

Maximizing Stochastic Monotone Submodular Functions

Arash Asadpour*

Hamid Nazerzadeh†

Amin Saberi*

Abstract

We study the problem of maximizing a stochastic monotone submodular function with respect to a matroid constraint. Our model can capture the effect of uncertainty in different problems, such as cascade effects in social networks, capital budgeting, sensor placement, etc. We study the adaptivity gap - the ratio between the values of optimal adaptive and non-adaptive policies - and show that it is equal to $\frac{e}{e-1}$. This result implies that the benefit of adaptivity is bounded.

We also study the myopic policy and show that it is a $\frac{1}{2}$ -approximation. Furthermore, when the matroid is uniform, approximation ratio of the myopic policy becomes $1 - \frac{1}{e}$ which is optimum.

1 Introduction

The problem of maximizing submodular functions has been extensively studied in operations research and computer science. For a set \mathcal{A} , the set function $f : 2^{\mathcal{A}} \rightarrow \mathbb{R}$ is submodular if for any two subsets $S, T \subset \mathcal{A}$ we have

$$f(S \cup T) + f(S \cap T) \leq f(S) + f(T)$$

An equivalent definition is that the inequality below holds for any $S \subseteq T \subseteq \mathcal{A}$ and $j \in \mathcal{A}$

$$f(T + j) - f(T) \leq f(S + j) - f(S)$$

where $f(\cdot + j)$ denotes $f(\cdot \cup \{j\})$. Also, function f is monotone if for any two subsets $S \subseteq T \subseteq \mathcal{A}$:

$$f(S) \leq f(T)$$

A wide range of optimization problem that arise in the real world can be modeled as maximizing a monotone submodular functions with respect to some constraints. One instance is the welfare maximization problem [9, 11, 23] which is to find an optimal allocation of resources to agents where the utilities of the agents are submodular. Submodularity corresponds to the law of diminishing return in economy.

Another application of this problem is capital budgeting in which a risk-averse investor with a limited budget is interested in finding the optimal investment in different projects [24, 2]. The utility function of a risk averse investor is submodular. It is also naturally non-negative and monotone.

Another example is the problem of viral marketing and maximizing influence through the network [14, 18], where the goal is to choose an initial “active” set of people, so as to maximize the spread of a technology or behavior in a social network. It is well-known that under many models of influence propagation in networks (e.g., cascade model [14]), the expected size of the final cascade is a submodular function of the set of initially activated individuals. Also, due to budget limitations, the number of people that we can activate in the beginning is bounded. Hence, the maximizing influence problem can be seen as a maximizing submodular function problem subject to cardinality constraints.

Yet another example is the problem of optimal placement of sensors for environmental monitoring [16, 17] where the objective is to place sensors in the environment in order to most effectively reduce uncertainty

*Management Science and Engineering Department, Stanford University, Stanford, CA. {asadpour, saberi}@stanford.edu

†Microsoft Research, Cambridge, MA. hamidnz@microsoft.com

in observations. This problem can be modeled by entropy minimization and, due to the concavity of the entropy function, it is a special case of submodular optimization.

For the above problems and many others, the constraints can be model by a matroid. A finite *matroid* \mathcal{M} is defined by a pair $(\mathcal{A}, \mathcal{I})$, where \mathcal{I} is a collection of subsets of \mathcal{A} (called the *independent sets*) with the following properties:

1. Every subset of an independent set is independent.
2. If S and T are two independent sets and T has more elements than S , then there exists an element in T which is not in S and when added to S still gives an independent set.

Two important special cases are *uniform matroid* and *partition matroid*. In a uniform matroid, all the subsets of \mathcal{A} of size at most k , for a given k , are independent. Uniform matroids represent cardinality constraints. A partition matroid is defined over a partition of set \mathcal{A} , where every independent set includes at most one element from each set in the partition.

The celebrated result of Nemhauser et al. [20] shows that for maximizing nonnegative monotone submodular functions over uniform matroids, the greedy algorithm gives a $(1 - \frac{1}{e} \approx 0.632)$ -approximation of the optimal solution. Later, they showed that for optimizing over matroids, the approximation ratio of the greedy algorithm is $\frac{1}{2}$. Recently, Calinescu et al [5] proposed a better approximation algorithm with ratio $1 - \frac{1}{e}$. It also has been shown that this factor is optimal (in the value oracle model), if only a polynomial number of queries is allowed [19, 10].

However, these algorithms are designed for deterministic environments. In practice, one must deal with the stochasticity caused by the uncertain nature of the problem, the incomplete information about the environment, etc. For instance, in welfare maximization, the quality of the resources may be unknown in advance, or in the capital budgeting problem some projects taken by an investor may fail due to unexpected events in the market. As another example, in viral marketing some people in the initial set might not adopt the behavior. Also, in the environmental monitoring example, it is expected that a non-negligible fraction of sensors might not work properly for various reasons.

All these possibilities motivate the problem of stochastic submodular maximization. In the stochastic setting, the outcome of the elements in the selected set are not known in advance and they will be only discovered after they are chosen.

1.1 Problem Definition

In defining the problem, we need to use some care to maintain generality. Consider a set $\mathcal{A} = \{X_1, \dots, X_n\}$ of n *independent* random variables over a domain Δ . The domain varies depending on the application. For instance, in the welfare maximization problem, X_i denotes the quality of the resource or in viral marketing X_i corresponds to the set of people who are influenced by person i . The distribution of each X_i is potentially different and is given by a function g_i .

Let x_i denote a realization of X_i . Also, let vector $s = \langle \hat{x}_1, \dots, \hat{x}_n \rangle$ denote a *realization* of set $S \subset \mathcal{A}$, where $\hat{x}_i = x_i$ for $X_i \in S$ and $\hat{x}_i = 0$ for $i \notin S$. For a given function $f : \Delta^n \rightarrow \mathbb{R}^+$, we can define the stochastic function $F : \mathcal{A} \rightarrow \mathbb{R}^+$ as $F(S) = \mathbb{E}[f(s)]$, where s is a realization of S and the expectation is taken with respect to the product distribution defined by g_i 's.

Also, consider a subset $T \in \mathcal{A}$, and a realization t of T . We can define a conditional expectation $\mathbb{E}[f(s)|t]$. In the distribution imposed by conditioning on t , $s_i = t_i$ if its corresponding random variable is in $S \cap T$. Otherwise s_i is chosen independently with respect to the distribution defined by g_i 's. Let us denote this conditional expectation with $F(S, t)$.

We call the set function F *stochastic monotone submodular* if $F(\cdot, t)$ is monotone submodular for every t . Observe that if f is monotone submodular then F is stochastic monotone submodular because it is a convex combination of monotone submodular functions.

Remark: We assume that either we can compute the value of function F up to a desired degree of accuracy explicitly, or F is given to us via an ‘‘oracle’’. This is a natural assumption for all the applications mentioned

in the paper. In fact, in most cases the expectations can be computed simply by using sampling. For example, sampling works when the probability distribution functions are constant Lipschitz continuous, or when their support is a polynomial size set of discrete values. In both cases, a small $o(1)$ error is introduced in the calculations that we ignore in the rest of the paper.

Definition: [*Maximizing a stochastic monotone submodular function*] The set $\mathcal{A} = \{X_1, \dots, X_n\}$ of n independent random variables, the matroid $\mathcal{M} = (\mathcal{A}, \mathcal{I})$, and the stochastic monotone submodular set function $F : 2^{\mathcal{A}} \rightarrow \mathbb{R}^+$ are given. Find a subset $S \in \mathcal{I}$ that maximizes F , i.e., $\max_{S \in \mathcal{I}} \mathbb{E}[F(S)]$ where the expectation is taken over the probability distribution of the sets chosen by the policy.

A special case of the above problem is the *stochastic max k -cover problem* which is defined as follows. Suppose a collection \mathcal{A} of random subsets of $N = \{1, 2, \dots, m\}$ are given. Each element $X_i \in \mathcal{A}$ is a random subset of N , and its distribution is denoted by a probability distribution g_i . In the stochastic maximum k -cover problem, the goal is to choose k elements of \mathcal{A} such that their union has the maximum cardinality. We discuss this problem in more detail in Section 2.1.

For the problem of maximizing a stochastic monotone submodular function, we study two types of policies: adaptive and non-adaptive. A *non-adaptive policy* is represented by a fixed subset of \mathcal{A} . An *adaptive policy* is a decision tree. It assumes that the value of each random variable can be observed as soon as it is chosen and it uses the observed values of the previously chosen elements to determine the next element in the subset.

We compare these policies by studying the adaptivity gap of the problem. The *adaptivity gap* is defined as the ratio between the expected values of optimal adaptive and non-adaptive policies. Adaptivity gap has been previously studied for stochastic maximization problems with respect to covering [13] and packing [7, 8] constraints.

1.2 Results

We present approximately optimal policies for the stochastic monotone submodular maximization problem. First, in Section 2, we compare the performance of the optimal adaptive and non-adaptive policies. Although non-adaptive policies may not perform as well as adaptive ones, they are particularly useful when it is difficult or time consuming to discover the outcome of an element. For example, in the capital budgeting problem, it is not possible for the investor to wait until the end of each project to measure the success, or in the environmental monitoring problem, it is not practical to measure the performance of sensors after placing each sensor in the environment.

Surprisingly, we learn that the adaptivity gap of the problem is equal to $\frac{e}{e-1} \approx 1.59$. In other words, there exists a non-adaptive policy which achieves at least $\frac{e-1}{e}$ fraction of the value of best adaptive policy. This result leads to a $(\frac{e-1}{e})^2 \approx 40\%$ approximation of the optimal adaptive policy by a non-adaptive policy that runs in polynomial time in n . We also give an example to show that our analysis of the adaptivity gap is tight. For that, we use a simple instance of the stochastic max k -cover problem.

In Section 3, we focus on natural myopic policies. We study the natural extension of the myopic policy studied in [5] in a stochastic environment. This policy iteratively chooses an element with the maximum *expected* marginal value, conditioned on the outcome of the previous elements.

We show that the approximation ratio of this policy with respect to the optimal adaptive policy is $\frac{1}{2}$ for *general matroids*. We also prove that over uniform matroid (i.e., subject to a cardinality constraint), the approximation ratio of this policy is $1 - \frac{1}{e}$.¹ Due to the results of [19, 10], the approximation ratio of $1 - \frac{1}{e}$ is optimal only if a polynomial number of oracle accesses is allowed.

The closest work to ours in the literature is by Chan and Farias [4]. They mainly study the problem of stochastic submodular optimization over partition matroids. In their model, there is an ordering over the partitions and any adaptive policy has to choose one element from each partition according to the given order. They present a $\frac{1}{2}$ -approximation of the optimal adaptive policy (that respects the ordering) using a myopic policy. In our setting, we do not have a fixed ordering. In addition, we prove most of our results for general matroids.

¹ The results for the uniform matroid has appeared in a preliminary version of this work [3].

2 The Adaptivity Gap of Stochastic Submodular Optimization Problem

In this section, we analyze the optimal adaptive and non-adaptive policies and compare the performance of the two. First, observe that since non-adaptive policies do not observe the realized value of the items until the end, they may choose all the elements in one step. In other words, any non-adaptive policy can be represented by the set of chosen elements.

On the other hand, an optimal adaptive policy selects the elements based on the realized values of the previously chosen elements. Note that the policy knows the probability distribution of the values of the elements that are not yet chosen, but not their actual values.

Although an adaptive policy can clearly perform better than a non-adaptive policy, we show that its advantage is limited. The main result of this section is as follows:

Theorem 1 *The adaptivity gap of the stochastic monotone submodular maximization problem is equal to $\frac{e}{e-1}$.*

In order to prove the above theorem, we start by establishing an upper bound on the adaptivity gap. In Section 2.1, we give an example that shows our analysis of the adaptivity gap is tight.

Before proving the theorem, observe that since $F(S)$ is a submodular function, we can use the following result of Calinescu et. al. [5]:

Theorem 2 (Calinescu et al [5]) *Given oracle access to F (see Remark 1), there exists a polynomial time algorithm that achieves an approximation ratio of $1 - \frac{1}{e} - o(1)$.*

The above theorem immediately implies that

Corollary 3 *A $(1 - \frac{1}{e} - o(1))$ -approximation of the optimal non-adaptive policy can be computed in polynomial time.*

Theorem 1 and the above corollary imply that:

Corollary 4 *There is a policy that is non-adaptive and also runs in polynomial time and computes a solution that is within $(\frac{e-1}{e})^2$ of the optimal adaptive policy.*

In the rest of this section, we prove Theorem 1. The proof is inspired by the techniques developed in Section 3.5 of [21] for submodular optimization (*in a non-stochastic setting*). For the sake of consistency, we use the same notation as [21] wherever possible.

We start by making a few observations about adaptive policies. First, any adaptive policy can be described by a (possibly randomized) decision tree in which at each step an element is being added to the current selection. Consider an arbitrary adaptive policy ADAPT. Each path from the root to a leaf of this tree corresponds to a realization $s \in \hat{I}$ of the *sequence* of elements chosen by ADAPT. Here, \hat{I} denotes the set of all possible realizations of sets in \mathcal{I} . Let $y = \langle y_1, \dots, y_n \rangle$ represent the probability that each element of \mathcal{A} is chosen by ADAPT, i.e., y_i is the probability of choosing X_i . These probabilities sum up to 1. Also, let β_s denote the probability density function for outcome $s \in \hat{I}$. Then, we have the following properties:

1. $\int_{s \in \hat{I}} \beta_s = 1$.
2. $\forall s : \beta_s \geq 0$.
3. $\forall i, dx : \int_{s, s_i \in dx_i} \beta_s ds = y_i g_i(x_i) dx_i$

The first two properties hold because β defines a probability measure on the space of all feasible outcomes. The third property implies that the probability that we observe outcome x_i (a realized value of X_i) among all possible outcome s is equal to the probability that X_i is chosen (i.e., y_i) multiplied by the probability that the outcome is equal to x_i . This property holds because of the independence among the random variables. Since every policy satisfies the above properties, we can establish an upper bound on the value of any adaptive policy. Hence, we define the function $f^+ : [0, 1]^n \rightarrow \mathbb{R}$ as follows:

$$f^+(y) = \sup_{\alpha} \left\{ \int_{s \in \hat{I}} \alpha_s f(s) : \int_s \alpha_s = 1, \alpha_s \geq 0, \forall i, dx_i : \int_{s, s_i \in dx_i} \alpha_s ds = y_i g_i(x_i) dx_i \right\}. \quad (1)$$

Another observation is that for an optimal adaptive policy, vector y described above is in the base polytope of \mathcal{M} (defined as follows). A set $S \in \mathcal{I}$ is called a basis for the matroid if $|S| = \max\{|T| : T \in \mathcal{I}\}$. The base polytope, $B(\mathcal{M})$, is defined as:

$$B(\mathcal{M}) = \text{conv}\{1_S | S \in \mathcal{I}, S \text{ is a basis}\}$$

Here ‘‘conv’’ denotes the convex hull and 1_S is the characteristic vector of S , i.e., 1 for elements in S and 0 for other elements.

Lemma 5 *The expected value of the optimal adaptive policy is at most $\max_{y \in B(\mathcal{M})} \{f^+(y)\}$,*

Proof : Note that an optimal adaptive policy only chooses independent sets. Due to monotonicity, all of these are independent sets are bases of the matroid. Hence, for an optimal adaptive policy vector y defined above is in $B(\mathcal{M})$. Moreover, the expected value of the adaptive policy is bounded by $f^+(y)$, because the policy has to satisfy the 3 properties mentioned earlier. \square

Now, we define an extension of set function $F(S)$ to the domain of real numbers. For vector $y \in [0, 1]^n$, let Y denoted a random set where Y includes $X_i \in \mathcal{A}$ with probability y_i . With abuse of notation, we define the extension $F : [0, 1]^n \rightarrow \mathbb{R}^+$ as follows:

$$F(y) = \mathbb{E}[F(Y)] = \sum_{Y \text{ is a basis of } \mathcal{I}} \left(\prod_{i \in Y} y_i \prod_{i \notin Y} (1 - y_i) \right) F(Y).$$

Function $f^+(y)$ sets an upper bound on the adaptive policies. We now establish a lower bound on the value of optimal non-adaptive policies via the following lemma from [21] (Lemma 3.4), which is based on pipage rounding [1].

Lemma 6 [21] *Any vector $y \in B(\mathcal{M})$ can be rounded to an integral solution $S \in \mathcal{I}$ of value $F(S) \geq F(y)$.*

To complete the proof we need to show that for any vector y , the values of $F(y)$ and $f^+(y)$ are within a constant factor of each other, which is established by combining Lemmas 3.7 and 3.8 from [21].

Lemma 7 [21] *For any monotone submodular function f and any vector y we have*

$$f^+(y) \leq \left(\frac{e}{e-1}\right) F(y)$$

Proof : [Theorem 1] Lemma 5 shows that $\max_{y \in B(\mathcal{M})} f^+(y)$ is an upper-bound on the performance of the optimal adaptive policy. Consider $y^* \in \text{argmax}_{y \in B(\mathcal{M})} f^+(y)$. By Lemma 7, we have $F(y^*)$ is at least a $(1 - \frac{1}{e})$ fraction of the expected value of an optimal adaptive policy. On the other hand, Lemma 6 implies that there exists a $S \in \mathcal{I}$ such that $F(S) \geq F(y^*)$. Note that $F(S)$ is in fact the expected value gained by a non-adaptive policy that selects set S . Hence, S is a $(1 - \frac{1}{e})$ -approximation of the optimal adaptive policy. By Proposition 10 in the next section, this factor is tight. \square

2.1 A Tight Example: Stochastic Maximum k -Cover

Given a collection \mathcal{A} of the subsets of $N = \{1, 2, \dots, n\}$, the goal of the *max k -cover* problem is to find k subsets from \mathcal{A} such that their union has the maximum cardinality [10]. In the stochastic version, the subset that an element of \mathcal{A} would cover is revealed only after choosing the element, according to a given probability distribution.

The following reduction shows that this problem is a special case of the stochastic monotone submodular maximization. For $S \in \mathcal{A}$, let $F(S)$ denote the expected number of elements covered by the subsets in S . Clearly, F is monotone. Consider two subsets $S \subseteq T \subseteq \mathcal{A}$, an element $X \in \mathcal{A}$, and a realization y of an arbitrary subset of \mathcal{A} . Because $\cup_{A \in S} A \subseteq \cup_{B \in T} B$, for every realization y , we have $F(S + X) - F(S) \geq F(T + X) - F(T)$. In addition, $\mathcal{M} = (\mathcal{A}, \{S \subseteq \mathcal{A} : |S| \leq k\})$ forms a uniform matroid. Therefore, the stochastic max k -cover problem is in fact a stochastic monotone submodular maximization problem.

In this section, we define an instance of stochastic max k -cover problem that gives a lower bound on the adaptivity gap. This example has been brought to our attention by Vondrák [22].

Consider the following instance: a ground set $N = \{1, 2, \dots, n\}$ and a collection $\mathcal{A} = \{X_j^{(i)} \mid 1 \leq i \leq n, 1 \leq j \leq n^2\}$ of its subsets are given. For every i, j , define $X_j^{(i)}$ to be the one-element subset $\{i\}$ with probability $\frac{1}{n}$ and the empty set with probability $1 - \frac{1}{n}$. The goal is to cover the maximum number of the elements of N by selecting at most $k = n^2$ subsets from \mathcal{A} .

Lemma 8 *The optimal non-adaptive policy is to pick n subsets from each of the collections $\mathcal{A}^{(i)} = \{X_j^{(i)} \mid 1 \leq j \leq n^2\}$ for every i . For large enough values of n , the expected value of this policy is (arbitrarily close to) $(1 - \frac{1}{e})n$.*

Proof : Consider an arbitrary non-adaptive policy which picks S , containing n^2 sets from \mathcal{A} . For each i , define $k_i = |S \cap \mathcal{A}^{(i)}|$. Moreover, each element $i \in N$ is covered if and only if at least one of its corresponding chosen subsets are realized as a non-empty subset. Hence, it will be covered with probability $1 - (1 - \frac{1}{n})^{k_i}$. Therefore, the expected value of this policy is $\sum_i 1 - (1 - \frac{1}{n})^{k_i}$. Note that $1 - (1 - \frac{1}{n})^x$ is a concave function with respect to x , and also $\sum_i k_i = n^2$. Hence, the expected value of the policy is maximized when $k_1 = k_2 = \dots = k_n = n$. In this case, the expected value is $(1 - (1 - \frac{1}{n})^n)n \approx (1 - \frac{1}{e})n$ for large n . \square

We now consider the following myopic adaptive policy \mathcal{P} : Start with $i = 1$ and pick the elements of $\mathcal{A}^{(i)}$ one by one until one of them is realized as $\{i\}$ or all of elements in $\mathcal{A}^{(i)}$ are chosen. Then increase i by one. Continue the iteration until $i = n + 1$.

The following lemma gives a lower bound on the number of elements in N covered by the adaptive policy.

Lemma 9 *The expected number of elements in N covered by \mathcal{P} described above is $(1 - o(1))n$.*

Proof : Let X_k be the indicator random variable corresponding to the event that the subset chosen at the k -th step is realized as a non-empty subset for any $1 \leq k \leq n^2$. Note that the number of elements covered by \mathcal{P} is $\sum_{k=1}^{n^2} X_k$. Moreover, all X_k 's are independent random variables.

By the description of \mathcal{P} , as long as $\sum_{i=1}^k X_k < n^2$, X_k will be one with probability $\frac{1}{n}$ and will be zero with probability $1 - \frac{1}{n}$. Also, when $\sum_{i=1}^t X_k = n$, we have already covered all the elements in N . Therefore, X_{t+1}, \dots, X_{n^2} will all be equal to zero. With this observation, we define i.i.d random variables Y_1, Y_2, \dots, Y_{n^2} , where each Y_i is set to be one with probability $\frac{1}{n}$ and zero with probability $\frac{n-1}{n}$. Observe that $\min\{n, Y = \sum_k Y_k\}$ has the same probability distribution as $\sum_k X_k$. Note that $E[Y] = n$. Using Chernoff bound, we have

$$\Pr[Y \leq n - n^{2/3}] \leq e^{-\frac{n^{4/3}}{2n}} = e^{-n^{1/3}}.$$

Thus, with probability at least $1 - e^{-n^{1/3}}$ we have $Y > n - n^{2/3}$. Hence,

$$E\left[\sum_{k=1}^{n^2} X_k\right] = E[\min\{n, Y\}] \geq (1 - e^{-n^{1/3}})(n - n^{2/3}) = n - o(n),$$

The myopic adaptive policy:Initialize $t = 0, S_0 = \emptyset, U_0 = \emptyset$ While $(\mathcal{A} \neq U_t \cup S_t)$ $t \leftarrow t + 1$ $S_t \leftarrow S_{t-1}$

Repeat

Select $X_i \in \operatorname{argmax}_{X_i \in \mathcal{A} \setminus (U_{t-1} \cup S_{t-1})} \mathbb{E}[F(S_{t-1} + X_i) | s_{t-1}]$ If $S_{t-1} \cup \{X_i\} \notin \mathcal{I}$ then $U_t \leftarrow U_{t-1} \cup \{X_i\}$

else

 $S_t \leftarrow S_{t-1} \cup \{X_i\}$ $U_t \leftarrow U_{t-1}$ Observe x_i and update s_t Until $(\mathcal{A} = U_{t-1} \cup S_{t-1})$ or $(S_t \neq S_{t-1})$

which completes the proof of the lemma. □

By combining the results of Lemmas 8 and 9 we have the following proposition:

Proposition 10 *For large enough n , the adaptivity gap of stochastic maximum coverage is at least $\frac{e}{e-1}$.*

3 Approximation Ratio of Simple Myopic Policies

In this section, we present an adaptive myopic policy with an approximation ratio of $\frac{1}{2}$ with respect to an optimal adaptive policy. In Section 3.1, we show that the myopic policy achieves the approximation ratio of $1 - \frac{1}{e}$ if the matroid is uniform. Note that even if the actual values were known, the problem of computing the optimal policy is intractable. As mentioned before, the maximum k -cover is a special case of our problem and Feige [10] has shown that it is not possible to find an approximation ratio better than $1 - \frac{1}{e}$ for the maximum k -cover problem, unless $NP \subset TIME(n^{O(\log \log n)})$.

The policy is given in the above figure. At each iteration, from the elements in \mathcal{A} that are not yet considered, the policy chooses an element with the maximum expected marginal value. We denote by S_t the set of elements chosen by the adaptive policy up to iteration t . Let s_t denote the realization of *all* these elements. Also, U_t is the set of elements considered but not chosen by the policy due to the matroid constraint. Here is the main result of this section.

Theorem 11 *For general matroids, the approximation ratio of the myopic adaptive policy with respect to any optimal adaptive policy is $\frac{1}{2}$.*

Define $\Delta_t = F(S_t) - F(S_{t-1})$. Also, let k be the number of elements chosen by the myopic policy (which is simply the rank of the matroid \mathcal{M}). The basic idea of the proof is similar to Fisher et al. [12]. But, the main difficulty is that the realized values of Δ_t are not always decreasing (due to the stochastic nature of the problem). In addition, the sequence of elements chosen by the optimal adaptive policy is random.

However, $\mathbb{E}[\Delta_t|s_{t-1}] \geq \mathbb{E}[\Delta_{t+1}|s_{t-1}]$ (Note that $\mathbb{E}[\Delta_t|s_t] \geq \mathbb{E}[\Delta_{t+1}|s_t]$ does not necessary hold). Based on this observation, we prove the theorem. We will also use the following lemma from [12].

Lemma 12 [12] For $t = 1, \dots, k$, we have $\sum_{i=1}^t |C_i| \leq t$.

Note that T , U , and S are random sets, but the lemma holds for every realization because it is a consequence of the matroid constraint, not the realizations of the element chosen by the policy.

We are now ready to prove the theorem.

Proof : [Theorem 11] Let P be the (random) set of elements chosen by the optimal adaptive policy. Also, for $t = 1, \dots, k$, define $C_t = P \cap (U_{t+1} \setminus U_t)$. Consider a realization s_t of S_t . Because F is stochastic monotone submodular we have

$$\mathbb{E}[F(P)|s_t] \leq \mathbb{E}\left[\sum_{l \in P \setminus S} F(S+l)|s_t\right]$$

The expectations, and in the rest of the proof, are taken over the probability distribution of all realizations of P such that the realized values of elements in $P \cap S_t$ are according to s_t . Since the above inequality holds for all s_t , we have

$$\begin{aligned} \mathbb{E}[F(P)] &\leq \mathbb{E}\left[\sum_{l \in P \setminus S} F(S+l)\right] \\ \mathbb{E}[F(P)] - \mathbb{E}[F(S)] &\leq \mathbb{E}\left[\sum_{l \in P \setminus S} (F(S+l) - F(S))\right] \end{aligned}$$

Note that $\bigcup_{t=1}^k C_t = P \setminus S$. Hence,

$$\mathbb{E}[F(P)] - \mathbb{E}[F(S)] \leq \sum_{t=1}^k \mathbb{E}\left[\sum_{l \in C_t} (F(S+l) - F(S))\right]$$

By expanding the expectation we have

$$\mathbb{E}[F(P)] - \mathbb{E}[F(S)] \leq \sum_{t=1}^k \int_{s_{t-1}: S_{t-1} \in \mathcal{I}} \mathbb{E}\left[\sum_{l \in C_t} F(S+l) - F(S) | s_{t-1}\right] \Pr[s_{t-1}] ds_{t-1} \quad (2)$$

Observe that conditioned on s_{t-1} , because the myopic policy chooses an element with the maximum marginal value, we have $\Delta_t \geq F(S+l) - F(S)$, $l \in C_t$. Therefore,

$$\mathbb{E}\left[\sum_{l \in C_t} F(S+l) - F(S) | s_{t-1}\right] \leq \mathbb{E}\left[\sum_{l \in C_t} \Delta_t | s_{t-1}\right]$$

By plugging the above inequality into (2), we get

$$\mathbb{E}[F(P)] - \mathbb{E}[F(S)] \leq \sum_{t=1}^k \int_{s_{t-1}: S_{t-1} \in \mathcal{I}} \mathbb{E}\left[\sum_{l \in C_t} \Delta_t | s_{t-1}\right] \Pr[s_{t-1}] ds_{t-1}$$

Using telescopic sums and the linearity of expectation we derive the following. Here Δ_{k+1} is defined 0.

$$\begin{aligned} \mathbb{E}[F(P)] - \mathbb{E}[F(S)] &\leq \sum_{t=1}^k \int_{s_{t-1}: S_{t-1} \in \mathcal{I}} \mathbb{E}\left[\sum_{l \in C_t} \sum_{j=t}^k (\Delta_j - \Delta_{j+1}) | s_{t-1}\right] \Pr[s_{t-1}] ds_{t-1} \\ &= \sum_{j=1}^k \sum_{t=1}^j \int_{s_{t-1}: S_{t-1} \in \mathcal{I}} \mathbb{E}\left[\sum_{l \in C_t} (\Delta_j - \Delta_{j+1}) | s_{t-1}\right] \Pr[s_{t-1}] ds_{t-1} \end{aligned}$$

Note that by using the Bayes' theorem and the law of total probability, for every t and j the integral term in the above is in fact equal to $\mathbb{E}[\sum_{l \in C_t} (\Delta_j - \Delta_{j+1})]$. Now, we can change the probability measure to calculate this expectation from s_{t-1} to s_{j-1} . Hence, we have

$$\begin{aligned} \mathbb{E}[F(P)] - \mathbb{E}[F(S)] &\leq \sum_{j=1}^k \sum_{t=1}^j \int_{s_{j-1}: S_{j-1} \in \mathcal{I}} \mathbb{E}[\sum_{l \in C_t} (\Delta_j - \Delta_{j+1}) | s_{j-1}] \Pr[s_{j-1}] ds_{j-1} \\ &= \sum_{j=1}^k \sum_{t=1}^j \int_{s_{j-1}: S_{j-1} \in \mathcal{I}} (E[|C_t| \mathbb{E}[\Delta_j - \Delta_{j+1} | s_{j-1}] | s_{j-1}]) \Pr[s_{j-1}] ds_{j-1} \end{aligned}$$

Note that conditioned on s_{j-1} , the term $\mathbb{E}[\Delta_j - \Delta_{j+1} | s_{j-1}]$ is by definition a constant and we can take it out from the outer expectation. Hence,

$$\mathbb{E}[F(P)] - \mathbb{E}[F(S)] \leq \sum_{j=1}^k \int_{s_{j-1}: S_{j-1} \in \mathcal{I}} \left(\mathbb{E}[\sum_{t=1}^j |C_t| | s_{j-1}] \mathbb{E}[\Delta_j - \Delta_{j+1} | s_{j-1}] \right) \Pr[s_{j-1}] ds_{j-1}$$

We now use Lemma 12 which implies that in every realization $\sum_{t=1}^j |C_t| \leq j$. We also use the fact that due to the submodularity and the rule of the policy, we have $\mathbb{E}[(\Delta_j - \Delta_{j+1}) | s_{j-1}] \geq 0$. We conclude that

$$\begin{aligned} \mathbb{E}[F(P)] - \mathbb{E}[F(S)] &\leq \sum_{j=1}^k \int_{s_{j-1}: S_{j-1} \in \mathcal{I}} j \mathbb{E}[(\Delta_j - \Delta_{j+1}) | s_{j-1}] \Pr[s_{j-1}] ds_{j-1} \\ &= \sum_{j=1}^k \int_{s_{j-1}: S_{j-1} \in \mathcal{I}} \mathbb{E}[\Delta_j | s_{j-1}] \Pr[s_{j-1}] ds_{j-1} \\ &= \sum_{j=1}^k \mathbb{E}[\Delta_j] \\ &= \mathbb{E}[F(S)] \end{aligned}$$

Therefore, $\mathbb{E}[F(P)] \leq 2\mathbb{E}[F(S)]$, as desired. \square

Fisher et al. [12] have shown that even in the non-stochastic setting, in the worst-case, the approximation ratio of the greedy algorithm (hence the myopic policy) is equal to $\frac{1}{2}$. Also, it is easy to see that if \mathcal{M} is an intersection of κ matroids, then the approximation ratio of the myopic policy is equal to $\frac{1}{1+\kappa}$.

3.1 Uniform Matroids

In this section we show that the myopic policy described in the previous section has a better approximation ratio if the matroid is uniform.

Theorem 13 *Consider the adaptive myopic policy that at each step selects an element with the maximum marginal value, conditioned on the realized value of the previously chosen elements. Over uniform matroids, the approximation ratio of this policy compared to the optimal adaptive policy is $1 - \frac{1}{e}$.*

The proof presented here is similar to the proof of Kleinberg et al.[15] for submodular set functions. The main technical difficulty in our case is that the optimal adaptive policy here is a random set whose distribution depends on the realized values of the elements of \mathcal{A} .

Proof : Let P denote the (random) set chosen by an optimal adaptive policy. Also, denote the marginal value of the t -th element chosen by the myopic policy by Δ_t , i.e.,

$$\Delta_t = F(S_t) - F(S_{t-1})$$

Consider a realization s_t of S_t . Because F is stochastic monotone submodular, we have

$$\mathbb{E}[F(P)|s_t] \leq \mathbb{E}[F(P \cup S_t)|s_t] \leq \mathbb{E}[F(S_t) + \sum_{l \in P} (F(S_t + l) - F(S_t))|s_t]. \quad (3)$$

The above expectations are taken over all realization of P such that the realized values of elements in $S_t \cap P$ are according to s_t . Because the myopic policy chooses the element with maximum marginal value, for every j , $1 \leq j \leq k$ (k is the rank of \mathcal{M}), we have

$$\mathbb{E}[\Delta_{t+1}|s_t] \geq \mathbb{E}[F(S_t + l) - F(S_t)|s_t]$$

Therefore, we get

$$\mathbb{E}[F(P)|s_t] \leq \mathbb{E}[F(S_t) + k\Delta_{t+1}|s_t]$$

Since the above inequality holds for every possible path in the history, by adding up all such inequalities for all t , $0 \leq t \leq k-1$, we have:

$$\begin{aligned} \mathbb{E}[F(P)] &\leq \mathbb{E}[F(S_t)] + k\mathbb{E}[\Delta_{t+1}] \\ &= \mathbb{E}[\Delta_1 + \dots + \Delta_t] + k\mathbb{E}[\Delta_{t+1}] \end{aligned}$$

We multiply the t -th inequality, $0 \leq t \leq k-1$, by $(1 - \frac{1}{k})^{k-1-t}$, and add them all up. The sum of the coefficients of $\mathbb{E}[F(P)]$ is equal to

$$\sum_{t=0}^{k-1} (1 - \frac{1}{k})^{k-1-t} = \sum_{t=0}^{k-1} (1 - \frac{1}{k})^t = \frac{1 - (1 - \frac{1}{k})^k}{1 - (1 - \frac{1}{k})} = k(1 - (1 - \frac{1}{k})^k) \quad (4)$$

On the right hand side, the sum of the coefficients corresponding to the term $\mathbb{E}[\Delta_t]$, $1 \leq t \leq k$, is equal to

$$\begin{aligned} k(1 - \frac{1}{k})^{k-t} + \sum_{j=t}^{k-1} (1 - \frac{1}{k})^{k-1-j} &= k(1 - \frac{1}{k})^{k-t} + \sum_{j=0}^{k-t-1} (1 - \frac{1}{k})^j \\ &= k(1 - \frac{1}{k})^{k-t} + k(1 - (1 - \frac{1}{k})^{k-t}) \\ &= k \end{aligned} \quad (5)$$

Thus, by inequalities (4) and (5) we conclude

$$(1 - (1 - \frac{1}{k})^k)\mathbb{E}[F(P)] \leq \sum_{t=1}^k \mathbb{E}[\Delta_t] = \mathbb{E}[F(S_k)].$$

Hence, the approximation ratio of the myopic policy is at least $1 - \frac{1}{e}$. \square

Acknowledgement We would like to thank Jan Vondrák for fruitful discussions.

References

- [1] A. Ageev and M. Sviridenko. Pipage rounding: A new method of constructing algorithms with proven performance guarantee. *J. Comb. Optim.*, 8(3):307–328, 2004.
- [2] S. Ahmed and A. Atamtürk. Maximizing a class of submodular utility functions. Research Report BCOL.08.02, IEOR, University of California-Berkeley, March 2008.

- [3] A. Asadpour, H. Nazerzadeh, and A. Saberi. Stochastic Submodular Maximization. *Workshop on Internet and Network Economics*, 477-489, 2008.
- [4] C. Chan and V. Farias. Stochastic depletion problems: Effective myopic policies for a class of dynamic optimization problems. *Mathematics of Operations Research*, 2009.
- [5] G. Calinescu, C. Chekuri, M. Pal and J. Vondrák. Maximizing a submodular set function subject to a matroid constraint. *to appear in SIAM Journal on Computing*.
- [6] G. Cornuejols, M. Fisher, and G. Nemhauser. Location of bank accounts to optimize float. *Management Science*, 23:789–810, 1977.
- [7] B. C. Dean, M. X. Goemans, and J. Vondrák. Approximating the stochastic knapsack problem: The benefit of adaptivity. *FOCS*, pages 208–217, 2004.
- [8] B. C. Dean, M. X. Goemans, and J. Vondrák. Adaptivity and approximation for stochastic packing problems. *SODA*, 395 - 404, 2005.
- [9] S. Dobzinski and M. Schapira. An improved approximation algorithm for combinatorial auctions with submodular bidders *SODA*, 2006.
- [10] U. Feige. A threshold of \ln for approximating set cover. *J. ACM*, 45(4):634–652, 1998.
- [11] U. Feige. On maximizing welfare when utility functions are subadditive *STOC*, 2006.
- [12] M. Fisher, G. Nemhauser, and L. Wolsey. An Analysis of Approximation for Maximizing Submodular Set Functions-II. *Mathematical Programming Study*, 8 , 7387, 1978.
- [13] M. X. Goemans and J. Vondrák. Stochastic covering and adaptivity. *LATIN*, pages 532–543, 2006.
- [14] D. Kempe, J. M. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *KDD*, pages 137–146. ACM, 2003.
- [15] J. M. Kleinberg, C. H. Papadimitriou, and P. Raghavan. Segmentation problems. *J. ACM*, 51(2):263–280, 2004.
- [16] A. Krause and C. Guestrin. Near-optimal nonmyopic value of information in models. *AAAI*, pages 324–331, 2005.
- [17] A. Krause and C. Guestrin. Near-optimal observation selection using submodular functions. *AAAI*, pages 1650–1654. AAAI Press, 2007.
- [18] E. Mossel and S. Roch. On the submodularity of influence in social networks. *Approx*, pages 128–134, 2007.
- [19] G. Nemhauser and L. Wolsey. Best algorithms for approximating the maximum of a submodular set function. *Mathematics of Operations Research*. 3:3, 177188, 1978.
- [20] G. Nemhauser, L. Wolsey, and M. Fisher. An analysis of the approximations for maximizing submodular set functions. *Mathematical Programming*, 14:265–294, 1978.
- [21] J. Vondrák. Submodularity in combinatorial optimization. *PhD thesis, Charles University, Prague*, 2007.
- [22] J. Vondrák. *Personal communication*, 2009.
- [23] J. Vondrák. Optimal approximation for the submodular welfare problem in the value oracle model. *STOC*, 2008.
- [24] H. Weingartner. Mathematical programming and the analysis of capital budgeting problems. *Prentice-Hall*, 1963.