

Adaptive Combinatorial Maximization: Beyond Approximate Greedy Policies

Shlomi Weitzman

Department of Computer Science, Ben-Gurion University of the Negev

SHLOMIWEITZMAN@GMAIL.COM

Sivan Sabato

Department of Computing and Software, McMaster University

Canada CIFAR AI Chair, Vector Institute

Department of Computer Science, Ben-Gurion University of the Negev

SABATOS@MCMASTER.CA

Editors: Claire Vernade and Daniel Hsu

Abstract

We study adaptive combinatorial maximization, which is a core challenge in machine learning, with applications in active learning as well as many other domains. We study the Bayesian setting, and consider the objectives of maximization under a cardinality constraint and minimum cost coverage. We provide new comprehensive approximation guarantees that subsume previous results, as well as considerably strengthen them. Our approximation guarantees simultaneously support the maximal gain ratio as well as near-submodular utility functions, and include both maximization under a cardinality constraint and a minimum cost coverage guarantee. In addition, we provided an approximation guarantee for a modified prior, which is crucial for obtaining active learning guarantees that do not depend on the smallest probability in the prior. Moreover, we discover a new parameter of adaptive selection policies, which we term the *maximal gain ratio*. We show that this parameter is strictly less restrictive than the greedy approximation parameter that has been used in previous approximation guarantees, and show that it can be used to provide stronger approximation guarantees than previous results. In particular, we show that the maximal gain ratio is never larger than the greedy approximation factor of a policy, and that it can be considerably smaller. This provides a new insight into the properties that make a policy useful for adaptive combinatorial maximization.

Keywords: Combinatorial maximization, adaptive submodularity, approximate greedy policies

1. Introduction

Adaptive combinatorial maximization is a core challenge in machine learning, spanning a wide range of applications, including the fundamental challenges of active learning (Golovin and Krause, 2011) and adaptive experiment design (Doppa, 2021), as well as many other applications, such as influence maximization in social networks (Seeman and Singer, 2013; Tong et al., 2016), movie recommendations (Mitrovic et al., 2019) and spectroscopy (Hino, 2020).

In this algorithmic problem, there are elements with hidden states. Elements are selected sequentially. Whenever an element is selected, its state is observed. Past observations can be used to make the next selection decision. The obtained utility from a specific selection process depends on the selected elements and their true states. For instance, consider a problem of selecting locations for placing radio towers (Asadpour et al., 2008; Golovin and Krause, 2011), where the goal is to ensure sufficient signal strength in the designated area. The states in this case are parameters of the transmission that can only be measured after a radio tower is placed. The next radio tower location can be selected based on the previous measurements.

A particularly important application of adaptive combinatorial application in machine learning is active learning (McCallum et al., 1998), in which the elements are examples, and the state of an element is its true label. The active learning algorithm selects which element to have labeled, so as to reveal its true label, and the objective is to identify the true classifier for the entire set of examples using a small number of labels.

We study the Bayesian setting, in which there is a known prior over the possible states of the elements, and the performance of the algorithm is measured with respect to the expectation over this prior. The goal of the algorithm is either to obtain a high expected utility using a limited number of element selections, also termed “maximization under a cardinality constraint” or alternatively, to obtain the maximal possible utility value using a small expected number of element selections. The latter is typically referred to as “minimum cost coverage”. This objective is of particular importance for active learning, since in this application, the goal is to reach a specific target of finding the correct classifier, which can be mapped to a maximal value of the utility function. Golovin and Krause (2011) were the first to introduce the adaptive Bayesian formulation.

Calculating an optimal selection policy for either of the objectives mentioned above is computationally hard in the general case (see, e.g. Nemhauser et al., 1978; Wolsey, 1982). Thus, we seek instead to provide approximation guarantees with respect to an optimal policy. Golovin and Krause (2011) defined a class of utility functions called *adaptive submodular*, which generalizes the non-adaptive notion of submodular set functions (Edmonds, 1970). They showed that if the functions are also *adaptive monotone*, then an *approximately greedy* policy obtains a bounded approximation factor for the objective of maximization under a cardinality constraint. A greedy policy selects, in each round, the element that will obtain the largest increase in utility in expectation. An approximate greedy policy may select an element that is maximal up to a given factor. Golovin and Krause (2017) provided an approximation guarantee for the minimum cost coverage objective. A smaller approximation factor was proved by Esfandiari et al. (2021).

The guarantees mentioned above require that the utility function is adaptive submodular. Fujii and Sakaue (2019) studied a relaxed version, using the *adaptive submodularity ratio*, a property that quantifies the closeness of a utility function to being submodular. They provided an approximation guarantee for maximization under a cardinality constraint for greedy policies, that is parameterized by the adaptive submodularity ratio of the utility function.

In this work, we provide a new perspective on the problem of adaptive combinatorial maximization, by providing comprehensive approximation guarantees that not only subsume previous results, but also considerably strengthen them. We discover a new parameter of adaptive selection policies, which we term the *maximal gain ratio*. We show that this parameter is strictly less restrictive than the greedy approximation parameter that has been used in previous approximation guarantees, and show that it can be used to provide stronger approximation guarantees than previous results. In particular, we show that the maximal gain ratio is never larger than the greedy approximation factor of a policy, and that it can be considerably smaller.

The maximal gain ratio is the maximal ratio between the expected marginal gain of remaining elements at termination and the expected marginal gain of selected elements. In this way, it can take into account also selections that would seem out of order for a greedy policy, but that accomplish a similar goal. Moreover, we show that even for greedy policies, the value of the maximal gain ratio can be arbitrarily small. Thus, this provides a more refined understanding of what makes an adaptive selection policy successful. Our contributions are the following:

- Defining the maximal gain ratio, a new policy parameter;

- Showing that this parameter is strictly less restrictive than the greedy approximation parameter;
- An approximation guarantee for utility maximization under a cardinality constraint that is parameterized by the maximal gain ratio of a policy as well as the adaptive submodularity ratio of the utility function;
- An approximation guarantee for the minimum cost coverage objective parameterized by the same properties;
- A version of the minimum cost coverage guarantee that allows using the problem parameters with a modified prior first suggested by [Golovin and Krause \(2011\)](#), leading to an improved approximation ratio for Bayesian active learning under priors with small probabilities.

We summarize the main difference between our work and previous works mentioned above in [Table 1](#). Our work is the first to support, at the same time, policies that are not greedy, utility functions that are not adaptive submodular, and the minimum cost coverage objective, as well as guarantees for the modified prior.

	Approximate greedy policies	Nearly adaptive submodular utility	Min cost coverage	Support for modified prior	Beyond approximate greedy
GK11	✓	✗	(✓)	(✓)	✗
FS19	✗	✓	✗	✗	✗
EKM21	✗	✗	✓	✗	✗
This Work	✓	✓	✓	✓	✓

Table 1: Comparison to previous works. GK11: [Golovin and Krause \(2011\)](#), FS19: [Fujii and Sakaue \(2019\)](#), EKM21: [Esfandiari et al. \(2021\)](#). (✓) indicates a looser guarantee.

Paper structure. In [Section 2](#), we complement the discussion above on related work. Definitions and necessary background are provided in [Section 3](#). We provide a summary of our main results in [Section 4](#). The maximal gain ratio is presented and analyzed in [Section 5](#). The proofs of the approximation guarantees for the true prior based on the maximal gain ratio are provided in [Appendix A](#). The modified prior and its relevance to active learning are discussed in [Section 6](#), where the proof of our guarantee for this prior is also provided, with some parts deferred to [Appendix B](#). We summarize in [Section 7](#).

2. Related Work

The concept of adaptive submodular functions was first introduced in [Golovin and Krause \(2011\)](#). They showed that this class of functions allows deriving guarantees for several settings, including adaptive submodular maximization with respect to a Bayesian prior. The latter setting was further studied in [Fujii and Sakaue \(2019\)](#) and [Esfandiari et al. \(2021\)](#), who provided additional guarantees, as discussed above. [Fujii and Sakaue \(2019\)](#) introduced the adaptive submodularity ratio, which generalized the submodularity ratio proposed by [Das and Kempe \(2011\)](#) for non-adaptive set functions to the adaptive setting, and provided guarantees parameterized by this quantity. The notion of

curvature, first introduced in Conforti and Cornuéjols (1984) for set functions in the non-adaptive combinatorial maximization setting, allows further improving the guarantees and algorithms for submodular maximization under a cardinality constraint (Vondrak, 2010; Balkanski and Singer, 2018).

Adaptive submodular maximization has been studied in other settings as well, such as a worst-case non-Bayesian setting with pointwise submodularity (Guillory and Bilmes, 2010) with applications to active learning (Cuong et al., 2014), outcome dependent costs (Sabato, 2018), adaptive maximization of non-monotone submodular functions (Gotovos et al., 2015; Amanatidis et al., 2020; Tang, 2021), streaming adaptive maximization (Tang and Yuan, 2023), and robust adaptive submodular maximization Tang (2022). It has been applied to many settings, including batch mode active learning (Chen and Krause, 2013), bandit algorithms (Gabillon et al., 2013), and reinforcement learning (Wu and Tseng, 2022).

3. Preliminaries

In this section, we give necessary preliminaries. Definitions and notation are provided in Section 3.1. Previously studied properties of utility functions and policies are provided in Section 3.2. Guarantees from previous works are given in Section 3.3.

3.1. Definitions and Notation

Let V be a finite ground set of elements. Let Y be a finite set of the possible states of each of the elements in V . A *realization* is a map $\phi : V \rightarrow Y$ that associates each ground element with a state. A *partial realization* $\psi : V \rightarrow Y$ is a map from a subset of V to Y . Its domain is denoted by $\text{dom}(\psi)$. We denote the set of all realizations by Φ_V and the set of all partial realizations by Ψ_V . For any realization $\phi \in \Phi_V$ and partial realization $\psi \in \Psi_V$, we write $\phi \sim \psi$ to indicate that for every $x \in \text{dom}(\psi)$, $\psi(x) = \phi(x)$. For any two partial realizations ψ and ψ' , we write $\psi \subseteq \psi'$ if $\text{dom}(\psi) \subseteq \text{dom}(\psi')$ and $\forall x \in \text{dom}(\psi)$, $\psi(x) = \psi'(x)$.

We assume that a true unknown realization $\phi \in \Phi_V$ is drawn according to a known prior over Φ_V . This prior is denoted by p unless explicitly stated otherwise. An adaptive algorithm interactively selects elements and observes their states: At time t , an element $x_t \in V$ is selected and its state $\phi(x_t)$ is observed. The previous observations at time t are represented by the partial realization $\psi_t \subseteq \phi$ whose domain is $\{x_1, \dots, x_t\}$. The utility of selecting a specific set of elements under a given true realization is measured using a *utility function* $f : 2^V \times \Phi_V \rightarrow \mathbb{R}$, which maps each pair of a set of elements and a realization to a real value representing the utility when this set is selected under the given realization. The goal is to obtain a high utility while selecting few elements.

A *policy* $\pi : \Psi_V \rightarrow V \cup \{\perp\}$ is a (possibly non-deterministic) map from partial realizations to elements that determines which element to select next, given the past observations represented by the input partial realization. $\pi(\psi) = \perp$ for a given $\psi \in \Psi_V$ indicates that π terminates after observing ψ . For a policy π and a realization $\phi \in \Phi_V$, we denote by $E(\pi, \phi) \subseteq V$ the (possibly random) set of elements that the policy π selects through the entire run of the policy under realization ϕ until termination. The *height* of a policy is the maximal number of elements it may select, for any true realization and any random bits. For an integer k , Π_k denotes the set of all policies with height at most k . For any policy π , a *sub-policy* of π is any policy that runs exactly like π , except that it possibly terminates earlier in some cases.

The expected utility of a policy under a given prior p for a given utility function f is given by

$$f_{\text{avg}}(\pi, p) := \mathbb{E}[f(E(\pi, \phi), \phi)].$$

Here and below, the expectation is over the true realization $\phi \sim p$ as well as the randomness of the policy. The *expected cost* of using a given policy is defined as the expected number of elements that it selects until termination, given by

$$c_{\text{avg}}(\pi, p) := \mathbb{E}[|E(\pi, \phi)|].$$

The *expected marginal gain* of an element $v \in V$ under a given prior p and a utility function f with respect to a partial realization $\psi \in \Psi_V$ is the expected contribution to the utility value if v is selected next, assuming that the elements selected so far and their observed states are as specified by $\psi \in \Psi_V$. Formally, the expected marginal gain is given by

$$\Delta_p^f(v | \psi) := \mathbb{E}[f(\{v\} \cup \text{dom}(\psi), \phi) - f(\text{dom}(\psi), \phi) | \phi \sim \psi].$$

The expected marginal gain of a policy π under p and f , assuming the previous observations in ψ is the expected contribution to the utility until termination, is denoted by

$$\Delta_p^f(\pi | \psi) := \mathbb{E}[f(E(\pi, \phi) \cup \text{dom}(\psi), \phi) - f(\text{dom}(\psi), \phi) | \phi \sim \psi].$$

3.2. Properties of utility functions and policies

A utility function f is *adaptive monotone* if selecting an element is never harmful in expectation. Formally, f is adaptive monotone if for any partial realization ψ and any $v \in V \setminus \text{dom}(\psi)$, $\Delta_p^f(v | \psi) \geq 0$. A utility function f is *adaptive submodular* if selecting an element satisfies an expected diminishing-returns property, that is, if selecting the same element later never contributes more in expectation to the utility than selecting the same element earlier in the same run. Formally, a utility function is adaptive submodular if for any two partial realizations ψ, ψ' such that $\psi \subseteq \psi'$ and for any $v \in V \setminus \text{dom}(\psi')$,

$$\Delta_p^f(v | \psi) \geq \Delta_p^f(v | \psi').$$

The *adaptive submodularity ratio* of a utility function quantifies how close it is to being submodular. For a given prior p and two integers n, k , the adaptive submodularity ratio of f is defined as

$$\gamma_{n,k}^s(f, p) := \min_{\psi \in \Psi_V: |\psi| \leq n, \psi' \subseteq \psi, \pi \in \Pi_k} \frac{\sum_{v \in V} \mathbb{P}_{\Phi \sim p}[v \in E(\pi, \Phi) | \Phi \sim \psi'] \Delta_p^f(v | \psi')}{\Delta_p^f(\pi | \psi')}.$$

The adaptive submodularity ratio $\gamma_{n,k}^s(f, p)$ takes values in $[0, 1]$, where a higher value indicates that a function is closer to being adaptive submodular, and a value of 1 indicates that the function is adaptive submodular.

Previous works have studied *greedy* or *approximate greedy* policies, which, in each round, select an element that maximizes or approximately maximizes the expected marginal gain. For $\alpha \geq 1$, a policy is considered α -approximate greedy if for any $\psi \in \Psi_V$ such that $\pi(\psi) \neq \perp$,

$$\Delta_p^f(\pi(\psi) | \psi) \geq \frac{1}{\alpha} \max_{v \in V} \Delta_p^f(v | \psi).$$

A greedy policy is a policy which is 1-approximate greedy. Denote by $\alpha_\pi(f, p)$ the smallest real value α such that π is an α -approximate greedy policy with respect to f and p . We term this the *greedy approximation ratio* of the policy.

3.3. Previous approximation guarantees

[Golovin and Krause \(2011\)](#) proved an approximation guarantee for the expected utility of an approximate greedy policy π of height at most l , assuming that f is an adaptive submodular utility function. They proved that for any policy π^* with height k ,

$$f_{\text{avg}}(\pi, p) > \left(1 - \exp\left(-\frac{l}{\alpha_\pi(f, p) \cdot k}\right)\right) f_{\text{avg}}(\pi^*, p). \quad (1)$$

[Fujii and Sakaue \(2019\)](#) considered the same setting, but allowed also utility functions that are not adaptive submodular, using the adaptive submodularity ratio of the utility function. They required the policy to be exactly greedy, thus not supporting α -approximate greedy policies for $\alpha > 1$. They proved that for any policy π^* with height k ,

$$f_{\text{avg}}(\pi, p) \geq \left(1 - \exp\left(-\frac{\gamma_{l,k}^s(f, p) \cdot l}{k}\right)\right) \cdot f_{\text{avg}}(\pi^*, p). \quad (2)$$

This coincides with Eq. (1) for adaptive submodular functions ($\gamma_{l,k}^s(f, p) = 1$) and greedy policies ($\alpha = 1$).

[Esfandiari et al. \(2021\)](#) provided a guarantee that depends on the average costs of the policies. This contrasts with the results discussed above, which depend on the height of the two policies being compared. They considered only greedy policies and adaptive submodular functions. They proved that for a greedy policy π , if it terminates when the expected marginal gain of all remaining elements is below some fixed threshold, and if its average cost is l , then for any policy π^* ,

$$f_{\text{avg}}(\pi, p) > \left(1 - \exp\left(-\frac{l}{c_{\text{avg}}(\pi^*, p) + 1}\right)\right) \cdot f_{\text{avg}}(\pi^*, p). \quad (3)$$

They further proved an approximation guarantee for the minimum cost coverage objective. This approximation guarantee holds for functions with a *discrete covering property* first proposed by [Golovin and Krause \(2011\)](#). A utility function f is said to satisfy the discrete covering property with parameters $Q, \eta > 0$ if there exists a value $\eta \in (0, Q]$ such that for all $\psi \in \Psi_V$ and $\phi \in \Phi_V$, if $f(\text{dom}(\psi), \phi) > Q - \eta$, then $f(\text{dom}(\psi), \phi) = Q$. [Esfandiari et al. \(2021\)](#) proved that for any such function, and any greedy policy π , if π^* is a policy that achieves covering, that is, $f(E(\pi^*, \phi), \phi) = Q$ for all $\phi \in \Phi_V$, then

$$c_{\text{avg}}(\pi, p) \leq (c_{\text{avg}}(\pi^*, p) + 1) \log\left(\frac{|V|Q}{\eta}\right) + 1. \quad (4)$$

This improves a previous bound by [Golovin and Krause \(2017\)](#), which had a squared logarithmic dependence. [Golovin and Krause \(2011\)](#) studied Generalized Binary Search (GBS), which is a specific strategy for Bayesian active learning with binary labels, and defined a *coverage utility function* for active learning that has the discrete covering property with parameters $Q = 1, \eta = \min_{\phi \in \Phi_V} p(\phi)$. Plugging these parameters into Eq. (4) may be disadvantageous if the minimal probability assigned by p is very small, due to the term $\log(1/\eta)$ in the approximation factor. [Golovin and Krause \(2011, 2017\)](#) defined a modified prior $p' \propto \max\{p(\phi), 1/|\Phi_V|^2\}$, and derived an $O(\log^2(|\Phi_V|))$ approximation guarantee based on their squared-logarithmic approximation bound, for a greedy policy that uses a coverage utility function based on this prior.

4. Main results

In this section, we provide an overview of our main results. In Section 5 below, we define a new policy parameter, which we term the *maximal gain ratio*. This is the maximal ratio between the expected marginal gain of remaining elements at termination and the expected marginal gain of elements selected by the policy. The maximal gain ratio of a policy π given a utility function f and a prior p is denoted by $\beta_\pi(f, p)$. Our approximation guarantees, provided below, are parameterized by $\beta_\pi(f, p)$, and are stronger when $\beta_\pi(f, p)$ is smaller. It is therefore desirable to have $\beta_\pi(f, p)$ as small as possible. In Section 5, we prove that the value of $\beta_\pi(f, p)$ is never larger than that of $\alpha_\pi(f, p)$, the greedy approximation ratio of the policy, and that it can actually be significantly smaller. This shows that a guarantee that uses $\beta_\pi(f, p)$ is strictly stronger than the same guarantee with $\alpha_\pi(f, p)$. We prove the following properties:

- For all policies π , utility functions f and priors p , $\beta_\pi(f, p) \leq \alpha_\pi(f, p)$.
- β_π can have a value of 1 even if the policy is not greedy. This contrasts with α_π , which is equal to 1 if and only if the policy π is greedy.
- β_π can be arbitrarily close to zero. This is in contrast to α_π , which is bounded below by 1.

Having defined the maximal gain ratio, we provide new approximation guarantees that use this parameter, instead of the greedy approximation ratio or an assumption that a policy is greedy, as in previous works. In addition, these approximation guarantees support multiple regimes that were previously studied separately: functions that are not necessarily adaptive submodular, policies that are not necessarily greedy, and guarantees for the minimum cost coverage objective for functions with the discrete covering property.

Our first approximation guarantee concerns utility maximization.

Theorem 1 *Let f be a non-negative utility function which is adaptive monotone with respect to a prior p . Let π be a policy with height n that terminates when the expected marginal gain of all remaining elements is below some fixed threshold and has an average cost of l .¹ Let π^* be a policy with height k . Then*

$$f_{\text{avg}}(\pi, p) \geq \left(1 - \exp\left(-l / \left(\frac{\beta_\pi(f, p)}{\gamma_{n,k}^s(f, p)} \cdot c_{\text{avg}}(\pi^*, p) + 1\right)\right)\right) \cdot f_{\text{avg}}(\pi^*, p).$$

This result includes as a special case the conditions of Eq. (3), which require f to be adaptive submodular and π to be greedy. In this case, $\gamma_{n,k}^s(f, p) = 1$ and $\beta_\pi(f, p) \leq 1$. Thus, even in this special case, the theorem provides a smaller approximation factor whenever $\beta_\pi(f, p) < 1$. Moreover, the same guarantee as Eq. (3) is obtained also for non-greedy policies if they have $\beta_\pi(f, p) = 1$.

Our next approximation guarantee considers the minimum cost coverage objective.

Theorem 2 *Let f be a non-negative utility function which is adaptive monotone with respect to a prior p , such that f has the discrete covering property with parameters Q, η . Let π be a policy with height n . Let π^* be a policy with height k such that $f(E(\pi^*, \phi), \phi) = Q$ for all $\phi \in \Phi_V$. Then,*

$$c_{\text{avg}}(\pi, p) \leq \left(\frac{\beta_\pi(f, p)}{\gamma_{n,k}^s(f, p)} \cdot c_{\text{avg}}(\pi^*, p) + 1\right) \cdot \log\left(\frac{nQ}{\eta}\right) + 2.$$

1. We provide a formal version of this condition in Appendix A.

Considering again the special case in which f is adaptive submodular and π is greedy, we recover Eq. (4) up to an additive term of 1, noting that $n \leq |V|$. Similarly to Theorem 1, here too this special case is in fact broader and can provide stronger guarantees than the corresponding inequality Eq. (4). The two theorems above are proved in Appendix A.

Lastly, we provide approximation guarantees that allow providing a stronger guarantee for priors with small probabilities, using the modified prior p' of Golovin and Krause (2011). We prove our results for a more general setting with general utility functions, and apply them to the coverage utility function proposed by Golovin and Krause (2011) for Bayesian active learning. We conclude that when using a greedy policy with a coverage utility function based on the modified prior, denoted $f_{p'}$, we have the following average cost approximation guarantee for any policy π , compared to the optimal policy π^* .

$$c_{\text{avg}}(\pi, p) \leq 2(\beta_{\pi}(f_{p'}, p') \cdot (c_{\text{avg}}(\pi^*, p) + 1) + 1) \log(2|\Phi_V|^2 n) + 4. \quad (5)$$

In particular, if π is greedy with respect to p' , then $\beta_{\pi}(f_{p'}, p') \leq 1$, and it could be significantly smaller, as discussed above. We provide the necessary definitions and prove Eq. (5) in Section 6.

Our results thus provide a comprehensive and improved set of approximation guarantees that encompass a wide range of problem parameters, and highlight the usefulness of the new maximal gain ratio parameter. In the next section, we define the maximal gain ratio and prove its properties.

5. The Maximal Gain Ratio

In this section, we define the maximal gain ratio and prove its properties as listed in Section 4. We start with necessary definitions. Given a utility function f , a prior p and a policy π , let $\tau \geq 0$ and $\rho \in [0, 1]$ be real numbers. We denote by $\pi^{\tau, \rho}$ a sub-policy of π that with probability ρ terminates if the expected marginal gain of every single remaining element is strictly smaller than τ , and otherwise (with probability $1 - \rho$), terminates if the expected marginal gain of every single remaining element is at most τ . We call τ a *threshold* and ρ a *tie-break probability*. To define the maximal gain ratio, we require τ, ρ such that $c_{\text{avg}}(\pi^{\tau, \rho}, p)$ is equal to a given integer. Esfandiari et al. (2021) studied these types of sub-policies when π is greedy, and provided a proof sketch for the existence of such τ, ρ in this context. We provide here a full unconditional existence proof for any policy π , and prove in addition that there is a unique such sub-policy for each possible integer.

Lemma 3 *Assume a utility function f , a policy π , and a prior p . For any integer $i \leq c_{\text{avg}}(\pi, p)$, there exist a threshold $\tau_i \geq 0$ and a tie-break probability $\rho_i \in [0, 1]$ such that $c_{\text{avg}}(\pi^{\tau_i, \rho_i}, p) = i$. Moreover, all such pairs τ_i, ρ_i induce the same policy.*

Proof For a threshold τ , denote $\mu^{\tau} := c_{\text{avg}}(\pi^{\tau, 1}, p)$. Define an equivalence relation between thresholds as follows: for every τ, τ' , say that $\tau \equiv \tau'$ if $\mu^{\tau} = \mu^{\tau'}$. Since the sets of elements and the set of possible states are finite, the set of all the marginal gains for a given prior is also finite. Hence, there is a finite number of equivalence classes. Denote this number s . Fix some representatives of the equivalence classes and denote them by $\tau^1 > \dots > \tau^s$. Note that $\mu^{\tau^s} = c_{\text{avg}}(\pi, p)$. Define $\mu^{\tau^0} = 0$, and note that the sequence $(\mu^{\tau^j})_{j \in \{0, \dots, s\}}$ is strictly increasing. Let $i \leq c_{\text{avg}}(\pi, p)$. Clearly, $i \leq \mu^{\tau^s}$. Let j be the smallest index in $[s]$ such that $i \leq \mu^{\tau^j}$. Denote $\tau_i := \tau^j$ and

$$\rho_i := \frac{i - \mu^{\tau_i^{j-1}}}{\mu^{\tau_i^j} - \mu^{\tau_i^{j-1}}}.$$

Note that $\rho_i \in [0, 1]$. Consider the policy π^{τ_i, ρ_i} . By the definition of thresholds and tie-break probabilities, it holds that

$$c_{\text{avg}}(\pi^{\tau_i, \rho_i}, p) = (1 - \rho_i)\mu^{\tau_i - 1} + \rho_i\mu^{\tau_i} = \mu^{\tau_i - 1} + \rho_i(\mu^{\tau_i} - \mu^{\tau_i - 1}) = i, \quad (6)$$

Where the last equality follows from the definition of ρ_i . This proves the existence of (τ_i, ρ_i) , as claimed in the first part of the lemma.

For the second part of the lemma, we show that any such pair would induce the same policy. Let τ' and ρ' be such that $c_{\text{avg}}(\pi^{\tau', \rho'}, p) = i$. First, note that one of $\pi^{\tau', \rho'}$ and π^{τ_i, ρ_i} must be a sub-policy of the other: If $\tau' < \tau_i$, then π^{τ_i, ρ_i} is a sub-policy of $\pi^{\tau', \rho'}$, since the former always terminates before or with the latter. If $\tau_i < \tau'$, then the symmetric claim holds. Lastly, if $\tau' = \tau_i$, then we may assume w.l.o.g that the policy with the larger ρ never terminates in cases where the other policy continues, by assuming a coordinated draw of random bits.

Next, assume for contradiction that $\pi^{\tau', \rho'} \neq \pi^{\tau_i, \rho_i}$. Therefore, there exists at least one partial realization ψ and random bits such that immediately after observing it, $\pi^{\tau', \rho'}$ terminates, while π^{τ_i, ρ_i} selects an element. Since no partial realization satisfies the opposite case, it follows that $i = c_{\text{avg}}(\pi^{\tau', \rho'}, p) < c_{\text{avg}}(\pi^{\tau_i, \rho_i}, p) = i$, a contradiction. Thus, in all cases, $\pi^{\tau', \rho'} = \pi^{\tau_i, \rho_i}$, as claimed. \blacksquare

Given a policy π and an integer $i \leq c_{\text{avg}}(\pi, p)$, denote by π_i the unique sub-policy of π using the threshold and tie-break construction and has an average cost of i . Denote by π_0 a policy that terminates before selecting any element. Denote by $\Delta_{\pi, i}^u$ the maximal expected marginal gain of remaining elements when π_i terminates in any of its possible paths. Formally, $\Delta_{\pi, i}^u$ is the smallest real value u such that

$$\mathbb{P}[\text{For } t \in \mathbb{N} \text{ such that } \pi_i \text{ terminates at time } t, \forall v \in V, \Delta_p^f(v \mid \psi_t) \leq u] = 1.$$

Denote by $\Delta_{\pi, i}^l$ the smallest possible expected marginal gain of elements that π_i selects at any time during its run. Formally, this is the largest real value u such that

$$\mathbb{P}[\text{For all } t \text{ until } \pi_i \text{ terminates, } \Delta_p^f(\pi_i(\psi_t) \mid \psi_t) \geq u] = 1.$$

Equipped with these definitions, we can now define the maximal gain ratio.

Definition 4 (maximal gain ratio) *The maximal gain ratio of a policy π with respect to a utility function f and a prior p is $\beta_\pi(f, p) := \max_{i \in \mathbb{N}, i \leq c_{\text{avg}}(\pi, p)} \Delta_{\pi, i}^u / \Delta_{\pi, i}^l$.*

The maximal gain ratio is thus based on comparing the possible gain of elements that were not selected, to the gain from selected elements. We now show that the maximal gain ratio never larger than the greedy approximation parameter, and is sometimes strictly smaller. This shows that Theorem 1, Theorem 2 and Eq. (5), are stronger due to their use of β_π than if they had relied on α_π instead.

First, we show that the maximal gain ratio of an approximate greedy policy is always upper bounded by the greedy approximation parameter of the same policy.

Theorem 5 *For any prior p , utility function f , and policy π , we have $\beta_\pi(f, p) \leq \alpha_\pi(f, p)$.*

Proof Let π be a policy such that $\alpha_\pi(f, p) < \infty$, and denote $\alpha := \alpha_\pi(f, p)$. Let $i \leq c_{\text{avg}}(\pi, p)$. Let τ_i and ρ_i be a threshold and a tie-break probability, respectively, such that $\pi_i := \pi^{\tau_i, \rho_i}$. By the definition of π^{τ_i, ρ_i} , we have that τ_i is an upper bound on the expected marginal gain of all the elements that remain when π^{τ_i, ρ_i} terminates. Therefore, $\Delta_{\pi, i}^u \leq \tau_i$.

At all times during the run of π_i up to the time of termination, there exists some element with an expected marginal gain of at least τ_i . Otherwise, by definition of $\pi_i = \pi^{\tau_i, \rho_i}$, the run would have terminated earlier. Therefore, since π is an α -approximate greedy policy, the expected marginal gain of the elements that π_i selects at any time during its run is at least τ_i/α . It follows that $\Delta_{\pi, i}^l \geq \tau_i/\alpha$. Therefore,

$$\beta_\pi(f, p) = \max_{i \leq c_{\text{avg}}(\pi, p)} \Delta_{\pi, i}^u / \Delta_{\pi, i}^l \leq \max_{i \leq c_{\text{avg}}(\pi, p)} \frac{\tau_i}{\tau_i/\alpha} = \alpha.$$

This completes the proof. \blacksquare

We now show that there are cases in which β_π is strictly smaller than α_π . First, we show that an approximate greedy policy that is not greedy can still have $\beta_\pi(f, p) = 1$.

Theorem 6 *For any integer $k \geq 3$, there exist a policy π with height k , a prior p and a utility function f such that $\beta_\pi(f, p) = 1$ while $\alpha_\pi(f, p) = 2$.*

Proof For a given integer $k \geq 3$, let $V = \{v_1, \dots, v_k\}$, and let $Y = \{0, 1\}$. For any $i \in [k]$, we denote $v_{[i]} := \{v_1, \dots, v_i\}$. We define a utility function f as follows. For a realization ϕ and a subset $A \subseteq V$, Let π be a policy that selects v_i in the i 'th step of its run and terminates immediately after selecting v_k . By the properties of $\Delta_p^f(v | \psi)$, if the observed partial realization is ψ such that $\text{dom}(\psi) = \emptyset$ or $\text{dom}(\psi) = v_{[k-1]}$, then the next element v that π selects satisfies $\Delta_p^f(v | \emptyset) = 1/k = \max_{v \in V} \Delta_p^f(v | \psi)$. Additionally, if $\text{dom}(\psi) = v_{[i-1]}$ for some $i \in \{2, \dots, k-1\}$, then

$$\Delta_p^f(v_i | \psi) = \frac{1}{k} = \frac{1}{2} \cdot \frac{2}{k} = \frac{1}{2} \cdot \Delta_p^f(v_{i+1} | \psi) = \frac{1}{2} \cdot \max_{v \in V} \Delta_p^f(v | \psi).$$

Therefore, $\alpha_\pi(f, p) = 2$.

On the other hand, we now show that $\beta_\pi(f, p) = 1$. Let $i \leq c_{\text{avg}}(\pi, p)$. First, we show that for $\tau = 1/k$ and $\rho = i/k$, it holds that $\pi_i = \pi^{\tau, \rho}$. Due to the value of τ , with probability ρ , $\pi^{\tau, \rho}$ selects all of the elements of V and terminates. Otherwise (with probability $1 - \rho$) $\pi^{\tau, \rho}$ terminates before selecting any element. Therefore,

$$c_{\text{avg}}(\pi^{\tau, \rho}, p) = \rho \cdot k + (1 - \rho) \cdot 0 = \frac{i}{k} \cdot k = i.$$

By Lemma 3, $\pi_i = \pi^{\tau, \rho}$. Now, the largest possible expected marginal gain of remaining elements when π_i terminates is $1/k$. Hence, $\Delta_{\pi, i}^u = 1/k$. The smallest possible expected marginal gain of any element selected by π_i is also $1/k$. Hence, $\Delta_{\pi, i}^l = 1/k$. Therefore,

$$\beta_\pi(f, p) = \max_{i \leq c_{\text{avg}}(\pi, p)} \Delta_{\pi, i}^u / \Delta_{\pi, i}^l = \max_{i \leq c_{\text{avg}}(\pi, p)} \frac{1/k}{1/k} = 1.$$

This completes the proof. \blacksquare

Lastly, we show that there exist policies π that are greedy with respect to a utility function f and a prior p that have an arbitrarily small $\beta_\pi(f, p)$. Thus, guarantees based on $\beta_\pi(f, p)$ are strictly stronger than guarantees based only on the fact that a policy is greedy.

Theorem 7 *For any integer k and any $\epsilon > 0$, there exist a prior p , a utility function f and a greedy policy π with respect to f that selects k elements, such that $\beta_\pi(f, p) = \epsilon$.*

Proof Let k be an integer and let $\epsilon > 0$. Let $V = \{v_1, v_2, \dots, v_k\}$, $Y = \{0, 1\}$. We define a utility function f as follows. For any realization ϕ and any subset $A \subseteq V$, $f(A, \phi) = \sum_{i=0}^{|A|} \epsilon^i$. Let the prior p be a uniform distribution over the realizations. By the definition of f , for any partial realization ψ , it holds that $\Delta_p^f(v \mid \psi) = \epsilon^{|\psi|+1}$.

Let π be a policy that selects v_i in the i 'th step of its run and terminates immediately after selecting v_k . For any partial realization ψ , all of the elements $v \in V$ has the same value of $\Delta_p^f(v \mid \psi)$. Therefore π is greedy with respect to f .

Recall that $\beta_\pi(f, p) = \max_{i \leq c_{\text{avg}}(\pi, p)} \Delta_{\pi, i}^u / \Delta_{\pi, i}^l$. To bound $\beta_\pi(f, p)$, consider the possible values of $i \leq c_{\text{avg}}(\pi, p) = k$. For $i = k$, the largest possible expected marginal gain of remaining elements when π_n terminates is 0. Therefore, $\Delta_{\pi, k}^u = 0$ and so $\Delta_{\pi, k}^u / \Delta_{\pi, k}^l = 0$. For $i \leq k - 1$, the largest possible expected marginal gain of remaining elements when π_i terminates is ϵ^{i+1} . Therefore, $\Delta_{\pi, i}^u = \epsilon^{i+1}$. The smallest possible expected marginal gain of elements that π_i selects at any time during its run is ϵ^i . Therefore, $\Delta_{\pi, i}^l = \epsilon^i$. Hence, $\Delta_{\pi, i}^u / \Delta_{\pi, i}^l = \epsilon^{i+1} / \epsilon^i = \epsilon$. It follows that $\beta_\pi(f, p) := \max_{i \leq c_{\text{avg}}(\pi, p)} \Delta_{\pi, i}^u / \Delta_{\pi, i}^l = \max\{0, \epsilon\} = \epsilon$, as claimed \blacksquare

We note that it has been shown in previous works (Conforti and Cornuéjols, 1984; Vondrak, 2010) that for some utility functions, stronger approximation guarantees can be provided for greedy policies than those that can be obtained for general submodular functions, based on the curvature property of the function. It is thus helpful to observe that guarantees based on curvature cannot be used to provide stronger approximation guarantees in the example provided in the proof above. Curvature-based guarantees are stronger if the utility function is closer to being modular (additive). In contrast, in the example above, the function is far from being modular, as adding the same element later in the run results in a significantly lower gain. More precisely, the curvature property takes values in $[0, 1]$ for submodular functions, and results concerning curvature are stronger than results for general submodular functions if the curvature is smaller. However, the curvature of the function in the proof above is $1 - \epsilon^{k-1}$, making it a poor candidate to gain from curvature-based bounds.

The properties proved above for the maximal gain ratio imply that Theorem 1 and Theorem 2 are strictly stronger than those in previous works, even for adaptive submodular utility functions and greedy policies. The proofs of these theorems are provided in Appendix A.

6. Using a modified prior

In this section, we present an approximation guarantee that uses the maximal gain ratio and the adaptive submodularity ratio for the modified prior, and use it to conclude Eq. (5), which provides an approximation guarantee that does not depend on the smallest probability in the true prior.

Theorem 8 *Let f be a non-negative utility function which is adaptive monotone with respect to a prior p , such that f has the discrete covering property with parameters Q, η . Let π be a policy with height n . Let π^* be a policy with height $k \leq |\Phi_V|$ such that $f(E(\pi^*, \phi), \phi) = Q$ for all $\phi \in \Phi_V$. Let $p'(\phi) \propto \max\{p(\phi), 1/|\Phi_V|^2\}$ be the modified prior of p . Then*

$$c_{\text{avg}}(\pi, p) \leq 2 \left(\frac{\beta_\pi(f, p')}{\gamma_{n, k}^s(f, p')} (c_{\text{avg}}(\pi^*, p) + 1) + 1 \right) \log\left(\frac{nQ}{\eta}\right) + 4.$$

This result improves over the previous guarantees of [Golovin and Krause \(2011, 2017\)](#) for the modified prior, since the latter had a square logarithmic factor and only applied to greedy policies. The proof of this theorem is provided in [Appendix B](#).

Before deriving [Eq. \(5\)](#), we provide background on the coverage utility function proposed by [Golovin and Krause \(2011\)](#) and its application to active learning. In the context of active learning, the observations in combinatorial maximization setting represent labeled examples. The coverage function measures the probability mass of realizations that have been disqualified due to the observations in the observed partial realization. To define the coverage function formally, denote the *version space* of $\psi \in \Psi$ by $\text{VS}(\psi) := \{\phi \in \Phi_V \mid \phi \sim \psi\}$. For $\mathcal{H} \subseteq \Phi$, denote $p(\mathcal{H}) := \sum_{\phi \in \mathcal{H}} p(\phi)$. For $A \subseteq V$, $\phi \in \Phi$, denote $\text{VS}_{A,\phi} = \text{VS}(\{(x, \phi(x)) \mid x \in A\})$.

Definition 9 (coverage utility function) *For a prior p over Φ_V , the coverage utility function is denoted $f_p : 2^V \times \Phi_V \rightarrow [0, 1]$ and defined such that for any realization ϕ and any set $A \subseteq V$, $f_p(A, \phi) := 1 - p(\text{VS}_{A,\phi}) + p(\phi)$.*

[Golovin and Krause \(2011\)](#) proved that for any prior p , f_p is adaptive monotone and adaptive submodular with respect to p . The latter implies that $\gamma_{n,k}^s(f_{p'}, p') = 1$. In addition, note that if $\text{VS}_{A,\phi} = \{\phi\}$, that is, the true realization is fully identified by observing the labels of the examples in A , then $f_p(A, \phi) = 1$, which is the function's maximal possible value. Thus, in this case we can set $Q = 1$ in [Theorem 8](#). In addition, for f_p , $\eta = \min_{\phi \in \Phi_V} p(\phi)$. When using $f_{p'}$ with the modified prior p' , we get

$$\eta = \min_{\phi \in \Phi_V} p'(\phi) \geq \frac{1/|\Phi_V|^2}{\sum_{\phi \in \Phi_V} \max\{p(\phi), 1/|\Phi_V|^2\}} \geq \frac{1/|\Phi_V|^2}{1 + 1/|\Phi_V|} \geq \frac{1}{2|\Phi_V|^2}.$$

By substituting these in [Theorem 8](#), we conclude that for any policy π^* with height at most $|\Phi_V|$, [Eq. \(5\)](#) holds. We have left to show that π^* can be set to an optimal policy, by proving that there exists an optimal policy with height at most $|\Phi_v|$. While this makes intuitive sense, we formally state and prove it for completeness in [Lemma 11](#) in [Appendix B](#). This completes the derivation of [Eq. \(5\)](#).

7. Conclusion

In this work, we provided new approximation guarantees for adaptive combinatorial maximization. We identified a new policy parameter, the maximal gain ratio, which is less restrictive than the greedy approximation parameter, and showed that using this parameter can lead to stronger approximation guarantees. Moreover, our guarantees are comprehensive, simultaneously supporting the new policy parameter as well as near-submodular utility functions, and include both maximization under a cardinality constraint and a minimum cost coverage guarantee. In addition, we provided an improved approximation guarantee for a modified prior, which is crucial for active learning guarantees that do not depend on the smallest probability in the prior.

The new maximal gain ratio parameter sheds new light on the properties a policy needs to have to support general approximation guarantees for adaptive combinatorial maximization. This opens an exciting direction of research on the relationship between policy properties and the approximations that can be obtained based on them.

Acknowledgments

Resources used in preparing this research were provided, in part, by the Province of Ontario, the Government of Canada through CIFAR, and companies sponsoring the Vector Institute; see <https://vectorinstitute.ai/partnerships/current-partners/>.

References

- Georgios Amanatidis, Federico Fusco, Philip Lazos, Stefano Leonardi, and Rebecca Reiffenhäuser. Fast adaptive non-monotone submodular maximization subject to a knapsack constraint. *Advances in neural information processing systems*, 33:16903–16915, 2020.
- Arash Asadpour, Hamid Nazerzadeh, and Amin Saberi. Stochastic submodular maximization. In *Internet and Network Economics: 4th International Workshop, WINE 2008, Shanghai, China, December 17-20, 2008. Proceedings 4*, pages 477–489. Springer, 2008.
- Eric Balkanski and Yaron Singer. Approximation guarantees for adaptive sampling. In *International Conference on Machine Learning*, pages 384–393. PMLR, 2018.
- Yuxin Chen and Andreas Krause. Near-optimal batch mode active learning and adaptive submodular optimization. In *International Conference on Machine Learning*, pages 160–168. PMLR, 2013.
- Michele Conforti and Gérard Cornuéjols. Submodular set functions, matroids and the greedy algorithm: tight worst-case bounds and some generalizations of the rado-edmonds theorem. *Discrete applied mathematics*, 7(3):251–274, 1984.
- Nguyen Viet Cuong, Wee Sun Lee, and Nan Ye. Near-optimal adaptive pool-based active learning with general loss. In *30th conference on Uncertainty in Artificial Intelligence (UAI)*, pages 122–131, 2014.
- Abhimanyu Das and David Kempe. Submodular meets spectral: Greedy algorithms for subset selection, sparse approximation and dictionary selection. *arXiv preprint arXiv:1102.3975*, 2011.
- Janardhan Rao Doppa. Adaptive experimental design for optimizing combinatorial structures. In *IJCAI*, pages 4940–4945, 2021.
- Jack Edmonds. Submodular functions, matroids, and certain polyhedra. *Combinatorial Structures and Their Applications*, pages 69–87, 1970.
- Hossein Esfandiari, Amin Karbasi, and Vahab Mirrokni. Adaptivity in adaptive submodularity. In *Conference on Learning Theory*, pages 1823–1846. PMLR, 2021.
- Kaito Fujii and Shinsaku Sakaue. Beyond adaptive submodularity: Approximation guarantees of greedy policy with adaptive submodularity ratio. In *International Conference on Machine Learning*, pages 2042–2051. PMLR, 2019.
- Victor Gabillon, Branislav Kveton, Zheng Wen, Brian Eriksson, and Shanmugavelayutham Muthukrishnan. Adaptive submodular maximization in bandit setting. *Advances in Neural Information Processing Systems*, 26, 2013.

- 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.
- Daniel Golovin and Andreas Krause. Adaptive submodularity: A new approach to active learning and stochastic optimization. *CoRR*, abs/1003.3967v5, 2017.
- Alkis Gotovos, Amin Karbasi, and Andreas Krause. Non-monotone adaptive submodular maximization. In *IJCAI*, pages 1996–2003, 2015.
- Andrew Guillory and Jeff Bilmes. Interactive submodular set cover. *arXiv preprint arXiv:1002.3345*, 2010.
- Hideitsu Hino. Active learning: Problem settings and recent developments. *arXiv preprint arXiv:2012.04225*, 2020.
- Andrew McCallum, Kamal Nigam, et al. Employing em and pool-based active learning for text classification. In *ICML*, volume 98, pages 350–358. Citeseer, 1998.
- Marko Mitrovic, Ehsan Kazemi, Moran Feldman, Andreas Krause, and Amin Karbasi. Adaptive sequence submodularity. *Advances in Neural Information Processing Systems*, 32, 2019.
- George L Nemhauser, Laurence A Wolsey, and Marshall L Fisher. An analysis of approximations for maximizing submodular set functions—i. *Mathematical Programming*, 14(1):265–294, 1978.
- Sivan Sabato. Submodular learning and covering with response-dependent costs. *Theoretical Computer Science*, 742:98 – 113, 2018.
- Lior Seeman and Yaron Singer. Adaptive seeding in social networks. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 459–468. IEEE, 2013.
- Shaojie Tang. Beyond pointwise submodularity: Non-monotone adaptive submodular maximization in linear time. *Theoretical Computer Science*, 850:249–261, 2021. ISSN 0304-3975.
- Shaojie Tang. Robust adaptive submodular maximization. *INFORMS Journal on Computing*, 34(6):3277–3291, 2022.
- Shaojie Tang and Jing Yuan. Streaming adaptive submodular maximization. *Theoretical Computer Science*, 944:113644, 2023.
- Guangmo Tong, Weili Wu, Shaojie Tang, and Ding-Zhu Du. Adaptive influence maximization in dynamic social networks. *IEEE/ACM Transactions on Networking*, 25(1):112–125, 2016.
- Jan Vondrak. Submodularity and curvature: the optimal algorithm. *RIMS Kôkyûroku Bessatsu*, 01 2010.
- Laurence A Wolsey. An analysis of the greedy algorithm for the submodular set covering problem. *Combinatorica*, 2(4):385–393, 1982.
- Ji-Jie Wu and Kuo-Shih Tseng. Adaptive submodular inverse reinforcement learning for spatial search and map exploration. *Autonomous Robots*, 46(2):321–347, 2022.

Appendix A. Proving the approximation guarantees for the given prior

In this section, we prove Theorem 1 and Theorem 2, stated in Section 4. To prove Theorem 1, we first prove the following lemma, which provides a lower bound for the difference between the average utility of two sub-policies.

Lemma 10 *Let π be a policy, p be a prior, and f be a utility function. For any integer $i \leq c_{\text{avg}}(\pi, p)$, define $\Delta_i := f_{\text{avg}}(\pi_i, p) - f_{\text{avg}}(\pi_{i-1}, p)$. Then $\Delta_i \geq \Delta_{\pi, i}^l$.*

Proof Let $i \leq c_{\text{avg}}(\pi, p)$. Let $\tau_i, \tau_{i-1}, \rho_i, \rho_{i-1}$ be parameters such that $\pi_i = \pi^{\tau_i, \rho_i}$ and $\pi_{i-1} = \pi^{\tau_{i-1}, \rho_{i-1}}$. These exist according to Lemma 3. Observe that necessarily, $\tau_i \leq \tau_{i-1}$. This is because if $\tau_{i-1} < \tau_i$, then the termination condition of π_i is stronger than that of π_{i-1} , implying that $i = c_{\text{avg}}(\pi_i, p) \leq c_{\text{avg}}(\pi_{i-1}, p) = i - 1$, a contradiction.

Observe that π_{i-1} must be a sub-policy of π_i : If $\tau_i < \tau_{i-1}$, then π_{i-1} is a sub-policy of π_i , since the former always terminates before or with the latter. If $\tau_i = \tau_{i-1}$, then $\rho_i < \rho_{i-1}$. In this case, π_i never terminates in cases where π_{i-1} continues, by assuming a coordinated draw of random bits.

For a set of elements S , a realization ϕ and a partial realization ψ , we define $\Delta_p^f(S | \psi) := \mathbb{E}[f(S \cup \text{dom}(\psi), \phi) - f(\text{dom}(\psi), \phi) | \phi \sim \psi]$. Let $A_{\psi, i, S}$ to be the event that the following hold:

1. The policy π_{i-1} selects exactly all of the elements of $\text{dom}(\psi)$,
2. The observations in ψ are realized, and
3. the set of elements selected by π_i coincides with $\text{dom}(\psi)$ at some point in its run, and immediately after that, π_i selects exactly all of the elements of S and terminates.

Since π_{i-1} is a sub-policy of π_i , we have

$$\begin{aligned}
 \Delta_i &= f_{\text{avg}}(\pi_i, p) - f_{\text{avg}}(\pi_{i-1}, p) \\
 &= \sum_{\psi \in \Psi_V} \sum_{S \subseteq V} \mathbb{P}[A_{\psi, i, S}] \Delta_p^f(S | \psi) \\
 &\geq \sum_{\psi \in \Psi_V} \sum_{S \subseteq V} \mathbb{P}[A_{\psi, i, S}] \cdot |S| \cdot \Delta_{\pi, i}^l \\
 &= (c_{\text{avg}}(\pi_i, p) - c_{\text{avg}}(\pi_{i-1}, p)) \cdot \Delta_{\pi, i}^l \\
 &= (i - (i - 1)) \cdot \Delta_{\pi, i}^l = \Delta_{\pi, i}^l,
 \end{aligned}$$

where the inequality follows from the definition of $\Delta_{\pi, i}^l$. This proves the claim. ■

Using this lemma, we can now prove Theorem 1.

The statement of Theorem 1 includes a condition that the policy π terminates when the expected marginal gain of all remaining elements is below some fixed threshold and has an average cost of l . Formally, we require π to be equal to π'_l for some policy π' (see Lemma 3 and the definition of π_i thereafter).

Proof [of Theorem 1] For any realization ϕ and any element $v \in V$, denote $\mathbb{1}_\phi^v := \mathbb{1}[v \in E(\pi^*, \phi)]$. For any partial realization ψ such that $|\psi| \leq n$, by the definition of $\gamma_{n,k}^s(f, p)$, we have

$$\begin{aligned} \gamma_{n,k}^s(f, p) \cdot \Delta_p^f(\pi^* | \psi) &\leq \sum_{v \in V} \mathbb{P}_{\Phi \sim p}[v \in E(\pi^*, \Phi) | \Phi \sim \psi] \Delta_p^f(v | \psi) \\ &= \sum_{v \in V} \mathbb{E}[\mathbb{1}[v \in E(\pi^*, \Phi)] | \Phi \sim \psi] \Delta_p^f(v | \psi) \\ &\leq \mathbb{E}[|E(\pi^*, \Phi)| | \Phi \sim \psi] \max_{v \in V} \Delta_p^f(v | \psi). \end{aligned}$$

Therefore,

$$\Delta_p^f(\pi^* | \psi) \leq \frac{\mathbb{E}[|E(\pi^*, \Phi)| | \Phi \sim \psi] \max_{v \in V} \Delta_p^f(v | \psi)}{\gamma_{n,k}^s(f, p)}. \quad (7)$$

For any integer $i \leq c_{\text{avg}}(\pi_i, p)$, let Δ_i be defined as in Lemma 10. For two policies π', π'' , we denote by $\pi' @ \pi''$ the policy that first runs π' and then runs π'' from the beginning, without taking into account any information collected during the execution of π' . Note that $E(\pi' @ \pi'', \phi) = E(\pi', \phi) \cup E(\pi'', \phi)$ for every realization ϕ . In particular, it holds that $E(\pi_i @ \pi^*, \phi) = E(\pi_i, \phi) \cup E(\pi^*, \phi)$ for any realization ϕ . Therefore, by the adaptive monotonicity of f , we have $f_{\text{avg}}(\pi^*, p) \leq f_{\text{avg}}(\pi_i @ \pi^*, p)$. For any $\psi \in \Psi_V$, we denote by A_ψ the event that π_i observes ψ and then terminates. For $i \in \{0, 1, \dots, l\}$ we define $\Delta_i^* = f_{\text{avg}}(\pi^*, p) - f_{\text{avg}}(\pi_i, p)$. For $i \in [l]$,

$$\begin{aligned} \Delta_i^* &= f_{\text{avg}}(\pi^*, p) - f_{\text{avg}}(\pi_i, p) \\ &\leq f_{\text{avg}}(\pi_i @ \pi^*, p) - f_{\text{avg}}(\pi_i, p) \\ &= \sum_{\psi \in \Psi_V} \mathbb{P}[A_\psi] \Delta_p^f(\pi^* | \psi) \\ &\leq \frac{\sum_{\psi \in \Psi_V} \mathbb{P}[A_\psi] \mathbb{E}[|E(\pi^*, \Phi)| | \Phi \sim \psi] \max_{v \in V} \Delta_p^f(v | \psi)}{\gamma_{n,k}^s(f, p)}, \end{aligned}$$

where the last inequality is by Eq. (7). By the definition of $\Delta_{\pi,i}^u$, it holds that $\max_{v \in V} \Delta_p^f(v | \psi) \leq \Delta_{\pi,i}^u$. By the definition of $\beta_\pi(f, p)$, we have $\Delta_{\pi,i}^u \leq \beta_\pi(f, p) \Delta_{\pi,i}^l \leq \beta_\pi(f, p) \Delta_i$, where the last inequality follows from Lemma 10. Therefore, $\max_{v \in V} \Delta_p^f(v | \psi) \leq \beta_\pi(f, p) \Delta_i$. It follows that

$$\Delta_i^* \leq \frac{\sum_{\psi \in \Psi_V} \mathbb{P}[A_\psi] \mathbb{E}[|E(\pi^*, \Phi)| | \Phi \sim \psi] \cdot \beta_\pi(f, p) \Delta_i}{\gamma_{n,k}^s(f, p)}.$$

In addition,

$$\sum_{\psi \in \Psi_V} \mathbb{P}[A_\psi] \mathbb{E}[|E(\pi^*, \Phi)| | \Phi \sim \psi] = \mathbb{E}[|E(\pi^*, \Phi)|] = c_{\text{avg}}(\pi^*, p).$$

Let $c^* = c_{\text{avg}}(\pi^*, p)$. Hence,

$$\Delta_i^* \leq \frac{\beta_\pi(f, p) \cdot c^*}{\gamma_{n,k}^s(f, p)} \cdot \Delta_i.$$

By the definitions of Δ_i and Δ_i^* , it holds that

$$\begin{aligned}\Delta_i &= f_{\text{avg}}(\pi_i, p) - f_{\text{avg}}(\pi_{i-1}, p) \\ &= f_{\text{avg}}(\pi^*, p) - f_{\text{avg}}(\pi_{i-1}, p) - (f_{\text{avg}}(\pi^*, p) - f_{\text{avg}}(\pi_i, p)) = \Delta_{i-1}^* - \Delta_i^*.\end{aligned}$$

Therefore,

$$\Delta_i^* \leq \frac{\beta_\pi(f, p) \cdot c^*}{\gamma_{n,k}^s(f, p)} (\Delta_{i-1}^* - \Delta_i^*).$$

Rearranging, we conclude that

$$\Delta_i^* \leq \left(1 - \frac{1}{\beta_\pi(f, p) \cdot c^* / \gamma_{n,k}^s(f, p) + 1}\right) \cdot \Delta_{i-1}^*.$$

By induction, for any integer l we obtain

$$\Delta_l^* \leq \left(1 - \frac{1}{\beta_\pi(f, p) \cdot c^* / \gamma_{n,k}^s(f, p) + 1}\right)^l \cdot \Delta_0^* \leq \exp\left(-\frac{l}{\beta_\pi(f, p) \cdot c^* / \gamma_{n,k}^s(f, p) + 1}\right) \cdot \Delta_0^*.$$

Since f is non-negative, we have $f_{\text{avg}}(\pi_0, p) \geq 0$. Therefore, by the definition of Δ_0^* , $\Delta_0^* \leq f_{\text{avg}}(\pi^*, p)$. Hence,

$$\Delta_l^* \leq \exp\left(-\frac{l}{\beta_\pi(f, p) \cdot c^* / \gamma_{n,k}^s(f, p) + 1}\right) \cdot f_{\text{avg}}(\pi^*, p).$$

Therefore, by the definition of Δ_l^* , it follows that

$$f_{\text{avg}}(\pi, p) \geq \left(1 - \exp\left(-\frac{l}{\beta_\pi(f, p) \cdot c^* / \gamma_{n,k}^s(f, p) + 1}\right)\right) \cdot f_{\text{avg}}(\pi^*, p).$$

This completes the proof. ■

Next, we prove Theorem 2.

Proof [of Theorem 2] Let $c^* = c_{\text{avg}}(\pi^*, p)$. Define

$$l := \lceil (\beta_\pi(f, p) \cdot c^* / \gamma_{n,k}^s(f, p) + 1) \cdot \log(nQ/\eta) \rceil.$$

By Theorem 1, we have

$$\begin{aligned}f_{\text{avg}}(\pi_l, p) &\geq \left(1 - \exp\left(-\frac{l}{\beta_\pi(f, p) \cdot c^* / \gamma_{n,k}^s(f, p) + 1}\right)\right) \cdot f_{\text{avg}}(\pi^*, p) \\ &\geq \left(1 - \exp\left(-\frac{(\beta_\pi(f, p) \cdot c^* / \gamma_{n,k}^s(f, p) + 1) \cdot \log(nQ/\eta)}{\beta_\pi(f, p) \cdot c^* / \gamma_{n,k}^s(f, p) + 1}\right)\right) \cdot f_{\text{avg}}(\pi^*, p) \\ &= (1 - \exp(-\log(nQ/\eta))) \cdot f_{\text{avg}}(\pi^*, p) \\ &= \left(1 - \frac{\eta}{nQ}\right) \cdot f_{\text{avg}}(\pi^*, p) \\ &= Q - \eta/n,\end{aligned}$$

Where the last inequality follows since $f_{\text{avg}}(\pi^*, p) = Q$ by the assumptions of the theorem.

Denote by F the (random) value of f when π_l terminates. By definition, $\mathbb{E}[F] = f_{\text{avg}}(\pi_l, p)$. By the definition of f and Q , $F \leq Q$ with probability one. Therefore, the random variable $X := Q - F$ is non-negative. By Markov's inequality, $\mathbb{P}[X \geq \eta] \leq \mathbb{E}[X]/\eta$. Hence,

$$\begin{aligned} \mathbb{P}[F > Q - \eta] &= 1 - \mathbb{P}[F \leq Q - \eta] = 1 - \mathbb{P}[X \geq \eta] \geq 1 - \mathbb{E}[X]/\eta \\ &= 1 - \frac{Q - f_{\text{avg}}(\pi_l, p)}{\eta} \geq 1 - \frac{Q - (Q - \eta/n)}{\eta} = 1 - \frac{1}{n}. \end{aligned}$$

By the assumptions of the theorem, if $f(\text{dom}(\psi), \phi) > Q - \eta$, then $f(\text{dom}(\psi), \phi) = Q$ for every $\psi \in \Psi_V$ and $\phi \in \Phi_V$. Therefore, with probability at least $1 - 1/n$, we have $F = Q$. Let $t := \mathbb{E}[|E(\pi_l, \phi)| \mid F = Q]$ be the expected cost of π_l conditioned on $F = Q$. Then $l = c_{\text{avg}}(\pi_l, p) \geq (1 - 1/n) \cdot t$, thus $t \leq l/(1 - 1/n)$. If $F \neq Q$, which happens with probability at most $1/n$, then π selects at most n elements. Hence,

$$c_{\text{avg}}(\pi, p) \leq \left(1 - \frac{1}{n}\right) \cdot t + \frac{1}{n} \cdot n \leq l + 1 \leq \left(\frac{\beta_\pi(f, p) \cdot c^*}{\gamma_{n,k}^s(f, p)} + 1\right) \cdot \log\left(\frac{nQ}{\eta}\right) + 2.$$

This completes the proof. ■

Appendix B. Proofs for Section 6

First, we prove Theorem 8.

Proof [of Theorem 8]

By the definition of p' , for any realization ϕ , we have $p'(\phi) := \max\{p(\phi), 1/|\Phi_V|^2\}/Z$, where Z is a normalizing constant. We have

$$Z = \sum_{\phi \in \Phi_V} \max\{p(\phi), 1/|\Phi_V|^2\} \leq \sum_{\phi \in \Phi_V} (p(\phi) + 1/|\Phi_V|^2) = 1 + 1/|\Phi_V|. \quad (8)$$

Therefore, for any realization ϕ ,

$$p'(\phi) = \max\{p(\phi), 1/|\Phi_V|^2\}/Z \geq p(\phi)/Z \geq p(\phi) \cdot \frac{|\Phi_V|}{|\Phi_V| + 1}.$$

Hence,

$$c_{\text{avg}}(\pi, p) = \sum_{\phi \in \Phi_V} p(\phi) |E(\pi, \phi)| \leq \frac{|\Phi_V| + 1}{|\Phi_V|} \sum_{\phi \in \Phi_V} p'(\phi) |E(\pi, \phi)| = \frac{|\Phi_V| + 1}{|\Phi_V|} c_{\text{avg}}(\pi, p').$$

Therefore,

$$c_{\text{avg}}(\pi, p) \leq \frac{|\Phi_V| + 1}{|\Phi_V|} \cdot c_{\text{avg}}(\pi, p') \leq 2c_{\text{avg}}(\pi, p').$$

By Theorem 2,

$$c_{\text{avg}}(\pi, p') \leq \left(\frac{\beta_\pi(f, p')}{\gamma_{n,k}^s(f, p')} \cdot c_{\text{avg}}(\pi^*, p') + 1\right) \log\left(\frac{nQ}{\eta}\right) + 2.$$

Therefore,

$$c_{\text{avg}}(\pi, p) \leq 2 \left(\frac{\beta_\pi(f, p')}{\gamma_{n,k}^s(f, p')} \cdot c_{\text{avg}}(\pi^*, p') + 1 \right) \log\left(\frac{nQ}{\eta}\right) + 4.$$

By the assumptions of the theorem, π^* has height at most $|\Phi_V|$. In other words, for any realization ϕ ,

$$|E(\pi^*, \phi)| \leq |\Phi_V|. \quad (9)$$

In addition, note that

$$Z = \sum_{\phi \in \Phi_V} \max\{p(\phi), 1/|\Phi_V|^2\} \geq \sum_{\phi \in \Phi_V} p(\phi) = 1.$$

Therefore, for any realization ϕ , if $p(\phi) \geq 1/|\Phi_V|^2$, then $p'(\phi) - p(\phi) = p(\phi)/Z - p(\phi) \leq 0$. Otherwise, $p(\phi) < 1/|\Phi_V|^2$, which implies $p'(\phi) = 1/(|\Phi_V|^2 Z) \leq 1/|\Phi_V|^2$. We conclude that for any realization ϕ , $p'(\phi) - p(\phi) \leq 1/|\Phi_V|^2$. Combining this with Eq. (9), we get

$$\begin{aligned} c_{\text{avg}}(\pi^*, p') - c_{\text{avg}}(\pi^*, p) &= \sum_{\phi \in \Phi_V} |E(\pi^*, \phi)|(p'(\phi) - p(\phi)) \\ &\leq \sum_{\phi \in \Phi_V} |\Phi_V|(p'(\phi) - p(\phi)) \\ &\leq \sum_{\phi \in \Phi_V} |\Phi_V| \frac{1}{|\Phi_V|^2} = 1. \end{aligned}$$

We conclude that $c_{\text{avg}}(\pi^*, p') \leq c_{\text{avg}}(\pi^*, p) + 1$. Therefore,

$$c_{\text{avg}}(\pi, p) \leq 2 \left(\frac{\beta_\pi(f, p')}{\gamma_{n,k}^s(f, p')} \cdot (c_{\text{avg}}(\pi^*, p) + 1) + 1 \right) \log\left(\frac{nQ}{\eta}\right) + 4.$$

This completes the proof. ■

Next, we state and prove Lemma 11.

Lemma 11 *Let π^* be a policy such that $f_p(E(\pi^*, \phi), \phi) = 1$ for all $\phi \in \Phi_V$ and such that $c_{\text{avg}}(\pi^*, p)$ is minimal. Then, the height of π^* is at most $|\Phi_V|$.*

Proof For simplicity, we denote below $\pi = \pi^*$. We prove the lemma by showing that π eliminates at least one realization from its version space in each query. This directly implies that the height of π is at most $|\Phi_V| - 1$.

We first assume that π is deterministic, and later extend the analysis to randomized policies. Assume for contradiction that there is a possible selection by π such that the resulting observation does not remove any realization from the version space. Formally, assume that there is some partial realization ψ , realization $\phi \sim \psi$ with some non-zero probability according to p , and an element $v \in V$, such that during the run of π under ϕ , the partial realization that is observed is ψ , $\pi(\psi) = v$, and for any realization $\phi' \sim \psi$ it holds that $\phi'(v) = \phi(v)$.

Let π' the policy defined as follows: For any partial realization ψ' :

$$\pi'(\psi') = \begin{cases} \pi(\psi' \cup \{(v, \phi(v))\}) & \psi \subseteq \psi', \\ \pi(\psi') & \text{otherwise.} \end{cases}$$

By the definition of π' , it holds that $c_{\text{avg}}(\pi', p) < c_{\text{avg}}(\pi, p)$. Thus, to reach a contradiction, it suffices to show that $f_p(E(\pi', \phi'), \phi') = 1$ for all $\phi' \in \Phi_V$.

Let $\phi' \in \Phi_V$. If ψ is not observed during the run of π' under ϕ' , then the behavior of π' is equivalent to that of π . Therefore,

$$f_p(E(\pi', \phi'), \phi') = f_p(E(\pi, \phi'), \phi') = 1.$$

Otherwise, ψ is observed during the run of π' under ϕ' . In this case, until ψ is observed, the behavior of π' under ϕ' is the same as its behavior under ϕ . Therefore, at some time during the run of π' , ψ is the partial realization that has been observed so far. By the definition of π' , it does not select v at this time, and then behaves the same as π . Hence, $E(\pi', \phi') = E(\pi, \phi') \setminus \{v\}$. For any deterministic policy $\bar{\pi}$, we define $\psi_{\bar{\pi}} = \{(v, \phi'(v)) \mid v \in E(\bar{\pi}, \phi')\}$ to be the partial realization that $\bar{\pi}$ observes under ϕ' just before it terminates. It holds that $\psi_{\pi} = \psi_{\pi'} \cup \{(v, \phi'(v))\}$.

Since ψ is observed during the run of π' under ϕ' , it holds that $\psi \subseteq \psi_{\pi'}$. Hence, any realization that is consistent with $\psi_{\pi'}$ is also consistent with ψ . We assumed that any realization that is consistent with ψ is also consistent with $\{(v, \phi'(v))\}$. Therefore, the same holds for $\psi_{\pi'}$. We conclude that any realization that is consistent with $\psi_{\pi'}$ is also consistent with ψ_{π} . Since $f_p(E(\pi, \phi'), \phi') = 1$, it holds that ϕ' is the only realization that is consistent with ψ_{π} . Hence, ϕ' is also the only realization that is consistent with $\psi_{\pi'}$. By the definition of f_p , it holds that $f_p(E(\pi', \phi'), \phi') = 1$. Thus, in all cases, $f_p(E(\pi', \phi'), \phi') = 1$ for all $\phi' \in \Phi_V$, as claimed. This shows the required contradiction.

Lastly, if π is not deterministic, then for any realization θ_{π} of the randomness of π , it holds that π given θ_{π} is deterministic. Assume for contradiction that there exists some realization θ_{π} of the randomness of π that has a non-zero probability, such that in some query, π under θ_{π} does not eliminate any realization from its version space. Since π under θ_{π} is a deterministic policy, similarly to the proof above, it follows that π under θ_{π} is not an optimal policy. Since θ_{π} has a positive probability, it follows that π is not an optimal policy, again resulting in a contradiction and proving the claim. \blacksquare