query exactly r elements, the expected value of the final solution is at most $\frac{k}{r}\big(\log^2 r + 1\big) \in O(\frac{k \log^2 r}{r})$.

Corollary 13 *Let π be a policy for adaptive stochastic submodular maximization that queries batches of size r. The approximation factor of π is upper bounded by $O(\frac{\log^2 r}{r})$.*

8. Conclusion

In this paper, we re-examined the required rounds of adaptive observations in order to maximize an adaptive submodular function. We proposed an efficient batch policy that with $O(\log n \times \log k)$ adaptive rounds of observations can achieve a $(1 - 1/e - \varepsilon)$ approximation guarantee with respect to an optimal policy that carries out k actions, from a set of n actions, in a fully sequential setting. We also extended our result to the case of adaptive stochastic minimum cost coverage and proposed a batch policy that provides the same guarantee in polylogarithmic adaptive rounds through a similar information-parallelism scheme. In the mean time, we also proved the conjecture by [24] that the greedy policy achieves the asymptotically tight logarithmic approximation guarantee for adaptive stochastic minimum cost coverage. One interesting future direction is to develop a semi adaptive policy for maximizing the value of information [13].

References

[1] Arpit Agarwal, Sepehr Assadi, and Sanjeev Khanna. Stochastic submodular cover with limited adaptivity. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, 2019.

[2] Arash Asadpour and Hamid Nazerzadeh. Maximizing stochastic monotone submodular functions. *Management Science*, 2016.

[3] Eric Balkanski and Yaron Singer. The adaptive complexity of maximizing a submodular function. In *ACM SIGACT Symposium on Theory of Computing, STOC*, pages 1138–1151, 2018.

[4] Eric Balkanski and Yaron Singer. Parallelization does not accelerate convex optimization: Adaptivity lower bounds for non-smooth convex minimization. *arXiv preprint arXiv:1808.03880*, 2018.

[5] Eric Balkanski, Adam Breuer, and Yaron Singer. Non-monotone Submodular Maximization in Exponentially Fewer Iterations. In *Advances in Neural Information Processing Systems*, pages 2359–2370, 2018.

[6] Eric Balkanski, Aviad Rubinstein, and Yaron Singer. An Exponential Speedup in Parallel Running Time for Submodular Maximization without Loss in Approximation. *CoRR*, abs/1804.06355, 2018.

[7] Sébastien Bubeck, Qijia Jiang, Yin-Tat Lee, Yuanzhi Li, and Aaron Sidford. Complexity of highly parallel non-smooth convex optimization. In *Advances in Neural Information Processing Systems*, 2019.

[8] L. Elisa Celis, Amit Deshpande, Tarun Kathuria, and Nisheeth K. Vishnoi. How to be Fair and Diverse? *CoRR*, abs/1610.07183, 2016.

[9] Lin Chen, Moran Feldman, and Amin Karbasi. Unconstrained submodular maximization with constant adaptive complexity. *CoRR*, abs/1811.06603, 2018.

[10] Yuxin Chen. *Near-optimal Adaptive Information Acquisition: Theory and Applications*. PhD thesis, ETH Zurich, 2017.

[11] Yuxin Chen and Andreas Krause. Near-optimal batch mode active learning and adaptive submodular optimization. *ICML*, 2013.

[12] Yuxin Chen, Hiroaki Shio, César Antonio Fuentes Montesinos, Lian Pin Koh, Serge Wich, and Andreas Krause. Active detection via adaptive submodularity. In *International Conference on International Conference on Machine Learning-Volume 32*, 2014.

[13] Yuxin Chen, S Hamed Hassani, Amin Karbasi, and Andreas Krause. Sequential information maximization: When is greedy near-optimal? In *Conference on Learning Theory*, 2015.

[14] Yuxin Chen, Shervin Javdani, Amin Karbasi, J Andrew Bagnell, Siddhartha S Srinivasa, and Andreas Krause. Submodular surrogates for value of information. In *AAAI*, 2015.

[15] Sanjoy Dasgupta and Daniel Hsu. Hierarchical sampling for active learning. In *Proceedings of the 25th international conference on Machine learning*, pages 208–215. ACM, 2008.

[16] Jelena Diakonikolas and Cristóbal Guzmán. Lower bounds for parallel and randomized convex optimization. *Journal of Machine Learning Research*, 2020.

[17] Ethan R. Elenberg, Alexandros G. Dimakis, Moran Feldman, and Amin Karbasi. Streaming Weak Submodularity: Interpreting Neural Networks on the Fly. In *Advances in Neural Information Processing Systems*, 2017.

[18] Alina Ene and Huy L. Nguyen. Submodular Maximization with Nearly-optimal Approximation and Adaptivity in Nearly-linear Time. *CoRR*, abs/1804.05379, 2018.

[19] Hossein Esfandiari, Amin Karbasi, Abbas Mehrabian, and Vahab Mirrokni. Batched multi-armed bandits with optimal regret. *arXiv preprint arXiv:1910.04959*, 2019.

[20] Matthew Fahrbach, Vahab S. Mirrokni, and Morteza Zadimoghaddam. Non-monotone Submodular Maximization with Nearly Optimal Adaptivity Complexity. *CoRR*, abs/1808.06932, 2018.

[21] Matthew Fahrbach, Vahab S. Mirrokni, and Morteza Zadimoghaddam. Submodular Maximization with Optimal Approximation, Adaptivity and Query Complexity. *CoRR*, abs/1807.07889, 2018.

[22] Zijun Gao, Yanjun Han, Zhimei Ren, and Zhengqing Zhou. Batched multi-armed bandits problem. In *Advances in Neural Information Processing Systems*, pages 503–513, 2019.

[23] Michel Goemans and Jan Vondrák. Stochastic covering and adaptivity. In *Latin American symposium on theoretical informatics*, 2006.

[24] Daniel Golovin and Andreas Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *Journal of Artificial Intelligence Research*, 42: 427–486, 2011.

[25] Daniel Golovin, Andreas Krause, and Debajyoti Ray. Near-optimal bayesian active learning with noisy observations. In *Advances in Neural Information Processing Systems*, pages 766–774, 2010.

[26] Alkis Gotovos, Amin Karbasi, and Andreas Krause. Non-monotone adaptive submodular maximization. In *IJCAI*, 2015.

[27] Andrew Guillory and Jeff Bilmes. Interactive Submodular Set Cover. In *International Conference on Machine Learning*, pages 415–422. Omnipress, 2010.

[28] Andrew Guillory and Jeff A. Bilmes. Active Semi-Supervised Learning using Submodular Functions. In *Uncertainty in Artificial Intelligence (UAI)*, pages 274–282, 2011.

[29] Anupam Gupta, Viswanath Nagarajan, and Sahil Singla. Adaptivity gaps for stochastic probing: Submodular and xos functions. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2017.

[30] Shervin Javdani, Yuxin Chen, Amin Karbasi, Andreas Krause, Drew Bagnell, and Siddhartha S Srinivasa. Near optimal bayesian active learning for decision making. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2014.

[31] Amin Karbasi, Stratis Ioannidis, and laurent Massoulie. Comparison-based learning with rank nets. In *Proceedings of the 29th International Conference on Machine Learning*, 2012.

[32] Amin Karbasi, Stratis Ioannidis, and Laurent Massoulié. Comparison-based learning with rank nets. *International Conference on Machine Learning*, 2012.

[33] Andreas Krause, Jure Leskovec, Carlos Guestrin, Jeanne VanBriesen, and Christos Faloutsos. Efficient sensor placement optimization for securing large water distribution networks. *Journal of Water Resources Planning and Management*, 134(6):516–526, November 2008. URL `files/krause08efficient.pdf`.

[34] Qi Lei, Lingfei Wu, Pin-Yu Chen, Alexandros G Dimakis, Inderjit S Dhillon, and Michael Witbrock. Discrete attacks and submodular optimization with applications to text classification. *arXiv preprint arXiv:1812.00151*, 2018.

[35] Hui Lin and Jeff Bilmes. A class of submodular functions for document summarization. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 510–520. Association for Computational Linguistics, 2011.

[36] Baharan Mirzasoleiman, Amin Karbasi, and Andreas Krause. Deletion-Robust Submodular Maximization: Data Summarization with "the Right to be Forgotten". In *International Conference on Machine Learning*, 2017.

[37] Marko Mitrovic, Ehsan Kazemi, Moran Feldman, Andreas Krause, and Amin Karbasi. Adaptive sequence submodularity. In *Advances in Neural Information Processing Systems*, 2019.

[38] Feng Nan and Venkatesh Saligrama. Comments on the proof of adaptive submodular function minimization. *arXiv preprint arXiv:1705.03771*, 2017.

[39] Christos H Papadimitriou and John N Tsitsiklis. The complexity of markov decision processes. *Mathematics of operations research*, 1987.

[40] Paola Sebastiani and Henry P Wynn. Bayesian experimental design and shannon information. In *Proceedings of the Section on Bayesian Statistical Science*, volume 44, pages 176–181. Citeseer, 1997.

[41] Yaron Singer. Influence maximization through adaptive seeding. *ACM SIGecom Exchanges*, 2016.

[42] Adish Singla, Ilija Bogunovic, Gábor Bartók, Amin Karbasi, and Andreas Krause. Near-optimally teaching the crowd to classify. In *International Conference on Machine Learning*, 2014.

[43] Ehsan Tohidi, Rouhollah Amiri, Mario Coutino, David Gesbert, Geert Leus, and Amin Karbasi. Submodularity in action: From machine learning to signal processing applications, 2020.

[44] Guangmo Tong, Weili Wu, Shaojie Tang, and Ding-Zhu Du. Adaptive influence maximization in dynamic social networks. *IEEE/ACM Transactions on Networking*, 2016.

[45] Kai Wei, Rishabh Iyer, and Jeff Bilmes. Submodularity in Data Subset Selection and Active Learning. In *International Conference on Machine Learning*, pages 1954–1963, 2015.

Appendix A. Omitted proofs

A.1. Proof of Corollary 3

Proof Let K be a random variable that indicates the number of items picked by π^*, i.e., $K = S(\pi^*, \Phi)$. Set $\ell = (\mathbb{E}[K]+1)\log(\frac{Q}{\delta\eta})$. Note that by definition of π^* we have $f(\pi^*) = Q$ for all ϕ, hence we have $f_{avg}(\pi^*) = Q$. By Lemma 1 we have

$$
\begin{aligned}
f_{avg}(\pi^{\tau_\ell}) &> \left(1 - e^{-\frac{\ell}{\mathbb{E}[K]+1}}\right) f_{avg}(\pi^*) && \text{By Lemma 1} \\
&= \left(1 - e^{-\log(\frac{Q}{\delta\eta})}\right) f_{avg}(\pi^*) && \text{Since } \ell = \big(\mathbb{E}[K]+1\big)\log(\frac{Q}{\delta\eta}) \\
&= \left(1 - \frac{\delta\eta}{Q}\right) f_{avg}(\pi^*) \\
&= Q - \delta\eta. && \text{Since } f_{avg}(\pi^*) = Q
\end{aligned}
$$

Recall that by definition we have $\mathbb{E}[f(\pi^{\tau_\ell})] = f_{avg}(\pi^{\tau_\ell}) = Q - \delta\eta$. Moreover, by adaptive monotonicity we have $f(\pi^{\tau_\ell}) \leq f(\phi) = Q$. Hence by Markov inequality with probability more than $1-\delta$ we have $f(\pi^{\tau_\ell}) > Q - \eta$. By definition of η this implies that with probability more than $1-\delta$ we have $f(\pi^{\tau_\ell}) = Q$. Equivalently, probability of $f(\pi^{\tau_\ell}) \neq Q$ is less than δ.

Note that the only source of randomness in π^{τ_ℓ} is from the randomness of the input. Hence, for any fixed ϕ we either have $f(\pi^{\tau_\ell}) = Q$ or $f(\pi^{\tau_\ell}) \neq Q$, deterministically. On the other hand, by definition $\delta = \min_\phi p(\phi)$. Hence, since the probability of $f(\pi^{\tau_\ell}) \neq Q$ is less than δ, the probability of $f(\pi^{\tau_\ell}) \neq Q$ must be 0. Therefore, the policy π^{τ_ℓ} reaches the utility Q, certainly, after selecting $\ell = \big(\mathbb{E}[K]+1\big)\log(nQ/\eta)$ items in expectation. ■

A.2. Proof of Lemma 6

Proof First we provide an upper bound on $f_{avg}(\pi^*)$. Pick an arbitrary $i \in \{0, \ldots, \ell\}$. Note that by adaptive monotonicity, we have $f_{avg}(\pi^*) \leq f_{avg}(\pi_{[i]}@\pi^*)$. Next we show that $f_{avg}(\pi_{[i]}@\pi^*) \leq f_{avg}(\pi_{[i]}) + \frac{k}{1-\varepsilon}\Delta_i(\pi)$, where $\Delta_i(\pi) = f_{avg}(\pi_{[i+1]}) - f_{avg}(\pi_{[i]})$. This implies that

$$
f_{avg}(\pi^*) \leq f_{avg}(\pi_{[i]}@\pi^*) \leq f_{avg}(\pi_{[i]}) + \frac{k}{1-\varepsilon}\Delta_i(\pi). \tag{3}
$$

Let Ψ_i be a random variable that indicates the partial realization of the elements selected by $\pi_{[i]}$. We use ψ_i to indicate a realization of Ψ_i. Let $\bar{\Psi}$ be a random variable that indicates the last partial realization that is queried by $\pi_{[i]}$ (ignoring the selected elements in the last batch that is not queried yet).

We have

$$
\begin{aligned}
& f_{avg}(\pi_{[i]}@\pi^*) - f_{avg}(\pi_{[i]}) \\
&= \sum_{\psi_i \in \Psi_i} \sum_{\psi_i \preccurlyeq \psi'} \sum_{e \notin \mathsf{dom}(\psi')} \sum_{\theta \in \Theta_{\pi^*}} p(\psi')p(\theta)\mathbb{1}_{\psi_i,\psi',e,\theta}\mathbb{E}\left[f(\psi' \cup \Phi(e)) - f(\psi')|\psi' \preccurlyeq \Phi \text{ and } \Theta = \theta\right] \\
& \qquad \text{By definition} \\
&= \sum_{\psi_i \in \Psi_i} \sum_{\psi_i \preccurlyeq \psi'} \sum_{e \notin \mathsf{dom}(\psi')} \sum_{\theta \in \Theta_{\pi^*}} p(\psi')p(\theta)\mathbb{1}_{\psi_i,\psi',e,\theta}\mathbb{E}\left[f(\psi' \cup \Phi(e)) - f(\psi')|\psi' \preccurlyeq \Phi\right] \\
& \qquad \text{By independency} \\
&= \sum_{\psi_i \in \Psi_i} \sum_{\psi_i \preccurlyeq \psi'} \sum_{e \notin \mathsf{dom}(\psi')} \sum_{\theta \in \Theta_{\pi^*}} p(\psi')p(\theta)\mathbb{1}_{\psi_i,\psi',e,\theta}\Delta(e|\psi') \\
& \qquad \text{By definition of } \Delta(e|\psi') \\
&\leq \sum_{\psi_i \in \Psi_i} \sum_{\psi_i \preccurlyeq \psi'} \sum_{e \notin \mathsf{dom}(\psi')} \sum_{\theta \in \Theta_{\pi^*}} p(\psi')p(\theta)\mathbb{1}_{\psi_i,\psi',e,\theta}\Delta(e|\psi_i) \\
& \qquad \text{Bu adaptive Submodularity.} \\
&\leq \sum_{\psi_i \in \Psi_i} \sum_{\psi_i \preccurlyeq \psi'} \sum_{e \notin \mathsf{dom}(\psi')} \sum_{\theta \in \Theta_{\pi^*}} p(\psi')p(\theta)\mathbb{1}_{\psi_i,\psi',e,\theta} \max_{e' \notin \mathsf{dom}(\psi_i)} \Delta(e'|\psi_i) \\
&= \sum_{\psi_i \in \Psi_i} \left(\sum_{\psi_i \preccurlyeq \psi'} \sum_{e \notin \mathsf{dom}(\psi')} \sum_{\theta \in \Theta_{\pi^*}} p(\psi')p(\theta)\mathbb{1}_{\psi_i,\psi',e,\theta} \right) \max_{e' \notin \mathsf{dom}(\psi_i)} \Delta(e'|\psi_i) \\
&\leq \sum_{\psi_i \in \Psi_i} p(\psi_i) k \max_{e' \notin \mathsf{dom}(\psi_i)} \Delta(e'|\psi_i) \\
&= k \sum_{\psi_i \in \Psi_i} p(\psi_i) \max_{e' \notin \mathsf{dom}(\psi_i)} \Delta(e'|\psi_i) \\
&= k\mathbb{E}_{\Psi_i}\left[\max_{e' \notin \mathsf{dom}(\Psi_i)} \Delta(e'|\Psi_i) \right] \\
&= k\mathbb{E}_{\bar{\Psi}}\left[\mathbb{E}_{\bar{\Psi} \preccurlyeq \Psi_i}\left[\max_{e' \notin \mathsf{dom}(\Psi_i)} \Delta(e'|\Psi_i) \right] \right] \\
&\leq k\mathbb{E}_{\bar{\Psi}}\left[\frac{1}{1-\varepsilon} \max_{e' \notin \mathsf{dom}(\Psi_i^\pi)} \mathbb{E}_{\bar{\Psi} \preccurlyeq \Psi_i^\pi}\left[\Delta(e'|\Psi_i^\pi)\right] \right] \\
& \qquad \text{By information gap bound} \\
&= \frac{k}{1-\varepsilon}\mathbb{E}_{\bar{\Psi}}\left[\max_{e' \notin \mathsf{dom}(\Psi_i^\pi)} \mathbb{E}_{\bar{\Psi} \preccurlyeq \Psi_i^\pi}\left[\Delta(e'|\Psi_i^\pi)\right] \right] \\
&= \frac{k}{1-\varepsilon}\Delta_i(\pi)
\end{aligned}
$$

This proves Inequality 3 as promised. Let us define

$$
\Delta_i^* = f_{avg}(\pi^*) - f_{avg}(\pi_{[i]}).
$$

Inequality 3 implies that

$$\Delta_i^* \leq \frac{k}{1-\varepsilon}(\Delta_i^* - \Delta_{i+1}^*).$$

By a simple rearrangement we have

$$\Delta_{i+1}^* \leq \big(1 - \frac{1-\varepsilon}{k}\big)\Delta_i^*.$$

By iteratively applying this inequality we have

$$\Delta_\ell^* \leq \big(1 - \frac{1-\varepsilon}{k}\big)^\ell \Delta_0^* \leq e^{-\frac{(1-\varepsilon)\ell}{k}}\Delta_0^*.$$

By applying the definition of Δ_i^* and some rearrangements we have

$$f_{avg}(\pi_{[\ell]}) > \big(1 - e^{-\frac{(1-\varepsilon)\ell}{k}}\big) f_{avg}(\pi^*) \geq \big(1 - e^{-\frac{\ell}{k}} - \varepsilon\big) f_{avg}(\pi^*)$$

as desired.

■

A.3. Proof of Lemma 14

Let us start with some definitions. Let π be the semi-adaptive greedy policy as defined in section 5, using restricted information gap (Definition 9). For an arbitrary number τ let π^τ be a policy that selects elements according to π and stops when the semi-adaptive value of all of the remaining elements is less than or equal to τ. We define τ_i to be a number such that the expected number of elements selected by π^{τ_i} is i.

Now we are ready to prove Lemma 14.

Lemma 14 *For any policy π^* and any positive integer ℓ we have*

$$f_{avg}(\pi^{\tau_\ell}) > \big(1 - e^{-\frac{(1-\varepsilon)\ell}{\mathbb{E}[K]+1}}\big) f_{avg}(\pi^*),$$

where K is a random variable that indicates the number of items picked by π^*, i.e. $K = |S(\pi^*, \Phi)|$.

Proof Let us define $\Delta^{\tau_i}(\pi) = f_{avg}(\pi^{\tau_i}) - f_{avg}(\pi^{\tau_{i-1}})$. Recall that in expectation π^{τ_i} picks one item more than $\pi^{\tau_{i-1}}$. Moreover note that by definition of π^{τ_i} the semi-adaptive value of all of the items selected by π^{τ_i} is at least τ_i. Hence

$$f_{avg}(\pi^{\tau_i}) - f_{avg}(\pi^{\tau_{i-1}}) = \Delta^{\tau_i}(\pi) \geq \tau_i. \tag{4}$$

Let Ψ_i be a random variable that indicates the partial realization of the elements selected by π^{τ_i}. We use ψ_i to indicate a realization of Ψ_i. Next we show that for any consistent partial realization $\psi_i \preccurlyeq \psi'$ and any element e we have

$$\Delta(e|\psi') \leq \frac{\tau_i}{1-\varepsilon}. \tag{5}$$

We have two cases based on the time that the policy π^{τ_i} stops.

- The policy π^{τ_i} queries a batch and then observe that the semi-adaptive value of all items drop below τ_i and then π^{τ_i} stops.
- While adding items to a batch (and before performing the query), the semi-adaptive value of all items drop below τ_i and then π^{τ_i} stops.

Note that in the first case the semi-adaptive values of all of the items are equal to their actual expected marginal benefit (i.e., $\Delta(e|\psi')$). Hence, we have $\Delta(e|\psi') \leq \Delta(e|\psi_i) \leq \tau_i$. In the second case, by the definition of the algorithm, the restricted information gap is at least $1-\epsilon$. This together with the fact that the semi-adaptive values of all items are below τ_i implies that the expected marginal benefit of the first item that was added to the last batch is at most $\frac{\tau_i}{1-\epsilon}$. This together with adaptive monotonicity implies $\Delta(e|\psi') \leq \frac{\tau_i}{1-\varepsilon}$ as desired.

We have

$$
\begin{aligned}
&f_{avg}(\pi^{\tau_i}@\pi^*) - f_{avg}(\pi^{\tau_i}) \\
&= \sum_{\psi_i\in\Psi_i}\sum_{\psi_i\preccurlyeq\psi'}\sum_{e\notin\mathsf{dom}(\psi')}\sum_{\theta\in\Theta_{\pi^*}} p(\psi')p(\theta)\mathbb{1}_{\psi_i,\psi',e,\theta}\mathbb{E}\left[f(\psi'\cup\Phi(e)) - f(\psi')|\psi'\preccurlyeq\Phi \text{ and } \Theta=\theta\right] \\
&\qquad\qquad\qquad\qquad\qquad\qquad \text{By definition} \\
&= \sum_{\psi_i\in\Psi_i}\sum_{\psi_i\preccurlyeq\psi'}\sum_{e\notin\mathsf{dom}(\psi')}\sum_{\theta\in\Theta_{\pi^*}} p(\psi')p(\theta)\mathbb{1}_{\psi_i,\psi',e,\theta}\mathbb{E}\left[f(\psi'\cup\Phi(e)) - f(\psi')|\psi'\preccurlyeq\Phi\right] \\
&\qquad\qquad\qquad\qquad\qquad\qquad \text{By independency} \\
&= \sum_{\psi_i\in\Psi_i}\sum_{\psi_i\preccurlyeq\psi'}\sum_{e\notin\mathsf{dom}(\psi')}\sum_{\theta\in\Theta_{\pi^*}} p(\psi')p(\theta)\mathbb{1}_{\psi_i,\psi',e,\theta}\Delta(e|\psi') \\
&\qquad\qquad\qquad\qquad\qquad\qquad \text{By definition of } \Delta(e|\psi') \\
&\leq \sum_{\psi_i\in\Psi_i}\sum_{\psi_i\preccurlyeq\psi'}\sum_{e\notin\mathsf{dom}(\psi')}\sum_{\theta\in\Theta_{\pi^*}} p(\psi')p(\theta)\mathbb{1}_{\psi_i,\psi',e,\theta}\frac{\tau_i}{1-\varepsilon} \\
&\qquad\qquad\qquad\qquad\qquad\qquad \text{By Inequality 5} \\
&\leq \sum_{\psi_i\in\Psi_i}\sum_{\psi_i\preccurlyeq\psi'}\sum_{e\notin\mathsf{dom}(\psi')}\sum_{\theta\in\Theta_{\pi^*}} p(\psi')p(\theta)\mathbb{1}_{\psi_i,\psi',e,\theta}\frac{\Delta^{\tau_i}(\pi)}{1-\varepsilon} \\
&\qquad\qquad\qquad\qquad\qquad\qquad \text{By Inequality 4} \\
&= \Big(\sum_{\psi_i\in\Psi_i}\sum_{\psi_i\preccurlyeq\psi'}\sum_{e\notin\mathsf{dom}(\psi')}\sum_{\theta\in\Theta_{\pi^*}} p(\psi')p(\theta)\mathbb{1}_{\psi_i,\psi',e,\theta}\Big)\frac{\Delta^{\tau_i}(\pi)}{1-\varepsilon} \\
&= \frac{\mathbb{E}\left[K\right]}{1-\varepsilon}\Delta^{\tau_i}(\pi).
\end{aligned}
$$

Now define $\Delta^*_i = f_{avg}(\pi^*) - f_{avg}(\pi^{\tau_i})$. The above inequality implies that

$$\Delta^*_i \leq \frac{\mathbb{E}\left[K\right]}{1-\varepsilon}(\Delta^*_{i-1} - \Delta^*_i).$$

By a simple rearrangement we have

$$\Delta^*_i \leq \big(1 - \frac{1}{\frac{\mathbb{E}[K]}{1-\varepsilon}+1}\big)\Delta^*_{i-1} \leq \big(1 - \frac{1-\varepsilon}{\mathbb{E}\left[K\right]+1}\big)\Delta^*_{i-1}.$$

By iteratively applying this inequality we have

$$\Delta_\ell^* \leq \big(1 - \frac{1-\varepsilon}{\mathbb{E}\left[K\right]+1}\big)^\ell \Delta_0^* \leq e^{-\frac{(1-\varepsilon)\ell}{\mathbb{E}[K]+1}} \Delta_0^*.$$

By applying the definition of Δ_i^* and some rearrangement we have

$$f_{avg}(\pi^{\tau_\ell}) > \big(1 - e^{-\frac{(1-\varepsilon)\ell}{\mathbb{E}[K]+1}}\big) f_{avg}(\pi^*)$$

as desired. ■

A.4. Proof of Theorem 10

Proof Let K be a random variable that indicates the number of items picked by π^*. Set

$$\ell = \frac{\mathbb{E}\left[K\right]+1}{1-\varepsilon}\log(nQ/\eta).$$

Note that by definition of π^* we have $f(\pi^*) = Q$ for all ϕ, hence we have $f_{avg}(\pi^*) = Q$. By Lemma 14 we have

$$\begin{aligned} f_{avg}(\pi^{\tau_\ell}) &> \left(1 - e^{-\frac{(1-\varepsilon)\ell}{\mathbb{E}[K]+1}}\right) f_{avg}(\pi^*) && \text{By Lemma 14} \\ &= \left(1 - e^{-log(nQ/\eta)}\right) f_{avg}(\pi^*) && \text{Since } \frac{\mathbb{E}\left[K\right]+1}{1-\varepsilon}\log(nQ/\eta) \\ &= (1 - \frac{\eta}{nQ}) f_{avg}(\pi^*) \\ &= Q - \frac{\eta}{n}. && \text{Since } f_{avg}(\pi^*) = Q \end{aligned}$$

Recall that, by definition $f_{avg}(\pi^{\tau_\ell}) = \mathbb{E}\left[f(\pi^{\tau_\ell})\right]$. Moreover, note that by adaptive monotonicity we have $f(\pi^{\tau_\ell}) \leq f(\phi) = Q$. Hence by Markov inequality with probability $1 - 1/n$ we have $f(\pi^{\tau_\ell}) > Q - \eta$. By definition of η this implies that with probability $1 - \frac{1}{n}$ we have $f(\pi^{\tau_\ell}) = Q$. Therefore, with probability $1 - 1/n$, π^{τ_ℓ} reaches Q after selecting $\ell = \big(\mathbb{E}\left[K\right]+1\big)\log(nQ/\eta)$ items in expectations. Otherwise, π picks at most all n items. Hence the expected number of items that π picks is upper bounded by

$$(1 - \frac{1}{n}) \times \frac{\mathbb{E}\left[K\right]+1}{1-\varepsilon}\log(\frac{nQ}{\eta}) + \frac{1}{n} \times n \leq (\frac{c_{avg}(\pi^*)+1}{1-\varepsilon}) \log\left(\frac{nQ}{\eta}\right) + 1.$$

■

A.5. Proof of Theorem 11

First let us start with a couple of definitions. We define random variable $\bar{\Psi}_t$ to be the partial realization obtained by the t-th query. Next we prove Theorem 11.

Proof of Theorem 11: In order to prove this theorem we show that $f_{avg}(\pi^T) > Q - \frac{\eta}{n}$. Note that $f(\pi^T) \leq Q$. This together with a Markov bound imply $f(\pi^T) > Q - \eta$, and hence $f(\pi^T) = Q$, with probability $1 - 1/n$. In this theorem we simply set $\varepsilon = 0.01$.

Let us set

$$\delta = \frac{\frac{\eta}{2Qn}}{\log_{\frac{1}{1-\varepsilon/2}} \frac{2Qn^2}{\eta}}.$$

By applying Lemma 7, $\log_{\frac{1}{1-\varepsilon/2}} \frac{2Qn^2}{\eta}$ times iteratively we have

$$\max_{e \notin \mathsf{dom}(\bar{\Psi}_T)} \Delta(e|\bar{\Psi}_T) \leq \big(1 - \frac{\varepsilon}{2}\big)^{\log_{\frac{1}{1-\varepsilon/2}} \frac{2Qn^2}{\eta}} \Delta_1(\pi) = \frac{\eta}{2Qn^2}\Delta_1(\pi),$$

with probability $1 - \delta \times \log_{\frac{1}{1-\varepsilon/2}} \frac{2Qn^2}{\eta} = 1 - \frac{\eta}{2Qn}$. This means that with probability $1 - \frac{\eta}{2Qn}$ the total expected marginal gain of the elements added after the T-th query is at most

$$n \times \frac{\eta}{2Qn^2}\Delta_1(\pi) = \frac{\eta}{2Qn}\Delta_1(\pi) \leq \frac{\eta}{2n},$$

where

$$\begin{aligned} T &= \log_{\frac{1}{1-\varepsilon/2}} \big(\frac{n}{\delta}\big) \times \log_{\frac{1}{1-\varepsilon/2}} \frac{2Qn^2}{\eta} \\ &\in O\big((\log n + \log\log(Qn/\eta)) \log(Qn/\eta)\big) \\ &\in O\big(\log n \log(Qn/\eta)\big).^{9} \end{aligned}$$

This implies that

$$f_{avg}(\pi^T) \geq (1 - \frac{\eta}{2Qn})(Q - \frac{\eta}{2n}) > Q - \frac{\eta}{n}.$$

This implies that $f(\pi^T) = Q$ with probability at least $1 - \frac{1}{n}$, as desired. □

Appendix B. Truncation

Consider the following simple adaptive submodular function $f(\cdot)$. We have three elements $\{x, y, z\}$ each of x and y are associated with an independent uniform random binary variable. The value of the empty set is zero. If element z exists in a set, it deterministically adds a value 1 to the set. If there is only one of x and y in the set, it adds a value 1 to the set. However, if both x and y are in the set, if their corresponding random variables match, they add a value 2 to the set, and otherwise add nothing.

Note that if one of x and y exists in a set, adding the other one does not change the value of the set in expectation. In all other cases the value of adding an element is 1. This implies that this function is adaptive submodular. However, $g(\psi) = \min\big(f(\psi), 1\big)$ is not adaptive submodular. For example the marginal gain of z on $g\big(\{(x, 1)\}\big)$ is 0 but the marginal gain of z on $g\big(\{(x, 1), (y, 0)\}\big)$ is 1.

9. We can assume $\log n > \log\log(Qn/\eta)$, since otherwise $n \leq \log(Qn/\eta)$ and hence trivially $T \in O(\log(Qn/\eta))$ as desired.