

Non-Monotone Adaptive Submodular Maximization

Alkis Gotovos
ETH Zurich

Amin Karbasi
Yale University

Andreas Krause
ETH Zurich

Abstract

A wide range of AI problems, such as sensor placement, active learning, and network influence maximization, require sequentially selecting elements from a large set with the goal of optimizing the utility of the selected subset. Moreover, each element that is picked may provide stochastic feedback, which can be used to make smarter decisions about future selections. Finding efficient policies for this general class of adaptive optimization problems can be extremely hard. However, when the objective function is adaptive monotone and adaptive submodular, a simple greedy policy attains a $1 - 1/e$ approximation ratio in terms of expected utility. Unfortunately, many practical objective functions are naturally non-monotone; to our knowledge, no existing policy has provable performance guarantees when the assumption of adaptive monotonicity is lifted. We propose the *adaptive random greedy policy* for maximizing adaptive submodular functions, and prove that it retains the aforementioned $1 - 1/e$ approximation ratio for functions that are also adaptive monotone, while it additionally provides a $1/e$ approximation ratio for non-monotone adaptive submodular functions. We showcase the benefits of adaptivity on three real-world network data sets using two non-monotone functions, representative of two classes of commonly encountered non-monotone objectives.

1 Introduction

Many practical problems in artificial intelligence boil down to selecting a number of elements from a large set of options in an initially unknown environment, so as to maximize some utility function defined over subsets of selected elements. Example problems include sensor placement (selecting sensor locations), active learning (selecting examples to label), and network influence maximization (selecting seed nodes). In contrast to the *non-adaptive* setting, where we commit to the selected elements all at once, in the *adaptive* setting the selection process is performed in a sequential manner, and each element that is picked provides some form of stochastic feedback, or, in other words, reveals part of the environment. Nat-




Figure 1: Birdwatching maximum cover example.

urally, we would like to leverage the acquired feedback to make smarter future selections.

As a running example, consider the toy instance of a stochastic maximum coverage problem shown in Figure 1. Suppose that square nodes represent birdwatching locations; by visiting each of them we observe a random number of bird species among those that are known to exist in that location.¹ We would like to plan a birdwatching trip to visit a number of these locations and maximize the number of observed species. It is intuitive that taking into account the already observed species in deciding which place to visit next can be greatly beneficial compared to committing to the full trip in advance.

Since the general class of adaptive optimization problems in partially observable environments does not admit efficient policies, previous work has focused on characterizing subclasses of problems that can be solved efficiently. Most notably, Golovin and Krause [2011] showed that, when the objective function under consideration is adaptive monotone and adaptive submodular, then a simple greedy policy attains a $1 - 1/e$ approximation ratio in terms of expected utility.

Adaptive submodularity and its intuitive meaning of “diminishing returns” is fundamental to a number of objective functions of interest in the adaptive setting. However, the additional assumption of adaptive monotonicity, i.e., that adding an element to our selection always leads to an increase in expected utility, is often not satisfied in practice. In our birdwatching example, the original objective is adaptive monotone submodular; however, if each subset of locations has an associated cost, e.g., the total required travel distance to visit all of them, our new objective will naturally be non-

¹<http://www.allaboutbirds.org/>

monotone. More broadly, as we will see, a common way of obtaining adaptive submodular objectives is by generalizing non-adaptive submodular functions under certain assumptions about the form of stochastic feedback. Despite the desire to generalize problems with non-monotone submodular objectives to the adaptive setting, no existing policies provide guarantees about non-monotone objectives in this setting.

To resolve this, we propose a new policy for maximizing non-monotone adaptive submodular functions, called *adaptive random greedy*. We prove that this policy retains the $1 - 1/e$ approximation ratio, if the function at hand is additionally adaptive monotone, while it also provides a $1/e$ -approximation for non-monotone functions. Our proposed policy generalizes the random greedy algorithm proposed by Buchbinder *et al.* [2014] for non-monotone submodular optimization to the adaptive setting. We further discuss two common ways in which non-monotone adaptive submodular objectives come about, and present a representative example objective for each of them. Finally, we evaluate our policy on these two objectives on some real-world network data sets, and, thereby, showcase the potential benefits of adaptivity.

2 Problem Statement and Background

Assume we are given a finite ground set E and a set O of observable states. Each item $e \in E$ is associated with a state $o \in O$ through a function $\phi : E \rightarrow O$, which is called a *realization* of the set of states. We also assume that the realization Φ is a random variable with known distribution $p(\phi)$. Furthermore, we are given an objective function $f : 2^E \times O^E \rightarrow \mathbb{R}_{\geq 0}$. For a set $A \subseteq E$ and a realization ϕ , the quantity $f(A, \phi)$ represents the utility of selecting subset A when the true realization is ϕ . In our birdwatching example, the ground set E consists of the three possible locations $\{\ell_1, \ell_2, \ell_3\}$ we may visit, and the state of each location is the (random) subset of bird species that we actually observe if we travel to that location; each realization maps all locations to subsets of actually observed species. Our objective function is the total number of observed species given a subset of visited locations.

Our goal is to come up with a sequential policy that—initially unaware of ϕ —builds up a set $A \subseteq E$, such that our utility $f(A, \phi)$ is maximized. That is, we iteratively select an item $e \in E$ to add to A and observe its state $\phi(e)$. In this setting, there are two factors that complicate matters compared to its non-adaptive counterpart. First, since utility depends on the random realization, the quantity of focus is the *expected* utility under the distribution of realizations $p(\phi)$. Second, the chosen set A itself is a random variable that depends on the realization, since the choices of our policy will change according to each observation $\phi(e)$, which is, of course, the whole point of adaptivity. In addition, the policy itself might make random decisions, which is an additional source of randomness for A .

To address the above complications, we define a *partial realization* as a set $\psi \subseteq E \times O$, which represents the item-observation pairs over a subset of E . In particular, we call this subset the *domain* of ψ , which is formally defined as $\mathcal{D}(\psi) := \{e \in E \mid \exists o \in O : (e, o) \in \psi\}$. Additionally, we

write $\psi(e) = o$, if $(e, o) \in \psi$, and call ψ *consistent* with realization ϕ (denoted by $\phi \sim \psi$), if $\psi(e) = \phi(e)$, for all $e \in \mathcal{D}(\psi)$, which means that the observations of a subset according to ψ agree with the assignments over the whole ground set according to ϕ .

We can now define a *policy* π as a function from partial realizations to a distribution $\mathcal{P}(E)$ over items that specifies which item to pick next, formally, $\pi : 2^{E \times O} \rightarrow \mathcal{P}(E)$. For birdwatching, our policy would specify which location to visit next (or, more generally, a distribution of next locations), given the already visited locations and the already observed species at each of them. The policy terminates when the current partial realization is not in its domain denoted by $\mathcal{D}(\pi) \subseteq 2^{E \times O}$. We also use the shorthand notation $\pi(e \mid \psi)$ for the probability of picking item e given partial realization ψ . We call $E(\pi, \phi) \subseteq E$ the set of items that have been selected upon termination of policy π under realization ϕ (our final set of visited birdwatching locations). Note that $E(\pi, \Phi)$ is a random variable that depends on both the randomness of the policy, as well as the randomness of realizations.

Finally, we can formally assess the performance of a policy π via its expected utility,

$$f_{\text{avg}}(\pi) := \mathbb{E}_{\Phi, \Pi}[f(E(\pi, \Phi), \Phi)].$$

Then, our goal is to come up with a policy that maximizes the expected utility, subject to a cardinality constraint on the number of items to be picked, $|E(\pi, \Phi)| \leq k$. In birdwatching terms, find a policy that maximizes the expected number of observed bird species, if we can visit at most k locations.

2.1 Monotonicity and Submodularity

Non-adaptive. Even in the non-adaptive setting, where we have to commit to a subset in advance, the problem of maximizing a set function $f : 2^E \rightarrow \mathbb{R}_{\geq 0}$, subject to a cardinality constraint $|A| \leq k$, is NP-hard in general. In this setting, the marginal gain of an element $e \in E$ given set $B \subseteq E$ is defined as $f(B \cup \{e\}) - f(B)$. Intuitively, the marginal gain quantifies the increase in utility if we add e to our selection, given that we have already picked the elements in B . Function f is called *monotone* if, for any $B \subseteq C \subseteq E$, it holds that $f(B) \leq f(C)$, which is equivalent to saying that the marginal gain is always non-negative. Furthermore, f is called *submodular*, if, for any $B \subseteq C \subseteq E$ and any $e \in E \setminus C$, it holds that $f(C \cup \{e\}) - f(C) \leq f(B \cup \{e\}) - f(B)$. This means that the marginal gain of any element decreases as the given set increases ($C \supseteq B$); in other words, submodularity expresses a property of “diminishing returns” as more and more elements are added to our selection.

In their seminal work, Nemhauser *et al.* [1978] showed that, if f is non-negative, monotone, and submodular, then constructing a subset A of size k by greedily picking elements according to their marginal gains, guarantees that $f(A)$ is a $(1 - 1/e)$ -approximation of the optimal value.

Adaptive. In the significantly more complex adaptive setting, the problem of computing an optimal policy is hard to approximate even for seemingly simple classes of objective functions (e.g., linear), as shown by Golovin and Krause [2011]. However, they also showed that the notions of

monotonicity and submodularity can be naturally generalized to this setting, and lead to similar performance guarantees to the non-adaptive setting.

More concretely, the *expected marginal gain* of an element $e \in E$ given partial realization ψ can be defined as

$$\Delta(e|\psi) := \mathbb{E}_{\Phi} [f(\mathcal{D}(\psi) \cup \{e\}, \Phi) - f(\mathcal{D}(\psi), \Phi) \mid \Phi \sim \psi].$$

The above expression is a conditional expectation, which only considers realizations that are consistent with ψ . In bird-watching terms, it quantifies how many new species we expect to observe if we visit a new location e given the already visited locations and already observed species in ψ . Then, the following properties can be defined analogously to their non-adaptive counterparts:

- f is called *adaptive monotone*, if $\Delta(e|\psi) \geq 0$, for all $e \in E$, and all ψ of positive probability,
- f is called *adaptive submodular*, if $\Delta(e|\psi') \leq \Delta(e|\psi)$, for all $e \in E \setminus \mathcal{D}(\psi')$, and all $\psi' \supseteq \psi$.

Given a number of previously selected elements and their corresponding observed states encoded in partial realization ψ , the adaptive greedy policy selects the element $e \in E \setminus \mathcal{D}(\psi)$ of highest marginal gain $\Delta(e|\psi)$, and continues to do so iteratively until k elements have been selected. Golovin and Krause [2011] showed that, if f is adaptive monotone submodular, then adaptive greedy is a $(1 - 1/e)$ -approximation in terms of expected utility f_{avg} .

3 Adaptive Random Greedy

While adaptive monotonicity is satisfied by many functions of interest, it is often the case that modeling practical problems naturally results in non-monotone objectives (see Section 4); no existing policy provides provable performance guarantees in this case. We now present our proposed adaptive random greedy policy (π^{rg}) for maximizing adaptive submodular functions, and prove approximation ratios irrespective of whether adaptive monotonicity is satisfied or not.

For technical reasons that will become apparent below, let us assume that we always add a set D of $2k - 1$ dummy elements to the ground set, such that, for any $d \in D$, and any partial realization ψ , it holds that $\Delta(d|\psi) = 0$. Obviously, these elements do not affect the optimal policy, and may be removed from the solution of any policy, without affecting its expected utility.

The detailed pseudocode of running the adaptive random greedy policy is presented in Algorithm 1. As discussed before, the algorithm is given a ground set and an objective function, as well as a known distribution over realizations Φ . At each iteration, the first step is to compute the expected marginal gain of each remaining element (line 4). The key difference compared to the original adaptive greedy policy is shown in lines 5–6; rather than selecting the element with the largest expected marginal gain, π^{rg} randomly selects an element from the set $\mathcal{M}_k(\psi)$, which contains the elements with the k largest gains. The dummy elements added to the ground set ensure that the policy never picks an element with negative expected marginal gain. Also, note that, although in Algorithm 1 the returned set A contains exactly k elements,

Algorithm 1 Adaptive random greedy

Input: ground set E , function f , distribution $p(\phi)$, cardinality constraint k

- 1: $A \leftarrow \emptyset$
- 2: $\psi \leftarrow \emptyset$
- 3: **for** $i = 1$ **to** k **do**
- 4: Compute $\Delta(e|\psi)$, for all $e \in E \setminus A$
- 5: $\mathcal{M}_k(\psi) \leftarrow \operatorname{argmax}_{S \subseteq E \setminus A, |S|=k} \left\{ \sum_{e \in S} \Delta(e|\psi) \right\}$
- 6: Sample m uniformly at random from $\mathcal{M}_k(\psi)$
- 7: $A \leftarrow A \cup \{m\}$
- 8: Observe $\Phi(m)$
- 9: $\psi \leftarrow \psi \cup \{(m, \Phi(m))\}$
- 10: **end for**
- 11: **Return** A

the actual selected set may very well contain less than k elements, since we implicitly assume that any dummy elements are removed from it after the policy terminates.

When running the original adaptive greedy policy, non-monotonicity can lead to situations, where selecting the element of maximum marginal gain leads to traps of low utility that cannot be escaped. In contrast, the randomization in the process of selecting each element that is introduced by adaptive random greedy helps dealing with such traps (on average) and, thus, leads to provable approximation guarantees for the expected utility, even for non-monotone objectives.

Theoretical analysis. More concretely, in this paper we show that the adaptive random greedy policy retains the $1 - 1/e$ approximation ratio for adaptive monotone submodular objectives, while, at the same time, it achieves a $1/e$ approximation ratio for non-monotone objectives, under the additional condition of submodularity for each realization. As we will see in Section 4, for the vast majority of non-monotone objectives used in practice this condition holds by way of construction. We, therefore, do not consider it a major restriction in the choice of objectives. Our proofs generalize the results of Buchbinder *et al.* [2014] for the random greedy algorithm, and are presented in detail in the long version of this paper; in what follows, we provide an outline of our analysis.

First, let us define the expected gain of running policy π after having obtained a partial realization ψ , as

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

Also, for any policy π and any positive integer k , we define the truncated policy $\pi_{[k]}$, which runs identically to π for k steps and then terminates. Finally, given any two policies π_1, π_2 , we denote by $\pi_1 @ \pi_2$ the policy that first runs π_1 until it terminates and then runs π_2 , discarding all observations made by π_1 . After $\pi_1 @ \pi_2$ terminates, the subset selected by it consists of the union of the subsets selected by each policy individually.

The following is a key lemma for both monotone and non-monotone objectives.

Lemma 1. *If f is adaptive submodular, then, for any policy π , and any partial realization ψ , it holds that*

$$\Delta(\pi_{[k]} | \psi) \leq \sum_{e \in \mathcal{M}_k(\psi)} \Delta(e | \psi).$$

It states the intuitive fact that, at any point, no k -step policy can give us a larger expected gain than the sum of the k currently largest expected marginal gains. This is a consequence of adaptive submodularity, which guarantees that the expected marginal gains of any element decrease as our selection grows larger.

Based on Lemma 1 and the fact that π^{rg} selects at each step one of the k elements in \mathcal{M}_k uniformly at random, we can show the following lemma, which also applies to both monotone and non-monotone objectives.

Lemma 2. *For any policy π , and any non-negative integer $i < k$, if f is adaptive submodular, then*

$$f_{\text{avg}}(\pi_{[i+1]}^{\text{rg}}) - f_{\text{avg}}(\pi_{[i]}^{\text{rg}}) \geq \frac{1}{k} \left(f_{\text{avg}}(\pi_{[i]}^{\text{rg}} @ \pi) - f_{\text{avg}}(\pi_{[i]}^{\text{rg}}) \right).$$

The lemma compares the expected gain at the i -th step of π^{rg} to the total gain of running any other policy (e.g., the optimal one) after the i -th step and, thereby, provides a means for obtaining approximation guarantees for π^{rg} , as long as we can bound the term $f_{\text{avg}}(\pi_{[i]}^{\text{rg}} @ \pi)$. The dichotomy between adaptive monotone and non-monotone objectives in terms of theoretical guarantees stems from the different approaches in bounding this term.

If f is adaptive monotone, we may use the trivial bound $f_{\text{avg}}(\pi_{[i]}^{\text{rg}} @ \pi) \geq f_{\text{avg}}(\pi)$, which immediately follows from the definition of adaptive monotonicity, to obtain the following theorem.

Theorem 1. *If f is adaptive monotone submodular, then for any policy π , and all integers $i, k > 0$ it holds that*

$$f_{\text{avg}}(\pi_{[i]}^{\text{rg}}) \geq \left(1 - e^{-i/k}\right) f_{\text{avg}}(\pi_{[k]}).$$

In particular, by setting $i = k$ we get the familiar $1 - 1/e$ approximation ratio for $\pi_{[k]}^{\text{rg}}$.

For the non-monotone case, we need to leverage the randomness of the selection process of π^{rg} to bound $f_{\text{avg}}(\pi_{[i]}^{\text{rg}} @ \pi)$. For that purpose, we generalize to the adaptive setting the following lemma shown by Buchbinder *et al.* [2014], which itself is based on a lemma by Feige *et al.* [2007] for the expected value of a submodular function under a randomly selected subset.

Lemma 3 ([Buchbinder *et al.*, 2014]). *If $f : 2^E \rightarrow \mathbb{R}_{\geq 0}$ is submodular and A is a random subset of E , such that each element $e \in E$ is contained in A with probability at most p , that is, $\mathbb{P}_A[e \in A] \leq p$, $\forall e \in E$, then*

$$\mathbb{E}_A[f(A)] \geq (1 - p)f(\emptyset).$$

Roughly speaking, the lemma states that a ‘‘random enough’’ subset A cannot have much worse value than that of the empty set. Note that f is not assumed to be monotone here.

The following lemma extends the above claim to the adaptive setting.

Lemma 4. *If f is adaptive submodular, and, additionally, $f(\cdot, \phi) : 2^E \rightarrow \mathbb{R}_{\geq 0}$ is submodular for all $\phi \in O^E$, then for any policy π such that each element of $e \in E$ is selected by it with probability at most p , that is, $\mathbb{P}_{\Pi}[e \in E(\pi, \phi)] \leq p$, $\forall \phi \in O^E$, $\forall e \in E$, the expected value of running π can be bounded as follows:*

$$f_{\text{avg}}(\pi) \geq (1 - p) f_{\text{avg}}(\pi_{[0]}).$$

As a consequence of the above lemma, we get that $f_{\text{avg}}(\pi_{[i]}^{\text{rg}} @ \pi) = f_{\text{avg}}(\pi @ \pi_{[i]}^{\text{rg}}) \geq (1 - p)f_{\text{avg}}(\pi)$, meaning that the elements added by adaptive random greedy cannot dramatically reduce the average value obtained by any other policy π . The probability p in this case can be bounded by using the fact that π^{rg} randomly selects one of k elements at each step, hence at the i -step we have $p \leq (1 - 1/k)^i$. Putting it all together, we obtain our main theorem for non-monotone objectives.

Theorem 2. *If f is adaptive submodular, and $f(\cdot, \phi) : 2^E \rightarrow \mathbb{R}_{\geq 0}$ is submodular for all $\phi \in O^E$, then, for any policy π , and all integers $i, k > 0$, it holds that*

$$f_{\text{avg}}(\pi_{[i]}^{\text{rg}}) \geq \frac{i}{k} \left(1 - \frac{1}{k}\right)^{i-1} f_{\text{avg}}(\pi_{[k]}).$$

By setting $i = k$, we get a $1/e$ approximation ratio for $\pi_{[k]}^{\text{rg}}$.

4 Examples of Non-Monotone Objectives

To underline the importance of non-monotonicity, we now present two different ways, in which non-monotone adaptive submodular functions commonly arise in practice. Both of the resulting classes of functions satisfy the assumptions of Theorem 2, and are, therefore, suitable to be maximized using adaptive random greedy. We also introduce two representative example objectives, one for each class, which are themselves of practical interest.

4.1 Objectives with a Modular Cost Term

Assume we are given an adaptive monotone submodular function $f_{\text{utility}}(A, \phi)$. In practice, apart from benefit, there might also be some associated cost with the selection of each element, which can be directly incorporated into the objective function via a modular cost term $f_{\text{cost}}(A) = \sum_{a \in A} c_a$. In this case, the resulting objective is of the form

$$f(A, \phi) = f_{\text{utility}}(A, \phi) - f_{\text{cost}}(A), \quad (1)$$

which is also adaptive submodular, but non-monotone.

A common alternative to the above is to introduce a knapsack constraint $f_{\text{cost}}(A) \leq C$, while retaining a monotone objective. Choosing a budget C , in this case, might not be obvious, whereas the formulation of equation (1) is often more straightforward, particularly when f_{cost} is expressed in the same units as f_{utility} .

Influence maximization. The concept of influence maximization in a social network was posed by Kempe *et al.* [2003], and has direct applications to problems such as viral marketing. Given a graph and a model of influence propagation, the goal is to select a subset of nodes that are initially active, in order to maximize the spread of influence measured

by the expected number of nodes that will ultimately be active according to the propagation model. We focus here on the independent cascade model, according to which, each edge of the graph is randomly set to be “live” or “blocked”, independently of any other edge in the network, and influence can only flow along “live” edges.

For this problem, the ground set E consists of network nodes, and each realization ϕ corresponds to a full outcome of the independent cascade model, that is, an assignment to each network edge of being either “live” or “blocked”. The objective $f_{\text{inf}}(A, \phi)$ is the number of ultimately active nodes under realization ϕ , if the nodes in the selected subset A are initially active. In the adaptive version of the problem, when a node v is selected, it reveals the status (“live” or “blocked”) of all outgoing edges of v and of any other node that can be reached from v through “live” edges. Kempe *et al.* [2003] showed that f_{inf} is non-negative, monotone, and submodular, for any realization ϕ ; Golovin and Krause [2011] showed that it is also adaptive monotone submodular.

Now, assume that each selected node incurs a unit cost, that is, we have a cost term $f_{\text{cost}}(A) = |A|$, which results in the following objective:

$$\tilde{f}_{\text{inf}}(A, \Phi) := f_{\text{inf}}(A, \Phi) - |A|.$$

Since our cost term is modular, \tilde{f}_{inf} has the form of equation (1), and, therefore, is non-monotone adaptive submodular. It is also non-negative, since $f_{\text{inf}}(A, \Phi) \geq |A|$.

4.2 Objectives with Factorial Realizations

Assume we are given a utility function $f(A, \phi)$, whose dependence of ϕ is constrained to the outcomes of the selected elements. More formally, there exists a function $g : 2^{E \times O} \rightarrow \mathbb{R}_{\geq 0}$, such that $f(A, \phi) = g(\{(e, \phi(e)) \mid e \in A\})$. Furthermore, assume that f is submodular in its first argument, for any realization ϕ , and that the distribution of realizations factorizes over the elements of the ground set E , that is, $\mathbb{P}_{\Phi}[\phi] = \prod_{e \in E} \mathbb{P}_{\Phi(e)}[\phi(e)]$. Given the above assumptions, it follows that f is adaptive submodular (see Theorem 6.1 of Golovin and Krause [2011]).

Maximum graph cut. The problem of finding the maximum cut in a graph $(\mathcal{V}, \mathcal{E})$ can be posed using the following objective:

$$f_{\text{cut}}(A) = \sum_{(v,w) \in \mathcal{E}} \mathbb{1}_{\{v \in A, w \in \mathcal{V} \setminus A \text{ or } v \in \mathcal{V} \setminus A, w \in A\}},$$

which is non-negative and submodular. Furthermore, it is symmetric, i.e., $f(A) = f(\mathcal{V} \setminus A)$, which implies that it is also non-monotone. We consider here an adaptive version of the maximum cut problem, where the selection of a node triggers either cutting that node itself, or a random neighbor of it with some prespecified probability.

Again, the ground set E consists of network nodes, and each realization ϕ corresponds to a function σ_{ϕ} that maps each node $v \in V$ to the node that would actually be cut, were v to be selected. Our adaptive max-cut objective can then be written as follows:

$$\tilde{f}_{\text{cut}}(A, \Phi) := f_{\text{cut}}\left(\bigcup_{v \in A} \sigma_{\Phi}(v)\right).$$

We can directly see that, in terms of the realization Φ , \tilde{f}_{cut} only depends on the outcomes of the selected elements. Furthermore, the distribution of realizations is factorial, since the outcome of each node cut is independent of all the others, i.e., $\sigma_{\phi}(v) = \hat{\sigma}_{\phi(v)}(v)$. Finally, since f_{cut} is submodular, it follows that \tilde{f}_{cut} is submodular in its first argument for any realization. We conclude that \tilde{f}_{cut} satisfies the properties described above, hence it is (non-monotone) adaptive submodular. It is also non-negative, because f_{cut} is non-negative.

5 Experiments

We have evaluated our proposed algorithm on the two objective functions described in the previous section, namely influence maximization and maximum cut, on a few real-world data sets. Since we have no competitor policy for the adaptive non-monotone submodular setting, we rather focus here on showcasing the potential benefits of adaptivity by comparing adaptive to non-adaptive random greedy on these two objectives.

5.1 Data Sets and Experimental Setup

For our experiments, we used networks from the KONECT² database, which accumulates network data sets from various other sources. The network sizes range from a few thousands to tens of thousands nodes.

Computing the (expected) marginal gains is at the heart of both the non-adaptive random greedy algorithm and our proposed adaptive random greedy policy. In terms of computational complexity, these computations may range from being completely straightforward to extremely demanding, depending on the specific objective at hand. For the influence optimization objective, the exact computation of the expected influence of a subset of nodes has been shown to be NP-hard [Kempe *et al.*, 2003]. To obtain an estimate we use Monte Carlo sampling over the outcomes of the independent cascades. Note that within adaptive random greedy we have to perform this simulation at every step, while conditioning on the observations obtained up to that step. For the maximum cut objective, it is very simple to compute the marginal gains in the adaptive case, since we already know at each step the outcomes of the previously selected nodes. In the non-adaptive setting, however, this is considerably harder, since we have to average over every possible outcome of the current selection; we again resort to sampling from the space of possible realizations to obtain estimates. To make the aforementioned computations more efficient, we subsample each network down to 2000 nodes, using a technique based on random walks proposed by Leskovec and Faloutsos [2006].

For both objectives, we select uniformly at random a subset of 100 nodes as the ground set E , and repeat the experiments for 50 such random ground sets. For each ground set instance, we evaluate the algorithms on 100 random realizations.

5.2 Results

We present here results for three data sets that represent ego networks from Facebook, Google+, and Twitter respectively

²<http://konect.uni-koblenz.de/>




Figure 2: Improvement in expected utility of using adaptive compared to non-adaptive random greedy for varying node budget k . (a)–(c) influence maximization; (d)–(f) maximum cut.

[McAuley and Leskovec, 2012]. Figure 2 shows the % relative improvement of adaptive random greedy over its non-adaptive counterpart in terms of expected utility for influence maximization (top) and maximum cut (bottom); each plot shows the improvement for varying values of the cardinality constraint k . For the influence maximization objective, the influence propagation probability of each edge is chosen to be $p = 0.1$, and for the maximum cut objective, selecting a node cuts that node or one of its neighbors with equal probability.

We can see that adaptivity is beneficial in general, while the improvement it provides varies substantially depending on the properties of each network. As an example, for networks containing a few nodes of very high degree, like the Google+ network in plots (b) and (e), adaptivity provides little benefit for influence maximization, since these nodes are the main source of influence, hence are almost always selected by the non-adaptive algorithm as well. On the other hand, adaptivity is much more beneficial for the maximum cut objective in such networks, since the feedback of whether such high degree nodes have already been cut by some of their neighbors helps making future selections more efficient.

Furthermore, if our goal is to reach a specific level of objective value using as few nodes as possible, then our gains due to adaptivity can be even more substantial in terms of the number of required nodes. For example, as shown in Figure 3(a), if we want to attain a maximum cut objective value of 1900 for the FACEBOOK network, a budget k of about 13 nodes is enough for adaptive random greedy, while a budget of almost 30 nodes is required for non-adaptive random greedy.

For the other two plots of Figure 3 we fix $k = 20$. In

plot (b) we show the improvement on FACEBOOK for varying edge probabilities p of the independent cascade model. At the extreme values of p , adaptivity provides no benefit, since the network is either disconnected ($p = 0$), or fully connected ($p = 1$). In plot (c) we show the improvement on TWITTER for varying cut distributions. The parameter β quantifies the probability of a node being cut when it is selected. A value of $\beta = 0$ corresponds to the setting we used in Figure 2, where the cutting probability is uniformly distributed among the selected node and each of its neighbors; $\beta = 1$ corresponds to deterministically cutting the selected node. We can see that, as the cutting distribution gets close to deterministic ($\beta \rightarrow 1$), the benefit of adaptivity diminishes.

Finally, we would like to comment on the behavior of the simple adaptive greedy algorithm with the additional modification to stop when the largest marginal gain becomes negative. In particular, for the specific objectives considered here, we have observed that its performance is very close to that of adaptive random greedy. This is presumably because both these objectives are approximately monotone for small values of k , and also fairly benign in the sense that they do not create traps that would severely diminish the performance of adaptive greedy. Intuitively, choosing one element cannot reduce the marginal gain of many other elements by a lot. However, even in the non-adaptive setting it is easy to come up with much harder non-monotone objectives for which simple greedy exhibits arbitrarily bad performance. The takeaway is that adaptive random greedy is comparable to adaptive greedy for the easier objectives that we have used here, while it also provides performance guarantees for the harder ones, a behavior that is completely analogous to how greedy vs. random greedy work in the non-adaptive setting.






Figure 3: (a) The budget k required by adaptive random greedy to reach a certain objective value (here 1900) is considerably smaller compared to its non-adaptive counterpart; (b) Improvement vs. independent cascade edge probability p ; (c) improvement vs. cut distribution parameter β .

6 Related Work

Compared to monotone submodular maximization, for which the $(1 - 1/e)$ -approximation of the greedy algorithm was shown by Nemhauser *et al.* [1978], constant-factor approximations for non-monotone submodular functions have been much more recent, for both the unconstrained case [Feige *et al.*, 2007], as well as under matroid and knapsack constraints [Lee *et al.*, 2009; Chekuri *et al.*, 2011]. Even more recently, Buchbinder *et al.* [2014] introduced the random greedy algorithm for maximizing non-monotone submodular functions under a cardinality constraint, from which we drew inspiration for our proposed adaptive random greedy policy.

The concepts of adaptive monotonicity and adaptive submodularity were introduced by Golovin and Krause [2011], who also showed that the greedy policy provides a $(1 - 1/e)$ -approximation under these assumptions. Example application domains, apart from those we present in this paper, include active learning [Chen *et al.*, 2014; 2015], interactive set coverage [Guillory and Bilmes, 2010], and incentive mechanism design [Singla and Krause, 2013].

The problem of influence maximization was originally proposed by Kempe *et al.* [2003] and was extended to the adaptive setting by Golovin and Krause [2011]. Various techniques have been proposed to make the computation of marginal gains feasible for large-scale networks using, for instance, more efficient sampling methods [Ohsaka *et al.*, 2014], and sketching-based approximations [Cohen *et al.*, 2014]. In this paper we chose to run experiments on smaller-scale networks, but these techniques could be applied to scale up adaptive random greedy as well. He and Kempe [2014] re-

cently considered the problem of assessing the robustness of influence maximization algorithms under network parameter misspecification, which interestingly leads to maximizing a non-monotone submodular objective.

Maximum graph cut has been a much-studied NP-complete problem with constant-factor SDP-based approximation algorithms for both the unconstrained [Goemans and Williamson, 1995] and cardinality-constrained [Feige and Langberg, 2001] cases. An interesting application of maximum cut objectives has been proposed by Lin and Bilmes [2010] and Lin and Bilmes [2011] for text summarization.

7 Conclusion

We proposed the adaptive random greedy policy for adaptive submodular maximization, the first policy with provable approximation guarantees for non-monotone objectives. We also presented two simple ways of constructing non-monotone objectives in practice, and observed the advantage of adaptivity by evaluating our policy on two network-related functions obtained this way. We believe that our work is a step towards understanding the class of functions amenable to adaptive optimization, and hope that it will encourage the broader use of non-monotone objectives in modeling and solving practical AI problems.

Acknowledgments

This work was funded in part by ERC StG 307036 and a Microsoft Research Faculty Fellowship.

References

- [Buchbinder *et al.*, 2014] Niv Buchbinder, Moran Feldman, Joseph Naor, and Roy Schwartz. Submodular maximization with cardinality constraints. In *Symposium on Discrete Algorithms (SODA)*, 2014.
- [Chekuri *et al.*, 2011] Chandra Chekuri, Jan Vondrak, and Rico Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. In *Symposium on Theory of Computing (STOC)*, 2011.
- [Chen *et al.*, 2014] Yuxin Chen, Hiroaki Shioi, Cesar Antonio Fuentes Montesinos, Lian Pin Koh, Serge Wich, and Andreas Krause. Active detection via adaptive submodularity. In *International Conference on Machine Learning (ICML)*, 2014.
- [Chen *et al.*, 2015] Yuxin Chen, Shervin Javdani, Amin Karbasi, J. Andrew Bagnell, Siddhartha Srinivasa, and Andreas Krause. Submodular surrogates for value of information. In *Conference on Artificial Intelligence (AAAI)*, 2015.
- [Cohen *et al.*, 2014] Edith Cohen, Daniel Delling, Thomas Pajor, and Renato Werneck. Sketch-based influence maximization and computation: Scaling up with guarantees. In *International Conference on Information and Knowledge Management (CIKM)*, 2014.
- [Feige and Langberg, 2001] Uriel Feige and Michael Langberg. Approximation algorithms for maximization problems arising in graph partitioning. *Journal of Algorithms*, 2001.
- [Feige *et al.*, 2007] Uriel Feige, Vahab S. Mirrokni, and Jan Vondrak. Maximizing non-monotone submodular functions. In *Symposium on Foundations of Computer Science (FOCS)*, 2007.
- [Goemans and Williamson, 1995] Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the Association for Computing Machinery (JACM)*, 1995.
- [Golovin and Krause, 2011] Daniel Golovin and Andreas Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *Journal of Artificial Intelligence Research (JAIR)*, 2011.
- [Guillory and Bilmes, 2010] Andrew Guillory and Jeff Bilmes. Interactive submodular set cover. In *International Conference on Machine Learning (ICML)*, 2010.
- [He and Kempe, 2014] Xinran He and David Kempe. Stability of influence maximization. In *Conference on Knowledge Discovery and Data Mining (KDD)*, 2014.
- [Kempe *et al.*, 2003] David Kempe, Jon Kleinberg, and Eva Tardos. Maximizing the spread of influence through a social network. In *Conference on Knowledge Discovery and Data Mining (KDD)*, 2003.
- [Lee *et al.*, 2009] Jon Lee, Vahab S. Mirrokni, Viswanath Nagarajan, and Maxim Sviridenko. Non-monotone submodular maximization under matroid and knapsack constraints. In *Symposium on the Theory of Computing (STOC)*, 2009.
- [Leskovec and Faloutsos, 2006] Jure Leskovec and Christos Faloutsos. Sampling from large graphs. In *Conference on Knowledge Discovery and Data Mining (KDD)*, 2006.
- [Lin and Bilmes, 2010] Hui Lin and Jeff Bilmes. Multi-document summarization via budgeted maximization of submodular functions. In *Human Language Technologies (HLT)*, 2010.
- [Lin and Bilmes, 2011] Hui Lin and Jeff Bilmes. A class of submodular functions for document summarization. In *Human Language Technologies (HLT)*, 2011.
- [McAuley and Leskovec, 2012] Julian McAuley and Jure Leskovec. Learning to discover social circles in ego networks. In *Conference on Neural Information Processing Systems (NIPS)*, 2012.
- [Nemhauser *et al.*, 1978] George L. Nemhauser, Laurence A. Wolsey, and Marshall L. Fisher. An analysis of approximations for maximizing submodular set functions. *Mathematical Programming*, 1978.
- [Ohsaka *et al.*, 2014] Naoto Ohsaka, Takuya Akiba, Yuichi Yoshida, and Ken-ichi Kawarabayashi. Fast and accurate influence maximization on large networks with pruned monte-carlo simulations. In *Conference on Artificial Intelligence (AAAI)*, 2014.
- [Singla and Krause, 2013] Adish Singla and Andreas Krause. Incentives for privacy tradeoff in community sensing. In *Conference on Artificial Intelligence (AAAI)*, 2013.