Minimum Cost Adaptive Submodular Cover

Yubing Cui* Viswanath Nagarajan†

Abstract

We consider the problem of minimum cost cover of adaptive-submodular functions, and provide a $4(\ln Q{+}1)$-approximation algorithm, where Q is the goal value. This bound is nearly the best possible as the problem does not admit any approximation ratio better than $\ln Q$ (unless $P = NP$). Our result is the first $O(\ln Q)$-approximation algorithm for this problem. Previously, $O(\ln Q)$-approximation algorithms were only known assuming either independent items or unit-cost items. Furthermore, our result easily extends to the setting where one wants to simultaneously cover *multiple* adaptive-submodular functions: we obtain the first approximation algorithm for this generalization.

1 Introduction

Adaptive stochastic optimization, where an algorithm makes sequential decisions while (partially) observing uncertainty, arises in numerous applications such as active learning [23], sensor placement [6] and viral marketing [13]. Often, these applications involve an underlying submodular function, and the framework of adaptive-submodularity (introduced by [9]) has been widely used to solve these problems. In this paper, we study a basic problem in this context: covering an adaptive-submodular function at the minimum expected cost. Our main result is an $O(\ln Q)$-approximation algorithm, where Q represents the maximal value of the adaptive-submodular function.

In some applications, such as sensor placement (or stochastic set cover [20]), the uncertainty just involves an *independent* random variable associated with each decision. However, there are also a number of applications where the random variables associated with different decisions are correlated. The adaptive-submodularity framework that we consider is also applicable in certain applications involving correlations.

As a motivating example, consider the viral marketing problem, where we are given a social network and target Q, and the goal is to influence at least Q users to adopt a new product. A user can be influenced in two ways (i) directly because the user is offered a promotion, or (ii) indirectly because some friend of the user was influenced *and* the friend influenced this user. We incur a cost only in case (i), which accounts for the promotional offer. A widely-used model for influence behavior is the *independent cascade model* [10]. Here, each arc (u, v) has a value $p_{uv} \in [0, 1]$ that represents the probability that user u will influence user v (if u is already influenced). A solution is a sequential process that in each step, selects one user w to influence directly, after which we get to observe which of w's friends were influenced (indirectly), and which of their friends were influenced, and so on. So, the solution can utilize these partial observations to make decisions *adaptively.* Such an adaptive solution can be represented by a decision tree; however, it may require exponential space to store explicitly. We will analyze simple solutions that can be implemented in polynomial time (and space), but our performance guarantees are relative to an optimal solution that can be very complex. Also, note that the random observations associated with different decisions (in the viral marketing problem) are highly correlated: the set of nodes that get (indirectly) influenced by any node w depends on the entire network (not just w).

1.1 Problem Definition

Random items. Let E be a finite set of n *items.* Each item $e \in E$ corresponds to a random variable $\Phi_e \in \Omega$, where Ω is the *outcome space* (for a single item). We use $\Phi = \langle \Phi_e : e \in E\rangle$ to denote the vector of all random variables (r.v.s). The r.v.s may be arbitrarily correlated across items. We use upper-case letters to represent r.v.s and the corresponding lower-case letters to represent realizations of the r.v.s. Thus, for any item e, $\phi_e \in \Omega$ is the

*Department of Mathematics, University of Michigan, Ann Arbor, USA.

†Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, USA. Research supported in part by NSF grants CMMI-1940766 and CCF-2006778.

Copyright © 2023 by SIAM
Unauthorized reproduction of this article is prohibited

realization of Φ_e; and $\phi = \langle \phi_e : e \in E \rangle$ denotes the realization of Φ. Equivalently, we can represent the realization ϕ as a subset $\{(e, \phi_e) : e \in E\} \subseteq E \times \Omega$ of item-outcome pairs.

A *partial realization* $\psi \subseteq E \times \Omega$ refers to the realizations of any *subset* of items; $\mathsf{dom}(\psi) \subseteq E$ denotes the items whose realizations are represented in ψ, and $\psi(e)$ denotes the realization of any item $e \in \mathsf{dom}(\psi)$. Note that a partial realization contains at most one pair of the form $(e, *)$ for any item $e \in E$. The (full) realization ϕ corresponds to a partial realization with $\mathsf{dom}(\phi) = E$. For two partial realizations $\psi, \psi' \subseteq E \times \Omega$, we say that ψ is a *subrealization* of ψ' (denoted $\psi \preccurlyeq \psi'$) if $\psi \subseteq \psi'$; in other words, $\mathsf{dom}(\psi) \subseteq \mathsf{dom}(\psi')$ and $\psi(e) = \psi'(e)$ for all $e \in \mathsf{dom}(\psi)$.[1] Two partial realizations $\psi, \psi' \subseteq E \times \Omega$ are said to be *disjoint* if there is no full realization ϕ with $\psi \preccurlyeq \phi$ and $\psi' \preccurlyeq \phi$; in other words, there is some item $e \in \mathsf{dom}(\psi) \cap \mathsf{dom}(\psi')$ such that the realization of Φ_e is different under ψ and ψ'.

We assume that there is a prior probability distribution $p(\phi) = \Pr[\Phi = \phi]$ over realizations ϕ. Moreover, for any partial realization ψ, we assume that we can compute the posterior distribution $p(\phi|\psi) = \Pr(\Phi = \phi|\psi \preccurlyeq \Phi)$.

Utility function. In addition to the random items (described above), there is a *utility function* $f : 2^{E\times\Omega} \to \mathbb{R}_{\geq 0}$ that assigns a value to any partial realization. We will assume that this function is monotone, i.e., having more realizations can not reduce the value. Formally,

Definition 1.1. (Monotonicity) *A function $f : 2^{E\times\Omega} \to \mathbb{R}_{\geq 0}$ is* monotone *if*

$$f(\psi) \leq f(\psi') \quad \text{for all partial realizations } \psi \preccurlyeq \psi'.$$

We also assume that the function f can always achieve its maximal value, i.e.,

Definition 1.2. (Coverable) *Let Q be the maximal value of function f. Then, function f is said to be* coverable *if this value Q can be achieved under every (full) realization, i.e.,*

$$f(\phi) = Q \text{ for all realizations } \phi \text{ of } \Phi.$$

Furthermore, we will assume that the function f along with the probability distribution $p(\cdot)$ satisfies a submodularity-like property. Before formalizing this, we need the following definition.

Definition 1.3. (Marginal benefit) *The* conditional expected marginal benefit *of an item $e \in E$ conditioned on observing the partial realization ψ is:*

$$\Delta(e|\psi) := \mathbb{E}\left[f(\psi \cup (e, \Phi_e)) - f(\psi) \,|\, \psi \preccurlyeq \Phi\right] \;=\; \sum_{\omega\in\Omega} \Pr[\Phi_e = \omega|\psi \preccurlyeq \Phi] \cdot \left(f(\psi \cup (e, \omega)) - f(\psi)\right).$$

We will assume that function f and distribution $p(\cdot)$ jointly satisfy the adaptive-submodularity property, defined as follows.

Definition 1.4. (Adaptive submodularity) *A function $f : 2^{E\times\Omega} \to \mathbb{R}_{\geq 0}$ is* adaptive submodular *w.r.t. distribution $p(\phi)$ if for all partial realizations $\psi \preccurlyeq \psi^{'}$, and all items $e \in E \setminus \mathsf{dom}(\psi^{'})$, we have*

$$\Delta(e|\psi) \geq \Delta(e|\psi^{'}).$$

In other words, this property ensures that the marginal benefit of an item never increases as we condition on more realizations.

Given any function f satisfying Definitions 1.1, 1.2 and 1.4, we can pre-process f by subtracting $f(\emptyset)$, to get an equivalent function (that maintains these properties), and has a smaller Q value.

Min-cost adaptive-submodular cover (ASC). In this problem, each item $e \in E$ has a positive cost c_e. The goal is to select items (and observe their realizations) sequentially until the observed realizations have function value Q. The objective is to minimize the expected cost of selected items.

Due to the stochastic nature of the problem, the solution concept here is much more complex than in the deterministic setting (where we just select a static subset). In particular, a solution corresponds to a "policy" that maps observed realizations to the next selection decision. The observed realization at any point corresponds

[1] We use the notation $\psi \preccurlyeq \psi'$ instead of $\psi \subseteq \psi'$ in order to be consistent with prior works.

Copyright © 2023 by SIAM
Unauthorized reproduction of this article is prohibited

to a partial realization (namely, the realizations of the items selected so far). Formally, a *policy* is a mapping $\pi : 2^{E\times\Omega} \to E$, which specifies the next item $\pi(\psi)$ to select when the observed realizations are ψ.[2] The policy π terminates at the first point when $f(\psi) = Q$, where $\psi \subseteq E \times \Omega$ denotes the observed realizations so far. The cost of policy π, denoted $c_{exp}(\pi)$, is the expected cost of all selected items until π terminates.

At any point in policy π, we refer to the cumulative cost incurred so far as the *time*. If $J_1, J_2, \cdots J_k$ denotes the (random) sequence of items selected by π then for each $i \in \{1, 2, \cdots k\}$, we view item J_i as being selected during the time interval $[\sum_{h=1}^{i-1} c(J_h)\,, \sum_{h=1}^{i} c(J_h))$ and the realization of J_i is only observed at time $\sum_{h=1}^{i} c(J_h)$. For any time $t \geq 0$, we use $\Psi(\pi, t) \subseteq E \times \Omega$ to denote the (random) realizations that have been observed by time t in policy π. We note that $\Psi(\pi, t)$ only contains the realizations of items that have been *completely* selected by time t. Note that the policy terminates at the earliest time t where $f(\Psi(\pi, t)) = Q$.

Given any policy π, we define its *cost k truncation* by running π and stopping it just before the cost of selected items exceeds k. That is, we stop the policy as late as possible while ensuring that the cost of selected items never exceeds k (for any realization).

Remark: Our definition of the utility function f is slightly more restrictive than the original definition [9]. In particular, the utility function in [9] is of the form $g : 2^E \times \Omega^E \to \mathbb{R}_{\geq 0}$, where the function value $g(\mathsf{dom}(\psi), \Phi)$ for any partial realization ψ is still random and can depend on the outcomes of unobserved items, i.e., those in $E \setminus \mathsf{dom}(\psi)$. Nevertheless, our formulation (ASC) still captures most applications of the formulation studied in [9]. See Section 3 for details.

1.2 Adaptive Greedy Policy

Algorithm 1 describes a natural greedy policy for min-cost adaptive-submodular cover, which has also been studied in prior works [8, 14, 17].

Algorithm 1 Adaptive Greedy Policy π.

1: selected items $A \leftarrow \emptyset$, observed realizations $\psi \leftarrow \emptyset$
2: **while** $f(\psi) < Q$ **do**
3: $\quad e^* = \arg\max_{e\in E\setminus A} \frac{\Delta(e|\psi)}{c_e}$
4: $\quad$ add e^* to the selected items, i.e., $A \leftarrow A \cup \{e^*\}$
5: $\quad$ select e^* and observe Φ_{e^*}
6: $\quad$ update $\psi \leftarrow \psi \cup \{(e^*, \Phi_{e^*})\}$

Remark: Note that the policy π remains the same if we replace the greedy choice by

$$e^* = \arg\max_{e\in E\setminus A} \frac{\Delta(e|\psi)}{c_e \cdot (Q - f(\psi))}. \tag{1.1}$$

This is because the additional term $Q - f(\psi)$ is the same for each item $e \in E \setminus A$ (note that at any particular step, ψ is a fixed partial realization). We will make use of this alternative greedy criterion in our analysis.

1.3 Our Contributions

Our main result is the following.

THEOREM 1.1. *Consider any instance of minimum cost adaptive-submodular cover, where the utility function $f : 2^{E\times\Omega} \to \mathbb{R}_{\geq 0}$ is monotone, coverable and adaptive-submodular w.r.t. the probability distribution $p(\cdot)$. Suppose that there is some value $\eta > 0$ such that $f(\psi) > Q - \eta$ implies $f(\psi) = Q$ for all partial realizations $\psi \subseteq E \times \Omega$. Then, the cost of the greedy policy is*

$$c_{exp}(\pi) \leq 4 \cdot (1 + \ln(Q/\eta)) \cdot c_{exp}(\sigma),$$

where σ denotes the optimal policy.

This is an asymptotic improvement over the $(1 + \ln(Q/\eta))^2$-approximation bound from [8] and the $(1 + \ln(\frac{nQc_{max}}{\eta}))$-approximation bound from [14]; the maximum item cost c_{max} can even be exponentially larger

[2] Policies and utility functions are not necessarily well-defined over all subsets $2^{E\times\Omega}$, but only over partial realizations; recall that a partial realization is of the form $\{(e, \phi_e) : e \in S\}$ where ϕ is some full-realization and $S \subseteq E$.

Copyright © 2023 by SIAM
Unauthorized reproduction of this article is prohibited

than Q. Our bound is the best possible (up to the constant factor of 4) because the set cover problem is a special case [15].

As a consequence, we obtain the first $O(\ln Q)$-approximation algorithm for the viral marketing application mentioned earlier. We also obtain an improved bound for the optimal decision tree problem with uniform priors. See Section 3 for details.

Our proof technique is very different from prior works [8, 14, 17]. We proceed by relating the non-completion probabilities in the optimal policy (at any time t) to that in the greedy policy (at a scaled time $\alpha \cdot t$). This suffices to bound the ASC objective because the expected cost of any policy is the integral of the non-completion probabilities over all times. In order to establish this relation, we consider the total value of the greedy criterion over a suitable time interval (called a "phase") and prove lower and upper bounds on this quantity. This high-level analysis has been used earlier for a number of stochastic covering problems, including the independent special case of ASC [24]. A key simplification/improvement is that, unlike prior work, we do not rely on non-completion probabilities at only power-of-two time points. We also make use of a stronger upper bound that combines multiple phases.

Moreover, our algorithm and analysis extend in a straightforward manner, to the setting with multiple adaptive-submodular functions, where the objective is the expected sum of "cover times" of the functions. We obtain the same approximation ratio even for this more general problem. Previous techniques [8, 14] do not seem to extend to this setting. In fact, the multiple ASC problem (with $Q = \eta = 1$) generalizes the min-sum set cover problem [25], which is NP-hard to approximate better than factor 4. So, our constant factor of 4 seems to be the best possible for this generalization.

1.4 Related Work

Adaptive submodularity was introduced by [9], where they considered both the maximum-coverage and the minimum-cost-cover problems. They showed that the greedy policy is a $(1 - \frac{1}{e})$ approximation for maximum coverage, where the goal is to maximize the expected value of an adaptive-submodular function subject to a cardinality constraint. They also claimed that the greedy policy is a $(1 + \ln(Q/\eta))$ approximation for min-cost cover of an adaptive-submodular function. However, this result had an error [12], and a corrected proof [8] only provides a double-logarithmic $(1 + \ln(Q/\eta))^2$ approximation. Recently, [14] obtained a single-logarithmic approximation bound of $(1 + \ln(\frac{nQc_{max}}{\eta}))$. However, this bound depends additionally on the number of items n and their maximum cost c_{max}. Our result shows that the greedy policy is indeed an $O(\ln(Q/\eta))$ approximation. As noted earlier, our definition of ASC is simpler and slightly more restrictive than the original one in [9], although most applications of adaptive-submodularity do satisfy our definition.

The special case of adaptive-submodularity where the random variables are independent across items, has also been studied extensively. For the maximum-coverage version, [1] obtained a $(1 - \frac{1}{e})$-approximation algorithm via a "non adaptive" policy (that fixes a subset of items to select upfront). Subsequent work [5, 7, 19] obtained constant factor approximation algorithms for a variety of constraints (beyond just cardinality). The minimum-cost cover problem (called *stochastic submodular cover*) was studied in [2, 17, 18, 22, 24]. In particular, an $O(\ln(Q/\eta))$ approximation algorithm follows from [24], and recently [17] proved that the greedy policy has a $(1 + \ln(Q/\eta))$ approximation guarantee. The latter guarantee is the best possible, even up to the constant factor: this also matches the best approximation ratio for the *deterministic* submodular cover problem [16].

The stochastic submodular cover problem with multiple submodular functions was studied in [24], for which an $O(\ln(Q/\eta))$ approximation algorithm was obtained. The analysis in this paper is similar, at a high level, to the analysis in [24]. However, we handle a more general setting (where items may be correlated), and we obtain a much better constant factor.

A different (scenario based) model for correlations in adaptive submodular cover was studied in [11, 21]. Here, the utility function f is just required to be submodular (not adaptive-submodular), but the algorithm requires an explicit description of the probability distribution $p(\cdot)$. In particular, [11] obtained a greedy-style policy with approximation ratio $O(\ln(mQ/\eta))$ where m is the support-size of distribution $p(\cdot)$, and Q and η are as before. Our proof technique (using non-completion probabilities at all times) can also be combined with [11] to improve the constant factor in their approximation ratio.

2 Proof of the Main Result

Let $L := 1 + \ln(Q/\eta)$ and $\beta > 1$ be some constant value (fixed later). Our analysis is based on relating the "non completion" probabilities at different times in the greedy policy π and the optimal policy σ. We first define these

Copyright © 2023 by SIAM
Unauthorized reproduction of this article is prohibited

quantities formally.

Definition 2.1. (Non-completion probabilities) *For any time $t \geq 0$, let*

$$o(t) := \Pr[\sigma \text{ does not terminate by time } t] = \Pr[f(\Psi(\sigma, t)) < Q].$$

Similarly, for any $t \geq 0$, let

$$a(t) := \Pr[\pi \text{ does not terminate by time } \beta L \cdot t] = \Pr[f(\Psi(\pi, \beta L t)) < Q].$$

See Figure 1 for an example of function $o(t)$. Notice that the function $a(t)$ corresponding to the greedy policy is defined w.r.t. *scaled* times $\beta L \cdot t$ rather than t.

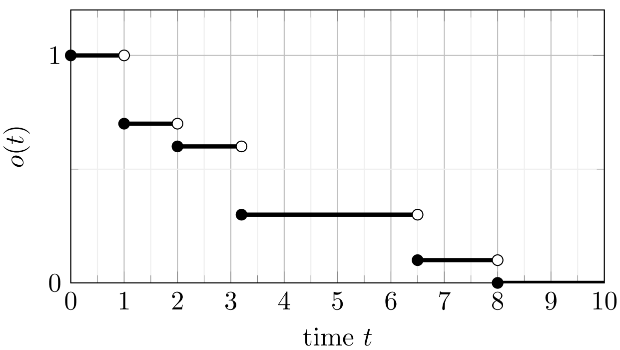

Figure 1: Graph of a simple $o(\cdot)$ function.

Clearly, $o(t)$ and $a(t)$ are non-increasing functions of t. Moreover, $o(0) = a(0) = 1$ and there exists some time t_0 such that $o(t) = a(t) = 0$ for all $t \geq t_0$. We can also express the expected policy costs as follows:

$$c_{exp}(\sigma) = \int_0^\infty \Pr[\sigma \text{ does not terminate by time } t]dt = \int_0^\infty o(t)dt.$$

$$c_{exp}(\pi) = \int_0^\infty \Pr[\pi \text{ does not terminate by time } y]dy = \int_0^\infty a(\frac{y}{\beta L})dy = \beta L \cdot \int_0^\infty a(t)dt.$$

We approximate the cost of π by considering only the "integral" time points $\{\beta L i : i \in \mathbb{Z}_{\geq 0}\}$. Define $A := \sum_{i \geq 0} a(i)$ to be the sum of non-completion probabilities at these time points.

$$c_{exp}(\pi) = \beta L \int_0^\infty a(t)dt \leq \beta L \sum_{i \geq 0} a(i) = \beta L \cdot A, \tag{2.2}$$

where the inequality uses the fact that $a(\cdot)$ is non-increasing.

Our analysis of π relies on tracking the "greedy criterion value" defined in (1.1). The following definition formalizes this.

Definition 2.2. (Greedy score) *For $t \geq 0$ and any partial realization ψ at time t, define*

$$score(t, \psi) := \begin{cases} \frac{\Delta(e|\psi)}{c_e[Q - f(\psi)]}, & \text{where } e \text{ is the item being selected in } \pi \text{ at time } t \text{ when } \psi \text{ was observed just before selecting } e. \\ 0, & \text{if no item is being selected in } \pi \text{ at time } t \text{ when } \psi \text{ was observed.} \end{cases}$$

Note that conditioned on ψ, the item e being selected in π at time t is deterministic.

The expression for *score* above is exactly the greedy criterion in (1.1). Moreover, the score may increase and decrease over time: see Figure 2 for an example.

In order to reduce notation, for any time t, we use $\Psi_t := \Psi(\pi, t)$ to denote the (random) partial realization observed by the greedy policy π at time t; recall that this only includes items that have been completely selected by time t.

Copyright © 2023 by SIAM
Unauthorized reproduction of this article is prohibited

Definition 2.3. (Gain of a phase) *For any integer $i \geq 0$, the time interval $[\beta Li, \beta L(i+1))$ is called* phase i. *For any phase $i \geq 0$, its* gain *is the expected total score accumulated during phase i,*

$$G_i := \int_{L\beta i}^{L\beta(i+1)} \mathbb{E}\left[score(t, \Psi_t)\right] dt,$$

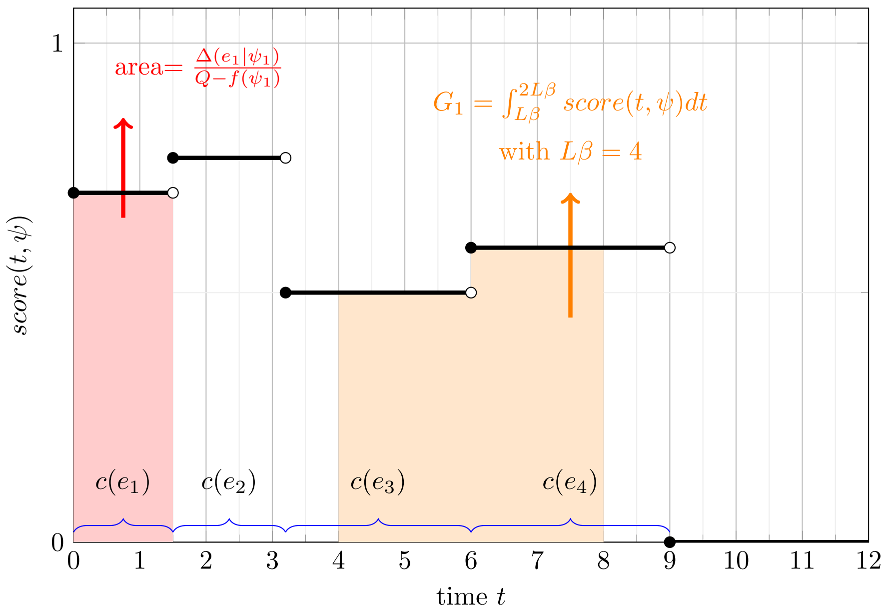

Figure 2: Graph of a simple $score(t, \psi)$ for illustration. $e_1, e_2, ...$ are greedy selections and ψ_i is the partial realization just before selecting e_i.

The key part of the analysis lies in upper and lower bounding the gains in all the phases.

Lemma 2.1. *For any $i \geq 0$, the total gain after phase i is*

$$\sum_{j \geq i} G_j \leq L \cdot a(i).$$

Proof. We start by re-expressing the score and gain in terms of the full realization. For any time $t \geq 0$ and full realization ϕ, let

$$S(t, \phi) := \begin{cases} \frac{f(\psi \cup (e, \phi_e)) - f(\psi)}{c_e[Q - f(\psi)]}, & \text{where } e \text{ is the item being selected in } \pi \text{ at time } t \text{ under } \phi, \\ & \text{and } \psi \preccurlyeq \phi \text{ is the partial realization just before selecting } e. \\ 0, & \text{if no item is being selected in } \pi \text{ at time } t \text{ under } \phi. \end{cases}$$

Then, for any $t \geq 0$ and any partial realization ψ at time t,

$$score(t, \psi) = \mathbb{E}_\Phi[S(t, \Phi) | \psi \preccurlyeq \Phi].$$

This uses the definition of $\Delta(e|\psi)$ and the fact that conditioned on ψ, the item e (being selected at time t) is fixed. Hence, for any phase $k \geq 0$, its gain

$$G_k = \int_{L\beta k}^{L\beta(k+1)} \mathbb{E}\left[score(t, \Psi_t)\right] dt = \int_{L\beta k}^{L\beta(k+1)} \mathbb{E}_{\Psi_t}\left[\mathbb{E}_\Phi[S(t, \Phi) | \Psi_t \preccurlyeq \Phi]\right] dt = \int_{L\beta k}^{L\beta(k+1)} \mathbb{E}\left[S(t, \Phi)\right] dt.$$

Now, fix phase $i \geq 0$ and any (full) realization ϕ.

Copyright © 2023 by SIAM
Unauthorized reproduction of this article is prohibited

Case 1: suppose that π under ϕ terminates before ($\leq$) time $L\beta i$. Then, $S(t,\phi)=0$ for all $t > L\beta i$, and so:

$$\sum_{j\geq i} G_j(\phi) = \sum_{j\geq i} \int_{L\beta j}^{L\beta(j+1)} S(t,\phi)dt = 0$$

Case 2: suppose that π under ϕ terminates after ($>$) time $L\beta i$.

$$\sum_{j\geq i} G_j(\phi) = \sum_{j\geq i} \int_{L\beta j}^{L\beta(j+1)} S(t,\phi)dt = \int_{L\beta i}^{\infty} S(t,\phi)dt \leq \int_{0}^{\infty} S(t,\phi)dt \leq L,$$

where the last inequality uses Lemma A.1 (proved in Appendix A).

Note that case 2 above happens exactly with probability $a(i)$. So,

$$\sum_{j\geq i} G_j = \mathbb{E}_\Phi\left[\sum_{j\geq i} G_j(\Phi) \,|\, \text{case 2 occurs under } \Phi\right] \cdot \Pr[\text{case 2 occurs}] \leq L\cdot a(i),$$

which completes the proof. ∎

LEMMA 2.2. *For any $i \geq 0$ and time $t \in [L\beta i, L\beta(i+1))$,*

$$\mathbb{E}[score(t,\Psi_t)] \geq \frac{a(i+1)-o(i+1)}{i+1}.$$

Hence, $G_i \geq \beta L \cdot \left(\frac{a(i+1)-o(i+1)}{i+1}\right)$ for each phase $i \geq 0$.

Proof. Note that the first statement in the lemma immediately implies the second statement. Indeed,

$$G_i = \int_{L\beta i}^{L\beta(i+1)} \mathbb{E}[score(t,\Psi_t)]dt \geq \beta L \cdot \left(\frac{a(i+1)-o(i+1)}{i+1}\right).$$

We now prove the first statement. Henceforth, fix phase $i \geq 0$ and time $t \in [L\beta i, L\beta(i+1))$.

Truncated optimal policy Let $\overline{\sigma}$ denote the cost $i+1$ truncation of policy σ. Note that the total cost of selected items in $\overline{\sigma}$ is always at most $i+1$. However, $\overline{\sigma}$ may not fully cover the utility function f (so it is not a feasible policy for min-cost adaptive submodular cover). We define the following random quantities associated with policy $\overline{\sigma}$:

$I_k :=$ set of first k items selected by $\overline{\sigma}$, for $k = 0, 1, \cdots$.
$I_\infty :=$ set of *all* items selected by the end of $\overline{\sigma}$.
$P_k := \{(e,\Phi_e) : e \in I_k\}$, i.e. partial realization of the first k items selected by $\overline{\sigma}$, for $k = 0, 1, \cdots$.
$P_\infty := \{(e,\Phi_e) : e \in I_\infty\}$, i.e. partial realization observed by the end of $\overline{\sigma}$.

Note that $\overline{\sigma}$ covers f exactly when $f(P_\infty) = Q$. Moreover, by definition of the function $o(\cdot)$, we have $\Pr[\overline{\sigma} \text{ covers } f] = 1 - o(i+1)$.

Conditioning on partial realizations in greedy Let ψ be any partial realization corresponding to Ψ_t with $f(\psi) < Q$. In other words, (i) ψ is the partial realization observed at time t in some execution of policy π, and (ii) the policy has not terminated (under realization ψ) by time t. Let $R(\pi,t)$ denote the collection of such partial realizations. Note that the partial realizations in $R(\pi,t)$ are mutually disjoint, and the total probability of these partial realizations equals the probability that π does not terminate by time t. We will show that:

$$\Pr[\psi \preccurlyeq \Phi]\cdot score(t,\psi) \geq \frac{1}{i+1}\cdot \Pr[(\psi \preccurlyeq \Phi) \wedge (\overline{\sigma} \text{ covers } f)], \qquad \forall \psi \in R(\pi,t). \tag{2.3}$$

We first complete the proof of the lemma assuming (2.3).

Copyright © 2023 by SIAM
Unauthorized reproduction of this article is prohibited

$$
\begin{aligned}
\mathbb{E}[score(t,\Psi_t)] \quad &\geq \quad \sum_{\psi\in R(\pi,t)} p(\psi)\cdot score(t,\psi) = \sum_{\psi\in R(\pi,t)} \Pr[\psi \preccurlyeq \Phi]\cdot score(t,\psi) \\
&\geq \quad \frac{1}{i+1}\sum_{\psi\in R(\pi,t)} \Pr[(\psi \preccurlyeq \Phi)\wedge(\overline{\sigma} \text{ covers } f)] \qquad (2.4)\\
&= \quad \frac{1}{i+1}\Pr\left[(\pi \text{ doesn't terminate by time } t)\wedge(\overline{\sigma} \text{ covers } f)\right] \qquad (2.5)\\
&\geq \quad \frac{1}{i+1}\left(\Pr[\pi \text{ doesn't terminate by time } t] - \Pr[\overline{\sigma} \text{ does not cover } f]\right) \qquad (2.6)\\
&= \quad \frac{1}{i+1}\cdot\left(a\left(\frac{t}{L\beta}\right) - o(i+1)\right) \qquad (2.7)\\
&\geq \quad \frac{a(i+1)-o(i+1)}{i+1}. \qquad (2.8)
\end{aligned}
$$

Inequality (2.4) is by (2.3). The equality in (2.5) uses the definition of $R(\pi,t)$. Inequality (2.6) is by a union bound. Equation (2.7) is by definition of the functions $a(\cdot)$ and $o(\cdot)$. Finally, (2.8) uses $t < \beta L(i+1)$ and that $a(\cdot)$ is non-increasing.

Proof of (2.3) Henceforth, fix any partial realization $\psi \in R(\pi,t)$. Our proof relies on the following quantity:

$$
Z := \mathbb{E}_\Phi\left[\mathbb{1}(\psi \preccurlyeq \Phi)\cdot\frac{f(\psi\cup P_\infty)-f(\psi)}{Q-f(\psi)}\right] \qquad (2.9)
$$

In other words, this is the expected increase in policy $\overline{\sigma}$'s function value (relative to the "remaining" target $Q-f(\psi)$) when restricted to (full) realizations Φ that agree with partial realization ψ.

For any partial realization ψ' such that $\psi \preccurlyeq \psi'$ and item $e \notin \mathsf{dom}(\psi')$, let $X_{e,\psi,\psi'}$ denote the indicator r.v. that policy $\overline{\sigma}$ selects item e at some point when its observed realizations are precisely $(\psi'\setminus\psi)\cup\chi$ where $\chi\subseteq\psi$. That is, $X_{e,\psi,\psi'}=1$ if policy $\overline{\sigma}$ selects e at a point where (i) all items in $\mathsf{dom}(\psi'\setminus\psi)$ have been selected and their realization is $\psi'\setminus\psi$, (ii) no item in $E\setminus\mathsf{dom}(\psi')$ has been selected, and (iii) if any item in $\mathsf{dom}(\psi)$ has been selected then its realization agrees with ψ. Note that conditioned on $\psi'\preccurlyeq\Phi$, $X_{e,\psi,\psi'}$ is a deterministic value: the realizations of items $\mathsf{dom}(\psi')$ are fixed by ψ' and if any item in $E\setminus\mathsf{dom}(\psi')$ is selected then $X_{e,\psi,\psi'}=0$ irrespective of its realization.

We can write Z as a sum of increments as follows:

$$
\begin{aligned}
Z &= \frac{1}{Q-f(\psi)}\mathbb{E}_\Phi\left[\mathbb{1}(\psi\preccurlyeq\Phi)\cdot\sum_{k\geq 1}[f(\psi\cup P_k)-f(\psi\cup P_{k-1})]\right]\\
&= \frac{1}{Q-f(\psi)}\sum_{\psi':\psi\preccurlyeq\psi'}\ \sum_{e\notin\mathsf{dom}(\psi')}\mathbb{E}_\Phi\left[\mathbb{1}(\psi'\preccurlyeq\Phi)\cdot X_{e,\psi,\psi'}\cdot[f(\psi'\cup(e,\Phi_e))-f(\psi')]\right]\\
&= \frac{1}{Q-f(\psi)}\sum_{\psi':\psi\preccurlyeq\psi'}\ \sum_{e\notin\mathsf{dom}(\psi')}\Pr[\psi'\preccurlyeq\Phi\wedge X_{e,\psi,\psi'}=1]\cdot\mathbb{E}_\Phi[f(\psi'\cup(e,\Phi_e))-f(\psi')\,|\,(\psi'\preccurlyeq\Phi)\wedge(X_{e,\psi,\psi'}=1)]\\
&= \frac{1}{Q-f(\psi)}\sum_{\psi':\psi\preccurlyeq\psi'}\ \sum_{e\notin\mathsf{dom}(\psi')}\Pr[\psi'\preccurlyeq\Phi\wedge X_{e,\psi,\psi'}=1]\cdot\mathbb{E}_\Phi[f(\psi'\cup(e,\Phi_e))-f(\psi')\,|\,\psi'\preccurlyeq\Phi] \qquad (2.10)\\
&= \frac{1}{Q-f(\psi)}\sum_{\psi':\psi\preccurlyeq\psi'}\ \sum_{e\notin\mathsf{dom}(\psi')}\Pr[\psi'\preccurlyeq\Phi\wedge X_{e,\psi,\psi'}=1]\cdot\Delta(e|\psi')
\end{aligned}
$$

Copyright © 2023 by SIAM
Unauthorized reproduction of this article is prohibited

$$\le \frac{1}{Q - f(\psi)} \sum_{\psi':\psi\preccurlyeq\psi'} \sum_{e\notin \mathsf{dom}(\psi')} \Pr[\psi' \preccurlyeq \Phi \wedge X_{e,\psi,\psi'} = 1] \cdot \Delta(e|\psi) \tag{2.11}$$

$$= \sum_{\psi':\psi\preccurlyeq\psi'} \sum_{e\notin \mathsf{dom}(\psi')} \Pr[\psi' \preccurlyeq \Phi \wedge X_{e,\psi,\psi'} = 1] \cdot c_e \cdot \frac{\Delta(e|\psi)}{c_e(Q - f(\psi))}$$

$$\le \sum_{\psi':\psi\preccurlyeq\psi'} \sum_{e\notin \mathsf{dom}(\psi')} \Pr[\psi' \preccurlyeq \Phi \wedge X_{e,\psi,\psi'} = 1] \cdot c_e \cdot \mathit{score}(t,\psi) \tag{2.12}$$

$$= \mathit{score}(t,\psi) \sum_{\psi':\psi\preccurlyeq\psi'} \sum_{e\notin \mathsf{dom}(\psi')} c_e \cdot \Pr[\psi' \preccurlyeq \Phi \wedge X_{e,\psi,\psi'} = 1]$$

$$= \mathit{score}(t,\psi) \cdot \sum_{e\in E\setminus \mathsf{dom}(\psi)} c_e \cdot \sum_{\substack{\psi':\psi\preccurlyeq\psi' \\ \mathsf{dom}(\psi')\not\ni e}} \Pr[\psi' \preccurlyeq \Phi \wedge X_{e,\psi,\psi'} = 1]$$

$$= \mathit{score}(t,\psi) \cdot \sum_{e\in E\setminus \mathsf{dom}(\psi)} c_e \cdot \Pr[\psi \preccurlyeq \Phi \wedge e \in I_\infty] \tag{2.13}$$

$$\le \mathit{score}(t,\psi) \cdot \mathbb{E}_\Phi \left[\mathbb{1}(\psi \preccurlyeq \Phi) \cdot \sum_{e\in I_\infty} c_e \right]$$

$$\le \mathit{score}(t,\psi) \cdot (i+1) \cdot \mathbb{E}_\Phi[\mathbb{1}(\psi \preccurlyeq \Phi)] \;=\; \mathit{score}(t,\psi) \cdot (i+1) \cdot \Pr[\psi \preccurlyeq \Phi]. \tag{2.14}$$

The equality (2.10) uses the fact that $X_{e,\psi,\psi'}$ is deterministic when conditioned on $\psi' \preccurlyeq \Phi$. Inequality (2.11) is by adaptive submodularity. (2.12) is by the greedy selection criterion. The inequality in (2.14) uses the fact that the total cost of $\overline{\sigma}$'s selections is always bounded above by $i+1$. Equation (2.13) uses the definition of I_∞ (all selected items in $\overline{\sigma}$) and the following identity:

$$\sum_{\substack{\psi':\psi\preccurlyeq\psi' \\ \mathsf{dom}(\psi')\not\ni e}} \mathbb{1}(\psi' \preccurlyeq \Phi) \cdot X_{e,\psi,\psi'} = \mathbb{1}(\psi \preccurlyeq \Phi \wedge e \in I_\infty), \qquad \forall e \in E \setminus \mathsf{dom}(\psi).$$

To see this, condition on any full realization ϕ. If $\psi \not\preccurlyeq \phi$ then both the left-hand-side (LHS) and right-hand-side (RHS) are 0. If $\psi \preccurlyeq \phi$ and e is not selected by $\overline{\sigma}$ under ϕ, then again $LHS = RHS = 0$. If $\psi \preccurlyeq \phi$ and e is selected by $\overline{\sigma}$ under ϕ, then $RHS = 1$ and LHS is the sum of $X_{e,\psi,\psi'}$ over ψ' such that $\psi \preccurlyeq \psi' \preccurlyeq \phi$ and $e \notin \mathsf{dom}(\psi')$. In this case, $X_{e,\psi,\psi'} = 1$ for exactly one such partial realization ψ', namely $\psi' = \psi \cup \kappa$ where $\kappa \preccurlyeq \phi$ is the partial realization immediately before e is selected. So, $LHS = RHS$ in all cases.

Note that whenever $\overline{\sigma}$ covers f, we have $f(P_\infty) = Q$. Combined with the monotone property of f, we have $f(\psi \cup P_\infty) = Q$ whenever $\overline{\sigma}$ covers f. So, we have:

$$Z \ge \Pr[(\psi \preccurlyeq \Phi) \wedge (\overline{\sigma} \text{ covers } f)].$$

Combining the above inequality with (2.14) finishes the proof of (2.3). ◻

Completing the proof of Theorem 1.1 Using Lemma 2.1 and Lemma 2.2, we get:

$$a(i)L \ge \sum_{j\ge i} G_j \ge L\beta \sum_{j\ge i} \frac{a(j+1) - o(j+1)}{j+1}, \qquad \forall i \ge 0 \tag{2.15}$$

Recall that $A = \sum_{i\ge 0} a(i)$. Similarly, define $O := \sum_{i\ge 0} o(i)$.
Using (2.15) and adding over all $i \ge 0$,

$$\frac{1}{\beta} A = \frac{1}{\beta} \sum_{i\ge 0} a(i) \ge \sum_{i\ge 0} \sum_{j\ge i} \frac{a(j+1) - o(j+1)}{j+1} = \sum_{j\ge 0} (a(j+1) - o(j+1)) = A - O.$$

The last equality uses $a(0) = 1 = o(0)$. It now follows that $A \le \frac{\beta}{\beta-1} O$.

Copyright © 2023 by SIAM
Unauthorized reproduction of this article is prohibited

Using $o(0) = 1$ and the non-increasing property of $o(\cdot)$, we know that

$$O - 1 = \sum_{i \geq 0} o(i+1) \leq c_{exp}(\sigma) = \int_0^\infty o(t)dt \leq \sum_{i \geq 0} o(i) = O$$

Finally, combined with (2.2), we have

$$c_{exp}(\pi) \leq L\beta A \leq \frac{L\beta^2}{\beta - 1} O \leq \frac{L\beta^2}{\beta - 1}(c_{exp}(\sigma) + 1).$$

In order to optimize the constant factor, we set $\beta = 2$, which implies:

$$c_{exp}(\pi) \ \leq \ 4L \cdot (c_{exp}(\sigma) + 1) \tag{2.16}$$

We now show that the extra +1 term can be eliminated by a scaling argument: this would complete the proof of Theorem 1.1. Let $b \geq 1$ be some parameter, and consider the ASC instance with scaled costs $b \cdot c_e$. Notice that the greedy policy remains the same, and its objective just scales up by b. Similarly, the optimal cost also scales up by b. So, using (2.16) on the scaled ASC instance, we get:

$$b \cdot c_{exp}(\pi) \leq 4L \cdot (b \cdot c_{exp}(\sigma) + 1) \implies c_{exp}(\pi) \leq 4L \cdot (c_{exp}(\sigma) + \frac{1}{b}).$$

Thus, by taking b to be arbitrarily large, we obtain $c_{exp}(\pi) \leq 4L \cdot c_{exp}(\sigma)$.

3 Applications

Here, we provide some concrete applications of our framework. These applications were already discussed in [8], but as noted in Section 1.1, the function definition in ASC is slightly more restrictive than the framework in [8].

Stochastic Submodular Cover. In this problem, there are n stochastic items (for example, corresponding to sensors). Each item e can be in one of many "states", and this state is observed only after selecting item e. E.g., the state of a sensor indicates the extent to which it is working. The states of different items are independent. There is a utility function $\hat{f} : 2^{E \times \Omega} \to \mathbb{R}_{\geq 0}$, where E is the set of items and Ω the set of states. It is assumed that $\hat{f}$ is monotone and submodular. For e.g., $\hat{f}$ quantifies the information gained from a set of sensors having arbitrary states. Each item e is also associated with a cost c_e. Given a quota Q, the goal is to select items sequentially to achieve utility at least Q, at the minimum expected cost. We assume that the quota Q can always be achieved by selecting adequately many items, i.e., $\hat{f}(\{(e, \phi_e) : e \in E\}) \geq Q$ for all possible states $\{\phi_e \in \Omega\}_{e \in E}$ for the items. This is a special case of ASC, where the items E and states (outcomes) Ω remain the same. We define a new utility function $f(\psi) = \min\left\{\hat{f}(\psi), Q\right\}$ for all $\psi \subseteq E \times \Omega$. Note that Q is the maximal value of function f and this value is achieved under every possible (full) realization. Moreover, f is also monotone and submodular. Clearly, the monotonicity property (Definition 1.1) holds. The adaptive-submodularity property also holds because the items are independent. Indeed, for any partial realizations $\psi \preccurlyeq \psi'$ and $e \in E \setminus \mathsf{dom}(\psi')$,

$$\begin{aligned}
\Delta(e|\psi) &= \sum_{\omega \in \Omega} \Pr[\Phi_e = \omega | \psi \preccurlyeq \Phi] \cdot (f(\psi \cup (e,\omega)) - f(\psi)) \\
&= \sum_{\omega \in \Omega} \Pr[\Phi_e = \omega] \cdot (f(\psi \cup (e,\omega)) - f(\psi)) \\
&\geq \sum_{\omega \in \Omega} \Pr[\Phi_e = \omega] \cdot (f(\psi' \cup (e,\omega)) - f(\psi')) \\
&= \sum_{\omega \in \Omega} \Pr[\Phi_e = \omega | \psi' \preccurlyeq \Phi] \cdot (f(\psi' \cup (e,\omega)) - f(\psi')) = \Delta(e|\psi').
\end{aligned}$$

In particular, when $\hat{f}$ is integer-valued, Theorem 1.1 implies a $4(1 + \ln Q)$-approximation algorithm. We note that [17] obtained a $(1 + \ln Q)$-approximation ratio, using a different analysis. The latter bound is the best possible, including the constant factor, as the problem generalizes set cover. Our analysis is simpler and also extends to the more general adaptive submodular setting.

Copyright © 2023 by SIAM
Unauthorized reproduction of this article is prohibited

Adaptive Viral Marketing. This problem is defined on a directed graph $G = (V, A)$ representing a social network [10]. Each node $v \in V$ represents a user. Each arc $(u, v) \in A$ is associated with a random variable $X_{uv} \in \{0, 1\}$. The r.v. $X_{uv} = 1$ if u will influence v (assuming u itself is influenced); we also say that arc (u, v) is *active* in this case. The r.v.s X_{uv} are independent, and we are given the means $\mathbb{E}[X_{uv}] = p_{uv}$ for all $(u, v) \in A$. When a node u is activated/influenced, all arcs (u, v) out of u are observed and if $X_{uv} = 1$ then v is also activated. This process then continues on u's neighbors to their neighbors and so on, until no new node is activated. We consider the "full feedback" model, where after activating a node w, we observe the X_{uv} r.v.s on all arcs (u, v) such that u is reachable from w via a path of active arcs. Further, each node v has a cost c_v corresponding to activating node v *directly*, e.g. by providing some promotional offer. Note that there is no cost incurred on v if it is activated (indirectly) due to a neighbor u with $X_{uv} = 1$. Given a quota Q, the goal is to activate at least Q nodes at the minimum expected cost.

To model this as ASC, the items $E = V$ are all nodes in G. We add self-loops $A_o = \{(v, v) : v \in V\}$ that represent whether a node is activated directly. So, the new set of arcs is $A' = A \cup A_o$. The outcome Φ_w of any node $w \in V$ is represented by a function $\phi_w : A' \to \{0, 1, ?\}$ where $\phi_w((w, w)) = 1$, $\phi_w((v, v)) = 0$ for all $v \in V \setminus w$, and for any $(u, v) \in A$:

- $\phi_w(u, v) = 1$ if there is a $w - u$ path of active arcs and $X_{uv} = 1$ (i.e., (u, v) is active).
- $\phi_w(u, v) = 0$ if there is a $w - u$ path of active arcs and $X_{uv} = 0$ (i.e., (u, v) is not active).
- $\phi_w(u, v) =?$ if there is no $w - u$ path of active arcs (so, the status of (u, v) is unknown).

Let Ω denote the collection of all such functions: this represents the outcome space. Note that Φ_w depends on the entire network (and not just node w). So, the r.v.s $\{\Phi_w\}_{w \in V}$ may be highly correlated. Observe that Φ_w is exactly the feedback obtained when node w is activated directly (by incurring cost c_w) at any point in a policy. Define function $\bar{f} : 2^{E \times \Omega} \to \mathbb{R}_{\geq 0}$ as:

$$\bar{f}(\psi) = \sum_{v \in V} \min \left\{ \sum_{u:(u,v) \in A'} |\{w \in \mathsf{dom}(\psi) : \psi_w(u, v) = 1\}|, 1 \right\}. \tag{3.17}$$

$\bar{f}$ is a sum of set-coverage functions, which is monotone and submodular. Then, utility function $f : 2^{E \times \Omega} \to \mathbb{R}_{\geq 0}$ is $f(\psi) = \min\{\bar{f}(\psi), Q\}$. Function f is clearly monotone (Definition 1.1). The adaptive-submodularity property also holds: see Theorem 19 in [8].

Hence, Theorem 1.1 implies a $4(1 + \ln Q)$-approximation algorithm for adaptive viral marketing. This is an improvement over previous approximation ratios of $(1 + \ln Q)^2$ [8] and $(1 + \ln(nQc_{max}))$ [14], where $n = |V|$ and c_{max} is the maximum cost.

Optimal Decision Tree (uniform prior). In this problem, there are m hypotheses H and n binary tests E. Each test $e \in E$ costs c_e, and has a positive outcome on some subset $T_e \subseteq H$ of hypotheses (its outcome is negative on the other hypotheses).[3] An unknown hypothesis h^* is drawn from H uniformly at random. The goal is to identify h^* by sequentially performing tests, at minimum expected cost. This is a special case of ASC, where the items correspond to tests E and the outcome space $\Omega = \{+, -\}$. The outcome Φ_e for any item e is the test outcome under the (unknown) hypothesis h^*. For any test $e \in E$, define subsets $S_{e,+} = H \setminus T_e$ and $S_{e,-} = T_e$, corresponding to the hypotheses that can be eliminated when we observe a positive or negative outcome on e. The utility function is

$$f(\psi) = \frac{1}{|H|} \cdot |\bigcup_{e \in \mathsf{dom}(\psi)} S_{e,\psi_e}|.$$

The quota $Q = 1 - \frac{1}{|H|}$. Achieving value Q means that $|H| - 1$ hypotheses have been eliminated, which implies that h^* is identified. The function f is again monotone and submodular. The monotonicity property (Definition 1.1) clearly holds. Moreover, using the fact that h^* has a *uniform distribution*, it is known that f is adaptive-submodular: see Lemma 23 in [8].

[3] Our results also extend to the case of multiway tests with non-binary outcomes.

Copyright © 2023 by SIAM
Unauthorized reproduction of this article is prohibited

So, Theorem 1.1 implies a $4(1+\ln(|H|-1))$-approximation algorithm for this problem; we use Q as above and $\eta = \frac{1}{|H|}$. The previous-best bounds for this problem were $(1+\ln(|H|-1))^2$ [8], $12 \cdot \ln|H|$ [3] and $(1+\ln(n|H|c_{max}))$ [14].

We note that [8] also obtained a $\left(\ln \frac{1}{p_{min}}\right)^2$-approximation for the optimal decision tree problem with arbitrary priors (where the distribution of h^* is not uniform); here $p_{min} \le \frac{1}{|H|}$ is the minimum probability of any hypothesis. This uses a different utility function that falls outside our ASC framework (as our definition of function f is more restrictive). Moreover, there are other approaches [4, 11] that provide a better $O(\ln|H|)$-approximation bound even for the problem with arbitrary priors.

4 Adaptive Submodular Cover with Multiple Functions

Here, we extend ASC to the setting of covering multiple adaptive-submodular functions. In the multiple adaptive-submodular cover (MASC) problem, there is a set E of items and outcome space Ω as before. Each item $e \in E$ has a cost c_e; we will view this cost as the item's *processing time*. Now, there are k different utility functions $f_r : 2^{E\times\Omega} \to \mathbb{R}_{\ge 0}$ for $r \in [k]$. We assume that each of these functions satisfies the monotonicity, coverability and adaptive-submodularity properties. We also assume, without loss of generality (by scaling), that the maximal value of each function $\{f_r\}_{r=1}^k$ is Q. As for the basic ASC problem, a solution to MASC corresponds to a policy $\pi : 2^{E\times\Omega} \to E$, that maps partial realizations to the next item to select. Given any policy π, the *cover time* of function f_r is defined as:

$$\mathsf{Cov}_r(\pi) := \text{the earliest time } t \text{ such that } f_r(\Psi(\pi,t)) = Q.$$

Recall that $\Psi(\pi,t) \subseteq E\times\Omega$ is the partial realization that has been observed by time t in policy π. Clearly, the cover time is a random quantity. The objective in MASC is to minimize the expected total cover time of all functions, i.e., $\sum_{r=1}^k \mathbb{E}[\mathsf{Cov}_r(\pi)]$. Note that MASC reduces to ASC when there is $k=1$ function.

Remark: One might also consider an alternative multiple-function formulation where we are interested in the expected *maximum* cover time of the functions (rather than total). This formulation can be directly solved as an instance of ASC where we use the single adaptive-submodular function $g = \sum_{r=1}^k f_r$ with maximal value $Q' = kQ$.

We extend the greedy policy for ASC to MASC, as described in Algorithm 2. For each $r \in [k]$ and item $e \in E$, we use $\Delta_r(e|\psi)$ to denote the marginal benefit of e under function f_r. Notice that the greedy selection criterion here involves a sum of terms corresponding to each un-covered function. A similar greedy rule was used earlier in the (deterministic) submodular function ranking problem [26] and in the independent special case of MASC in [24].

Algorithm 2 Adaptive Greedy Policy $\tilde{\pi}$ for MASC.

selected items $A \leftarrow \emptyset$, observed realizations $\psi \leftarrow \emptyset$
while there exists function f_r with $f_r(\psi) < Q$ **do**
    select item

$$e^* = \arg\max_{e\in E\setminus A} \frac{1}{c_e} \cdot \sum_{r\in[k]: f_r(\psi)<Q} \frac{\Delta_r(e|\psi)}{Q - f_r(\psi)},$$

    and observe Φ_{e^*}
    add e^* to the selected items, i.e., $A \leftarrow A \cup \{e^*\}$
    update $\psi \leftarrow \psi \cup \{(e^*, \Phi_{e^*})\}$

THEOREM 4.1. *Consider any instance of adaptive-submodular cover with k utility functions, where each function $f_r : 2^{E\times\Omega} \to \mathbb{R}_{\ge 0}$ is monotone, coverable and adaptive-submodular w.r.t. the same probability distribution $p(\cdot)$. Suppose that there is some value $\eta > 0$ such that $f_r(\psi) > Q - \eta$ implies $f_r(\psi) = Q$ for all partial realizations $\psi \subseteq E\times\Omega$ and $r \in [k]$. Then, the cost of the greedy policy is*

$$c_{total}(\tilde{\pi}) \le 4 \cdot (1 + \ln(Q/\eta)) \cdot c_{total}(\tilde{\sigma}),$$

where $\tilde{\sigma}$ denotes the optimal policy.

Copyright © 2023 by SIAM
Unauthorized reproduction of this article is prohibited

The proof of Theorem 4.1 is a straightforward extension of Theorem 1.1. We use the same notations and definitions if not mentioned explicitly.

Definition 4.1. (MASC non-completion probabilities) *For any time $t \geq 0$ and $r \in [k]$, let*

$$\tilde{o}_r(t) := \Pr[\tilde{\sigma} \text{ does not cover } f_r \text{ by time } t] = \Pr[\mathsf{Cov}_r(\tilde{\sigma}) > t] = \Pr\left[f_r(\Psi(\tilde{\sigma}, t)) < Q\right].$$
$$\tilde{a}_r(t) := \Pr[\tilde{\pi} \text{ does not cover } f_r \text{ by time } \beta L \cdot t] = \Pr[\mathsf{Cov}_r(\tilde{\pi}) > \beta L \cdot t] = \Pr\left[f_r(\Psi(\tilde{\pi}, \beta L t)) < Q\right].$$

Also, define for any $t \geq 0$, $\tilde{o}(t) := \sum_{r=1}^{k} o_r(t)$ and $\tilde{a}(t) := \sum_{r=1}^{k} a_r(t)$.

The two functions $\tilde{o}(\cdot)$ and $\tilde{a}(\cdot)$ share the same properties with $o(\cdot)$ and $a(\cdot)$ respectively, except that $\tilde{o}(0) = \tilde{a}(0) = k$. We can express the expected total cover times of policies as follows:

$$c_{total}(\tilde{\sigma}) = \int_0^\infty \tilde{o}(t)dt = \sum_{r \in [k]} \int_0^\infty o_r(t)dt.$$

$$c_{total}(\tilde{\pi}) = \int_0^\infty \tilde{a}(\frac{y}{\beta L})dy = \beta L \cdot \int_0^\infty \tilde{a}(t)dt = \beta L \cdot \sum_{r \in [k]} \int_0^\infty a_r(t)dt.$$

Definition 4.2. (MASC greedy score) *For any $t \geq 0$ and partial realization ψ at time t, define*

$$\widetilde{score}(t, \psi) := \frac{1}{c_e} \sum_{r \in [k]: f_r(\psi) < Q} \frac{\Delta_r(e|\psi)}{Q - f_r(\psi)} \quad {\scriptstyle \begin{array}{c}\text{where } e \text{ is the item being selected in } \tilde{\pi} \text{ at time } t \\ \text{when } \psi \text{ was observed just before selecting } e\end{array}},$$

and $\widetilde{score}(t, \psi) := 0$ if no item is being selected in $\tilde{\pi}$ at time t when ψ was observed.

Definition 4.3. (MASC gain in a phase) *For any phase $i \geq 0$, its* gain *is the expected total score accumulated during phase i,*

$$\tilde{G}_i := \int_{L\beta i}^{L\beta(i+1)} \mathbb{E}\left[\widetilde{score}(t, \Psi_t)\right] dt$$

The next two lemmas lower and upper bound the gains.

Lemma 4.1. *For any $i \geq 0$, the total gain after phase i is*

$$\sum_{j \geq i} \tilde{G}_j \leq \sum_{r \in [k]} L \cdot \tilde{a}_r(i) = L \cdot \tilde{a}(i).$$

To prove Lemma 4.1, we just apply Lemma 2.1 to each f_r and sum the results.

Lemma 4.2. *For any $i \geq 0$ and time $t \in [L\beta i, L\beta(i+1))$,*

$$\mathbb{E}[\widetilde{score}(t, \Psi_t)] \geq \sum_{r \in [k]} \frac{\tilde{a}_r(i+1) - \tilde{o}_r(i+1)}{i+1} = \frac{\tilde{a}(i+1) - \tilde{o}(i+1)}{i+1}.$$

Hence, $\tilde{G}_i \geq \beta L \cdot \left(\frac{\tilde{a}(i+1) - \tilde{o}(i+1)}{i+1}\right)$ for each phase $i \geq 0$.

To prove Lemma 4.2, we replicate all the steps in Lemma 2.2, except that we redefine Z to be

$$\tilde{Z} := \mathbb{E}_\Phi\left[\mathbb{1}(\psi \preccurlyeq \Phi) \cdot \sum_{r \in [k]: f_r(\psi) < Q} \frac{f_r(\psi \cup P_\infty) - f_r(\psi)}{Q - f_r(\psi)}\right],$$

which takes all k functions into account. Here, ψ is any partial realization observed at time t. We then obtain:

$$\widetilde{score}(t, \psi) \cdot (i+1) \cdot \Pr[\psi \preccurlyeq \Phi] \geq \tilde{Z} \geq \sum_{r \in [k]: f_r(\psi) < Q} \Pr[(\psi \preccurlyeq \Phi) \wedge (\overline{\sigma} \text{ covers } f_r)].$$

Copyright © 2023 by SIAM
Unauthorized reproduction of this article is prohibited

Let $R(\tilde{\pi}, t)$ denote *all* the possible partial realizations observed at time t in policy $\tilde{\pi}$. Then, we get

$$\begin{aligned}
\mathbb{E}[\widetilde{score}(t, \Psi_t)] &= \sum_{\psi \in R(\tilde{\pi},t)} \Pr[\psi \preccurlyeq \Phi] \cdot \widetilde{score}(t, \psi) \\
&\geq \frac{1}{i+1} \sum_{\psi \in R(\tilde{\pi},t)} \sum_{r \in [k]: f_r(\psi) < Q} \Pr[(\psi \preccurlyeq \Phi) \wedge (\overline{\sigma} \text{ covers } f_r)] \\
&= \frac{1}{i+1} \sum_{r \in [k]} \Pr\left[(\tilde{\pi} \text{ doesn't cover } f_r \text{ by time } t) \wedge (\overline{\sigma} \text{ covers } f_r)\right] \\
&\geq \frac{1}{i+1} \sum_{r \in [k]} \left(\Pr[\tilde{\pi} \text{ doesn't cover } f_r \text{ by time } t] - \Pr[\overline{\sigma} \text{ doesn't cover } f_r]\right) \\
&= \frac{1}{i+1} \sum_{r \in [k]} \left(\tilde{a}_r\left(\frac{t}{L\beta}\right) - o_r(i+1)\right) \geq \frac{\tilde{a}(i+1) - \tilde{o}(i+1)}{i+1}.
\end{aligned}$$

Finally, we combine Lemma 4.1 and Lemma 4.2, as in Theorem 1.1, to obtain:

$$c_{total}(\tilde{\pi}) \leq \frac{L\beta^2}{\beta - 1}(c_{total}(\tilde{\sigma}) + k).$$

Note that the $+k$ term comes from $\sum_{i \geq 0} \tilde{o}(i) - \tilde{o}(0) \leq \sum_{i \geq 0} \tilde{o}(i+1) \leq \int_0^\infty \tilde{o}(t)dt$ and $\tilde{o}(0) = k$. Following the same argument as in the single-function case, the extra $+k$ term can be removed. Setting $\beta = 2$, we obtain Theorem 4.1.

We now list some applications of the MASC result.

- When items are deterministic, MASC reduces to the deterministic submodular ranking problem, for which an $O(\ln(Q/\eta))$ was obtained in [26]. We note that the result in [26] was only for unit costs, whereas our result holds for arbitrary costs. This problem generalizes the min-sum set cover problem, which is NP-hard to approximate better than factor 4 [25]. For min-sum set cover, the paramters $Q = \eta = 1$: so Theorem 4.1 implies a *tight* 4-approximation algorithm for it. This also suggests that the leading constant of 4 in our approximation ratio is best possible for MASC.

- When the outcomes are independent across items, MASC reduces to stochastic submodular cover with multiple functions, which was studied in [24]. We obtain an $4 \cdot (1 + \ln(Q/\eta))$ approximation ratio that improves the bound of $56 \cdot (1 + \ln(Q/\eta))$ in [24] by a constant factor. Although [24] did not try to optimize the constant factor, their approach seems unlikely to provide such a small constant factor.

- Consider the following generalization of adaptive viral marketing. Instead of a single quota on the number of influenced nodes, there are k different quotas $Q_1 \leq Q_2 \leq \cdots Q_k$. Now, we want a policy such that the *average* expected cost for achieving these quotas is minimized. Recall the function $\bar{f}$ defined in (3.17) for the single-quota problem. Then, corresponding to the different quotas, define functions $f_r(\psi) = \frac{1}{Q_r} \cdot \min\{\bar{f}(\psi), Q_r\}$ for each $r \in [k]$. Each of these functions is monotone, adaptive-submodular and has maximal value $Q = 1$. The parameter $\eta = 1/Q_k$, so we obtain a $4(1 + \ln Q_k)$-approximation algorithm.

References

[1] A. Asadpour and H. Nazerzadeh, *Maximizing Stochastic Monotone Submodular Functions*, Manag. Sci., 62.8 (2016), pp. 2374–2391.

[2] A. Agarwal, S. Assadi, and S. Khanna, *Stochastic Submodular Cover with Limited Adaptivity*, in Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms, 2019, pp. 323–342.

[3] A. Guillory and J. A. Bilmes, *Average-Case Active Learning with Costs*, in R. Gavaldà, G. Lugosi, T. Zeugmann, and S. Zilles, editors, Algorithmic Learning Theory, Lecture Notes in Computer Science, Springer, 5809 (2009), pp. 141–155.

Copyright © 2023 by SIAM
Unauthorized reproduction of this article is prohibited

[4] A. Gupta, V. Nagarajan, and R. Ravi, *Approximation Algorithms for Optimal Decision Trees and Adaptive TSP Problems*, Math. Oper. Res., 42.3 (2017), pp. 876–896.
[5] A. Gupta, V. Nagarajan, and S. Singla, *Adaptivity Gaps for Stochastic Probing: Submodular and XOS Functions*, in Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms, 2017, pp. 1688–1702.
[6] C. Guestrin, A. Krause, and A. P. Singh, *Near-Optimal Sensor Placements in Gaussian Processes*, in Proceedings of the 22nd International Conference on Machine Learning, 2005, pp. 265–272.
[7] D. Bradac, S. Singla, and G. Zuzic, *(Near) Optimal Adaptivity Gaps for Stochastic Multi-Value Probing*, in Approximation, Randomization, and Combinatorial Optimization, 145 (2019), pp. 49:1–49:21.
[8] D. Golovin and A. Krause, *Adaptive Submodularity: A New Approach to Active Learning and Stochastic Optimization*, CoRR, abs/1003.3967 (2017).
[9] D. Golovin and A. Krause, *Adaptive Submodularity: Theory and Applications in Active Learning and Stochastic Optimization*, J. Artif. Intell. Res. (JAIR), 42 (2011), pp. 427-486.
[10] D. Kempe, J. M. Kleinberg, and É. Tardos, *Maximizing the Spread of Influence through a Social Network*, Theory Comput., 11 (2015), pp. 105–147.
[11] F. Navidi, P. Kambadur, and V. Nagarajan, *Adaptive Submodular Ranking and Routing*, INFORMS Oper. Res., 68.3 (2020), pp. 856–877.
[12] F. Nan and V. Saligrama, *Comments on the Proof of Adaptive Stochastic Set Cover Based on Adaptive Submodularity and Its Implications for the Group Identification Problem in "Group-Based Active Query Selection for Rapid Diagnosis in Time-Critical Situations"*, IEEE Transactions on Information Theory, 63.11 (2017), pp. 7612–7614.
[13] G. Tong, W. Wu, S. Tang, and D. Du, *Adaptive Influence Maximization in Dynamic Social Networks*, IEEE/ACM Transactions on Networking, 25.1 (2017), pp. 112-125.
[14] H. Esfandiari, A. Karbasi, and V. Mirrokni, *Adaptivity in Adaptive Submodularity*, in Proceedings of 34th Conference on Learning Theory, PMLR, 134 (2021), pp. 1823–1846.
[15] I. Dinur and D. Steurer, *Analytical approach to parallel repetition*, in Symposium on Theory of Computing, 2014, pp. 624–633.
[16] L. A. Wolsey, *An analysis of the greedy algorithm for the submodular set covering problem*, Combinatorica, 2.4 (1982), pp. 385–393.
[17] L. Hellerstein, D. Kletenik, and S. Parthasarathy, *A Tight Bound for Stochastic Submodular Cover*, J. Artif. Intell. Res., 71 (2021), pp. 347–370.
[18] L. Hellerstein and D. Kletenik, *Revisiting the Approximation Bound for Stochastic Submodular Cover*, J. Artif. Intell. Res., 63 (2018), pp. 265–279.
[19] M. Adamczyk, M. Sviridenko, and J. Ward, *Submodular Stochastic Probing on Matroids*, Math. Oper. Res., 41.3 (2016), pp. 1022–1038.
[20] M. Goemans and J. Vondrák, *Stochastic Covering and Adaptivity*, in LATIN 2006: Theoretical Informatics, Springer Berlin Heidelberg, 2006, pp. 532–543.
[21] N. Grammel, L. Hellerstein, D. Kletenik, and P. Lin, *Scenario submodular cover*, in International Workshop on Approximation and Online Algorithms, Springer, 2016, pp. 116–128.
[22] R. Ghuge, A. Gupta, and V. Nagarajan, *The Power of Adaptivity for Stochastic Submodular Cover*, in Proceedings of the 38th International Conference on Machine Learning, 139 (2021), pp. 3702–3712.
[23] S. Dasgupta, *Analysis of a greedy active learning strategy*, in Advances in Neural Information Processing Systems, 2004, pp. 337–344.
[24] S. Im, V. Nagarajan, and R. Zwaan, *Minimum Latency Submodular Cover*, ACM Trans. Algorithms, 13.1 (2016), pp. 13:1–13:28.
[25] U. Feige, L. Lovász, and P. Tetali, *Approximating min sum set cover*, Algorithmica, 40.4 (2004), pp. 219–234.
[26] Y. Azar and I. Gamzu, *Ranking with Submodular Valuations*, in D. Randall, editor, Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, 2011, pp. 1070–1079.

A Missing Proofs

Below, we upper bound the total score for a single realization ϕ. A similar fact and proof was used previously in [17, 24, 26].

Lemma A.1. *For any (full) realization ϕ,*

$$\int_0^\infty S(t,\phi)dt \le L = \ln(Q/\eta) + 1$$

Proof. Under ϕ, let $e_1, e_2, ..., e_k$ be the sequence of items selected by π, let ψ_i be the partial realization just before selecting e_i, and define $f_i := f(\psi_i)$. Note that $0 \le f_1 \le f_2 \le ... \le f_{k+1} = Q$ by monotonicity and the assumption

Copyright © 2023 by SIAM
Unauthorized reproduction of this article is prohibited

that f is always covered by π. Moreover, we have $f_k \leq Q - \eta$ by the definition of η: otherwise we would have $f_k = Q$ and π would terminate before selecting e_k.

Define a function $g : [0, \infty) \to \mathbb{R}$ by

$$g(x) := \begin{cases} \frac{1}{Q-f_i}, & \text{over } [f_i, f_{i+1}) \text{ for } i = 1, 2, ..., k-1, \\ 0, & \text{otherwise.} \end{cases} \tag{A.1}$$

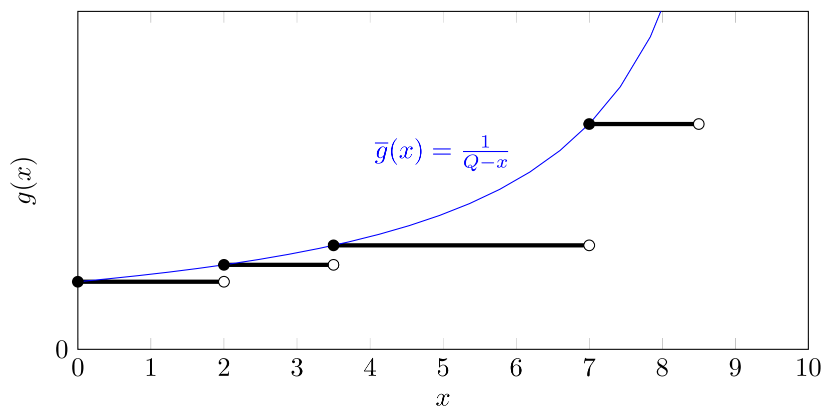

Figure 3: Example of $g(x)$, where $k = 5, Q = 10, \eta = 1$ and $f_1 = 0, f_2 = 2, f_3 = 3.5, f_4 = 7, f_5 = 8.5$.

Note that $g(x) \leq \frac{1}{Q-x}$ for all $0 \leq x \leq Q - \eta$. So,

$$\int_0^\infty S(t, \phi) dt = \sum_{i=1}^{k} \frac{f_{i+1} - f_i}{Q - f_i} \leq \sum_{i=1}^{k-1} \frac{f_{i+1} - f_i}{Q - f_i} + 1 = \int_0^\infty g(x) dx + 1 \leq \int_0^{Q-\eta} \frac{1}{Q-x} dx + 1 = L$$

$\square$

Copyright © 2023 by SIAM
Unauthorized reproduction of this article is prohibited