Maximizing non-monotone submodular functions[*]

Uriel Feige [†]
Dept. of Computer Science and Applied Mathematics
The Weizmann Institute
Rehovot, Israel
uriel.feige@weizmann.ac.il

Vahab S. Mirrokni [‡]
Google Research
New York, NY
mirrokni@google.com

Jan Vondrák [§]
IBM Alamaden Research Center
San Jose, CA
jvondrak@us.ibm.com

January 29, 2011

Abstract

Submodular maximization generalizes many important problems including Max Cut in directed and undirected graphs and hypergraphs, certain constraint satisfaction problems and maximum facility location problems. Unlike the problem of minimizing submodular functions, the problem of maximizing submodular functions is NP-hard.

In this paper, we design the first constant-factor approximation algorithms for maximizing nonnegative (non-monotone) submodular functions. In particular, we give a deterministic local-search $\frac{1}{3}$-approximation and a randomized $\frac{2}{5}$-approximation algorithm for maximizing nonnegative submodular functions. We also show that a uniformly random set gives a $\frac{1}{4}$-approximation. For symmetric submodular functions, we show that a random set gives a $\frac{1}{2}$-approximation, which can be also achieved by deterministic local search.

These algorithms work in the value oracle model where the submodular function is accessible through a black box returning $f(S)$ for a given set S. We show that in this model, a $(\frac{1}{2} + \epsilon)$-approximation for symmetric submodular functions would require an exponential number of queries for any fixed $\epsilon > 0$. In the model where f is given explicitly (as a sum of nonnegative submodular functions, each depending only on a constant number of elements), we prove NP-hardness of $(\frac{5}{6} + \epsilon)$-approximation in the symmetric case and NP-hardness of $(\frac{3}{4} + \epsilon)$-approximation in the general case.

---

1 Introduction

We consider the problem of *maximizing a nonnegative submodular function.* This means, given a submodular function $f : 2^X \to \mathbb{R}_+$, we want to find a set $S \subseteq X$ maximizing $f(S)$.

Definition 1.1. *A function $f : 2^X \to \mathbb{R}$ is submodular if for any $S, T \subseteq X$,*

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

An alternative definition of submodularity is the property of *decreasing marginal values*: For any $A \subseteq B \subseteq X$ and $x \in X \setminus B$, $f(B \cup \{x\}) - f(B) \leq f(A \cup \{x\}) - f(A)$. This can be deduced from the first definition by substituting $S = A \cup \{x\}$ and $T = B$; the reverse implication also holds.

We assume *value oracle* access to the submodular function; i.e., for a given set S, an algorithm can query an oracle to find its value $f(S)$.

Background. Submodularity, a discrete analog of convexity, has played an essential role in combinatorial optimization [35]. It appears in many important settings including cuts in graphs [21, 42, 18], rank function of matroids [10, 19], set covering problems [12], and plant location problems [8, 9]. In many settings such as set covering or matroid optimization, the relevant submodular functions are *monotone*, meaning that $f(S) \leq f(T)$ whenever $S \subseteq T$. Here, we are more interested in the general case where $f(S)$ is not necessarily monotone. A canonical example of such a submodular function is $f(S) = \sum_{e \in \delta(S)} w(e)$, where $\delta(S)$ is a cut in a graph (or hypergraph) induced by a set of vertices S and $w(e)$ is the weight of edge e. Cuts in undirected graphs and hypergraphs yield *symmetric* submodular functions, satisfying $f(S) = f(\bar{S})$ for all sets S. Symmetric submodular functions have been considered widely in the literature [17, 42]. It appears that symmetry allows better/simpler approximation results, and thus deserves separate attention.

The problem of maximizing a submodular function is of central importance, with special cases including Max Cut [21], Max Directed Cut [26], hypergraph cut problems, maximum facility location [1, 8, 9], and certain restricted satisfiability problems [27, 11]. While the Min Cut problem in graphs is a classical polynomial-time solvable problem, and more generally it has been shown that any submodular function can be *minimized* in polynomial time [45, 18], maximization turns out to be more difficult and indeed all the aforementioned special cases are NP-hard.

A related problem is Max-k-Cover, where the goal is to choose k sets whose union is as large as possible. It is known that a greedy algorithm provides a $(1 - 1/e)$-approximation for Max-k-Cover and this is optimal unless $P = NP$ [12]. More generally, this problem can be viewed as maximization of a monotone submodular function under a cardinality constraint, i.e. $\max\{f(S) : |S| \leq k\}$, assuming f submodular and $0 \leq f(S) \leq f(T)$ whenever $S \subseteq T$. Again, the greedy algorithm provides a $(1 - 1/e)$-approximation for this problem [39] and this is optimal in the oracle model [40]. More generally, a $(1 - 1/e)$-approximation can be achieved for monotone submodular maximization under a knapsack constraint [46]. For the problem of maximizing a monotone submodular function subject to a matroid constraint, the greedy algorithm gives only a $\frac{1}{2}$-approximation [16]. Recently, this has been improved to an optimal $(1 - 1/e)$-approximation using the *multilinear extension* of a submodular function [49, 6].

In contrast, here we consider the unconstrained maximization of a submodular function which is not necessarily monotone. We only assume that the function is nonnegative.[1] Typical examples of such a problem are Max Cut and Max Directed Cut. Here, the best approximation factors have been achieved using semidefinite programming: 0.878 for Max Cut [21] and 0.874 for Max Di-Cut [11, 31]. The approximation factor for Max Cut has been proved optimal, assuming the Unique Games Conjecture [29, 38]. Without the use of semidefinite programming, only $\frac{1}{2}$-approximation for Max Cut was known for a long time. For Max Di-Cut, a combinatorial $\frac{1}{2}$-approximation was presented in [26]. Recently, Trevisan gave a 0.53-approximation algorithm for Max Cut using a spectral partitioning method [48].

---

[1]For submodular functions without any restrictions, verifying whether the maximum of the function is greater than zero is NP-hard and requires exponentially many queries in the value oracle model. Thus, no efficient approximation algorithm can be found for general submodular maximization. For a general submodular function f with minimum value f_0, we can design an approximation algorithm to maximize a *normalized submodular function g* where $g(S) = f(S) - f_0$.

Model	Rnd. Set	Non-adapt.	Det. adaptive	Rnd. adaptive	Oracle hardness	NP-hardness
Symmetric	1/2	1/2	1/2	1/2	$1/2 + \epsilon$	$5/6 + \epsilon$
Asymmetric	1/4	1/3	1/3	2/5	$1/2 + \epsilon$	$3/4 + \epsilon$

Figure 1: Summary of our results: see *Our results* below for details.

More generally, submodular maximization encompasses such problems as Max Cut in hypergraphs and Max SAT with no mixed clauses (every clause contains only positive or only negative literals). Tight results are known for Max Cut in k-uniform hypergraphs for any fixed $k \geq 4$ [27, 24] where the optimal approximation factor $(1 - 2^{-k+1})$ is achieved by a random solution (and the same result holds for Max $(k-1)$-SAT with no mixed clauses [24, 25]). The lowest approximation factor $(\frac{7}{8})$ is achieved for $k = 4$; for $k < 4$, better than random solutions can be found by semidefinite programming.

Submodular maximization also appears in maximizing the difference of a monotone submodular function and a linear function. An illustrative example of this type is the maximum facility location problem in which we want to open a subset of facilities and maximize the total profit from served clients plus the cost of facilities we did *not* open. In a series of papers, approximation algorithms have been developed for a variant of this problem which is a special case of maximizing nonnegative submodular functions [8, 9]. Greedy and local-search methods have been the basic techniques of choice for these problems [8, 9, 39, 16]. The best approximation factor known for this problem is 0.828 [1], using linear programming.

In the general case of non-monotone submodular functions, the maximization problem has been studied in the operations research community. Many efforts have been focused on designing heuristics for this problem, including data-correcting search methods [22, 23, 28], accelatered greedy algorithms [43], and polyhedral algorithms [33]. Prior to our work, to the best of our knowledge, no guaranteed approximation factor was known for maximizing non-monotone submodular functions.

Our results. We design several constant-factor approximation algorithms for maximization of non-negative submodular functions. We also prove negative results, in particular a query complexity result matching our algorithmic result in the symmetric case.

Non-adaptive algorithms. A non-adaptive algorithm is allowed to generate a (possibly random) sequence of polynomially many sets, query their values and then produce a solution. In this model, we show that a $\frac{1}{4}$-approximation is achieved in expectation by a uniformly random set. For symmetric submodular functions, this gives a $\frac{1}{2}$-approximation. This coincides with the approximation factors obtained by random sets for Max Di-Cut and Max Cut. We prove that these factors cannot be improved, assuming that the algorithm returns one of the queried sets. However, we also design a non-adaptive algorithm which performs a polynomial-time computation on the obtained values and achieves a $\frac{1}{3}$-approximation. In the symmetric case, we prove that the $\frac{1}{2}$-approximation is optimal even among adaptive algorithms (see below).

Adaptive algorithms. An adaptive algorithm is allowed to perform a polynomial-time computation including a polynomial number of queries to a value oracle. In this (most natural) model, we develop a local-search $\frac{1}{2}$-approximation in the symmetric case and a $\frac{1}{3}$-approximation in the general case. Then we improve this to a $\frac{2}{5}$-approximation using a randomized "smooth local search". This is perhaps the most noteworthy of our algorithms; it proceeds by locally optimizing a smoothed variant of $f(S)$, obtained by biased sampling depending on S. This auxiliary function is in fact equal to the *multilinear extension* of $f(S)$ [5, 6], evaluated on a special subset of $[0, 1]^X$. The approach of locally optimizing a modified function has been referred to as "non-oblivious local search" in the literature; e.g., see [2] for a non-oblivious local-search $\frac{2}{5}$-approximation for the Max Di-Cut problem. Another (simpler) $\frac{2}{5}$-approximation algorithm for Max Di-Cut appears in [26]. However, these algorithms do not generalize naturally to ours and the re-appearance of the same approximation factor seems coincidental.

Hardness results. We show that it is impossible to improve the $\frac{1}{2}$-approximation algorithm for maximizing symmetric nonnegative submodular functions in the value oracle model. We prove that for any fixed $\epsilon > 0$, a $(\frac{1}{2} + \epsilon)$-approximation algorithm would require exponentially many queries. This settles

the status of symmetric submodular maximization in the value oracle model. Note that this query complexity lower bound does not assume any computational restrictions. In contrast, in the special case of Max Cut, polynomially many value queries suffice to infer all edge weights in the graph, and thereafter an exponential time computation (involving no further queries) would actually produce the optimal cut.

For explicitly represented submodular functions, known inapproximability results for Max Cut in graphs and hypergraphs provide an obvious limitation to the best possible approximation ratio. We prove stronger limitations. For any fixed $\epsilon > 0$, it is NP-hard to achieve an approximation factor of $\frac{3}{4} + \epsilon$ (or $\frac{5}{6} + \epsilon$) in the general (or symmetric) case, respectively. These results are valid even when the submodular function is given as a sum of polynomially many nonnegative submodular functions, each depending only on a constant number of elements, which is the case for the aforementioned special cases such as Max Cut.

Follow-up research. Following the conference version of this paper, several developments have been made in the area of submodular maximization, some of them inspired by the results and techniques in this paper. In particular, local-search constant-factor approximation algorithms have been developed for maximizing a non-monotone submodular function subject to multiple knapsack (linear) constraints or multiple matroid constraints [32], and these algorithms have been further improved in [34, 50]. For maximizing monotone submodular functions, optimal approximation algorithms have been developed in the cases of a matroid constraint [49, 6], and multiple knapsack constraints [30]. For the problem considered in this paper, maximization of a non-monotone submodular function, an improved 0.41-approximation has been found [41], and hardness of 0.695-approximation has been proved for explicit instances, assuming the unique games conjecture [4].

The idea behind the information-theoretic lower bounds in this paper has been generalized to prove hardness results for other problems, including social welfare maximization in combinatorial auctions with submodular bidders [36], and non-monotone submodular maximization subject to matroid base constraints [50]. A more general connection has been made between the approximability of submodular maximization problems and properties of their *multilinear relaxation* [50], which can be seen as a generalization of the hardness construction in this paper (Section 4.2).

It has been also shown that submodular functions can be approximated point-wise within an $\tilde{O}(\sqrt{n})$ factor using a polynomial number of value queries. This is optimal up to logarithmic factors [20]. Related to this phenomenon are several $\tilde{O}(\sqrt{n})$-approximation algorithms for submodular minimization under various constraints [47, 20].

Preliminaries. We use the following form of the Chernoff bound (see [3, Theorem A.1.16]).

Theorem 1.2. *Let $Y_1, \ldots, Y_t$ be independent random variables in $[-1, 1]$ such that $\mathbf{E}[Y_i] = 0$. Then*

$$\Pr[\sum_{i=1}^{t} Y_i > \lambda] \leq e^{-\lambda^2/2t}.$$

2 Non-adaptive algorithms

It is known that simply choosing a random cut is a good choice for Max Cut and Max Di-Cut, achieving an approximation factor of $1/2$ and $1/4$ respectively. We show the natural role of submodularity here by presenting the same approximation factors in the general case of submodular functions.

The Random Set Algorithm: RS.
Let $X(p)$ denote a random subset of X, where each element is sampled independently with probability p.

- Return $R = X(1/2)$.

Theorem 2.1. *Let $f : 2^X \rightarrow \mathbb{R}_+$ be a submodular function, $OPT = \max_{S \subseteq X} f(S)$ and let R denote a uniformly random subset $R = X(1/2)$. Then $\mathbf{E}[f(R)] \geq \frac{1}{4}OPT$. In addition, if f is symmetric ($f(S) = f(X \setminus S)$ for every $S \subseteq X$), then $\mathbf{E}[f(R)] \geq \frac{1}{2}OPT$.*

Before proving this result, we show a useful probabilistic property of submodular functions (extending the considerations of [13, 15]). This property will be essential in the analysis of our improved randomized algorithm as well.

Lemma 2.2. *Let $g : 2^X \to \mathbb{R}$ be submodular. Denote by $A(p)$ a random subset of A where each element appears with probability p. Then*

$$\mathbf{E}[g(A(p))] \geq (1-p) \ g(\emptyset) + p \ g(A).$$

Proof. By induction on the size of A: For $A = \emptyset$, the lemma is trivial. So assume $A = A' \cup \{x\}, x \notin A'$. Then $A(p) \cap A'$ is a subset of A' where each element appears with probability p; hence we denote it $A'(p)$. By submodularity, $g(A(p)) - g(A'(p)) \geq g(A(p) \cup A') - g(A')$, and therefore

$$
\begin{aligned}
\mathbf{E}[g(A(p))] &\geq \mathbf{E}[g(A'(p)) + g(A(p) \cup A') - g(A')] \\
&= \mathbf{E}[g(A'(p))] + \mathbf{E}[g(A(p) \cup A') - g(A')] \\
&= \mathbf{E}[g(A'(p))] + p \ (g(A) - g(A')).
\end{aligned}
$$

Applying the inductive hypothesis, $\mathbf{E}[g(A'(p))] \geq (1-p) \ g(\emptyset) + p \ g(A')$, we get the statement of the lemma. $\qquad \square$

By a double application of Lemma 2.2, we obtain the following.

Lemma 2.3. *Let $f : 2^X \to \mathbb{R}$ be submodular, $A, B \subseteq X$ two (not necessarily disjoint) sets and $A(p), B(q)$ their independently sampled subsets, where each element of A appears in $A(p)$ with probability p and each element of B appears in $B(q)$ with probability q. Then*

$$\mathbf{E}[f(A(p) \cup B(q))] \geq (1-p)(1-q) \ f(\emptyset) + p(1-q) \ f(A) + (1-p)q \ f(B) + pq \ f(A \cup B).$$

Proof. Condition on $A(p) = R$ and define $g(T) = f(R \cup T)$. This is a submodular function as well and Lemma 2.2 implies $\mathbf{E}[g(B(q))] \geq (1-q) \ f(R) + q \ f(R \cup B)$. Also, $\mathbf{E}[g(B(q))] = \mathbf{E}[f(A(p) \cup B(q)) \mid A(p) = R]$, and by unconditioning: $\mathbf{E}[f(A(p) \cup B(q))] \geq \mathbf{E}[(1-q) \ f(A(p)) + q \ f(A(p) \cup B)]$. Finally, we apply Lemma 2.2 once again: $\mathbf{E}[f(A(p))] \geq (1-p) \ f(\emptyset) + p \ f(A)$, and by applying the same to the submodular function $h(S) = f(S \cup B)$, $\mathbf{E}[f(A(p) \cup B)] \geq (1-p) \ f(B) + p \ f(A \cup B)$. This implies the claim. $\qquad \square$

This lemma gives immediately the performance of Algorithm RS.

Proof. Denote the optimal set by S and its complement by $\bar{S}$. We can write $R = S(1/2) \cup \bar{S}(1/2)$. Using Lemma 2.3, we get

$$\mathbf{E}[f(R)] \geq \frac{1}{4}f(\emptyset) + \frac{1}{4}f(S) + \frac{1}{4}f(\bar{S}) + \frac{1}{4}f(X).$$

Every term is nonnegative and $f(S) = OPT$, so we get $\mathbf{E}[f(R)] \geq \frac{1}{4}OPT$. In addition, if f is symmetric, we also have $f(\bar{S}) = OPT$ and then $\mathbf{E}[f(R)] \geq \frac{1}{2}OPT$. $\qquad \square$

As we show in Section 4.2, $\frac{1}{4}$-approximation is optimal for non-adaptive algorithms that are required to return one of the queried sets. Here, we show that it is possible to design a $\frac{1}{3}$-approximation algorithm which queries a polynomial number of sets non-adaptively and then returns a possibly different set after a polynomial-time computation. The intuition behind the algorithm comes from the Max Di-Cut problem: When does a random cut achieve only 1/4 of the optimum? This is if and only if the optimum contains all the directed edges of the graph, i.e. the vertices can be partitioned into $V = A \cup B$ so that all edges of the graph are directed from A to B. However, in this case it is easy to find the optimal solution, by a local test on the in-degree and out-degree of each vertex. In the language of submodular function maximization, this means that elements can be easily partitioned into those whose inclusion in S always increases the value of $f(S)$, and those which always decrease $f(S)$. Our generalization of this local test is the following.

4

Definition 2.4. *Let* $R = X(1/2)$ *denote a uniformly random subset of* X. *For each element* x, *define*

$$\omega(x) = \mathbf{E}[f(R \cup \{x\}) - f(R \setminus \{x\})].$$

Note that the expectation $\mathbf{E}[f(R)]$, for any distribution of R, can be estimated by random sampling up to an additive error polynomially small relative to OPT. More precisely, we use the following lemma.

Lemma 2.5. *If* $f : 2^X \to \mathbb{R}$, $R_1, \ldots, R_t$ *are independent samples from some distribution on* 2^X, *and* $\mathbf{E}[f(R_i)] = \mu$ *for all* i, *then*

$$\left| \frac{1}{t} \sum_{i=1}^{t} f(R_i) - \mu \right| \leq \epsilon \max f(S)$$

with probability at least $1 - 2e^{-t\epsilon^2/2}$, *for any* $\epsilon \in [0, 1]$.

Proof. Let $M = \max f(S)$; $f(R)$ is a random variable in the range $[0, M]$. Let $Y_i = \frac{1}{M}(f(R_i) - \mu)$ where R_i is the i-th random sample. We have $Y_i \in [-1, 1]$, these random variables are independent, and $\mathbf{E}[Y_i] = 0$. By Theorem 1.2,

$$\Pr\left[\left| \frac{1}{t} \sum_{i=1}^{t} f(R_i) - \mu \right| > \epsilon M \right] = \Pr\left[\left| \sum_{i=1}^{t} Y_i \right| > t\epsilon \right] \leq 2e^{-t\epsilon^2/2}.$$

$\square$

This is sufficient for our purposes; we use $\epsilon = 1/poly(n)$, $t = n/\epsilon^2$, and obtain with high probability estimates with additive error $OPT/poly(n)$. In the following, we assume that we have estimates for quantities like $\mathbf{E}[f(R)]$ and $\omega(x)$ within additive error $OPT/poly(n)$.

The Non-Adaptive Algorithm: NA.

- Use random sampling to find $\tilde{\omega}(x)$ for each $x \in X$, such that $|\tilde{\omega}(x) - \omega(x)| < \frac{1}{n^2}OPT$ w.h.p.

- Independently, sample a random set $R = X(1/2)$.

- With prob. $8/9$, return R.

- With prob. $1/9$, return $A = \{x \in X : \tilde{\omega}(x) > 0\}$.

Theorem 2.6. *For any nonnegative submodular function, Algorithm NA achieves expected value at least* $(1/3 - o(1)) OPT$.

Proof. Let $A = \{x \in X : \tilde{\omega}(x) > 0\}$ and $B = X \setminus A = \{x \in X : \tilde{\omega}(x) \leq 0\}$. Therefore we have $\omega(x) \geq -OPT/n^2$ for any $x \in A$ and $\omega(x) \leq OPT/n^2$ for any $x \in B$. We shall keep in mind that (A, B) is a partition of all the elements, and so we have $(A \cap T) \cup (B \cap T) = T$ for any set T, etc.

Denote by C the optimal set, $f(C) = OPT$. Let $f(A) = \alpha$, $f(B \cap C) = \beta$ and $f(B \cup C) = \gamma$. By submodularity, we have

$$\alpha + \beta = f(A) + f(B \cap C) \geq f(\emptyset) + f(A \cup (B \cap C)) \geq f(A \cup C)$$

and

$$\alpha + \beta + \gamma \geq f(A \cup C) + f(B \cup C) \geq f(X) + f(C) \geq OPT.$$

Therefore, either α, the value of A, is at least $OPT/3$, or else one of β and γ is at least $OPT/3$; we prove that then $\mathbf{E}[f(R)] \geq OPT/3$ as well.

Let us start with $\beta = f(B \cap C)$. Instead of $\mathbf{E}[f(R)]$, we show that it is enough to estimate $\mathbf{E}[f(R \cup (B \cap C))]$. Recall that for any $x \in B$, we have $\omega(x) = \mathbf{E}[f(R \cup \{x\}) - f(R \setminus \{x\})] \leq OPT/n^2$. Consequently, we

5

also have $\mathbf{E}[f(R\cup\{x\})-f(R)]=\frac{1}{2}\omega(x)\leq OPT/(2n^2)$. Let us order the elements of $B\cap C=\{b_1,\ldots,b_\ell\}$ and write

$$f(R\cup(B\cap C))=f(R)+\sum_{j=1}^{\ell}(f(R\cup\{b_1,\ldots,b_j\})-f(R\cup\{b_1,\ldots,b_{j-1}\})).$$

By the property of decreasing marginal values, we get

$$f(R\cup(B\cap C))\leq f(R)+\sum_{j=1}^{\ell}(f(R\cup\{b_j\})-f(R))=f(R)+\sum_{x\in B\cap C}(f(R\cup\{x\})-f(R))$$

and therefore

$$
\begin{aligned}
\mathbf{E}[f(R\cup(B\cap C))] &\leq \mathbf{E}[f(R)]+\sum_{x\in B\cap C}\mathbf{E}[f(R\cup\{x\})-f(R)]\\
&\leq \mathbf{E}[f(R)]+|B\cap C|\frac{OPT}{2n^2}\leq \mathbf{E}[f(R)]+\frac{OPT}{2n}.
\end{aligned}
$$

So it is enough to lower-bound $\mathbf{E}[f(R\cup(B\cap C))]$. We do this by defining a new submodular function, $g(R)=f(R\cup(B\cap C))$, and applying Lemma 2.3 to $\mathbf{E}[g(R)]=\mathbf{E}[g(C(1/2)\cup\bar{C}(1/2))]$. The lemma implies that

$$
\begin{aligned}
\mathbf{E}[f(R\cup(B\cap C))] &\geq \frac{1}{4}g(\emptyset)+\frac{1}{4}g(C)+\frac{1}{4}g(\bar{C})+\frac{1}{4}g(X)\\
&\geq \frac{1}{4}g(\emptyset)+\frac{1}{4}g(C)=\frac{1}{4}f(B\cap C)+\frac{1}{4}f(C)=\frac{\beta}{4}+\frac{OPT}{4}.
\end{aligned}
$$

Note that $\beta\geq OPT/3$ implies $\mathbf{E}[f(R\cup(B\cap C))]\geq OPT/3$. Symmetrically, we show a similar analysis for $\mathbf{E}[f(R\cap(B\cup C))]$. Now we use the fact that for any $x\in A$, $\mathbf{E}[f(R)-f(R\setminus\{x\})]=\frac{1}{2}\omega(x)\geq -OPT/(2n^2)$. Let $A\setminus C=\{a_1,a_2,\ldots,a_k\}$ and write

$$
\begin{aligned}
f(R) &= f(R\setminus(A\setminus C))+\sum_{j=1}^{k}(f(R\setminus\{a_{j+1},\ldots,a_k\})-f(R\setminus\{a_j,\ldots,a_k\}))\\
&\geq f(R\setminus(A\setminus C))+\sum_{j=1}^{k}(f(R)-f(R\setminus\{a_j\}))
\end{aligned}
$$

using the condition of decreasing marginal values. Note that $R\setminus(A\setminus C)=R\cap(B\cup C)$. By taking the expectation,

$$
\begin{aligned}
\mathbf{E}[f(R)] &\geq \mathbf{E}[f(R\cap(B\cup C))]+\sum_{j=1}^{k}\mathbf{E}[f(R)-f(R\setminus\{a_j\})]\\
&= \mathbf{E}[f(R\cap(B\cup C))]-|A\setminus C|\frac{OPT}{2n^2}\geq \mathbf{E}[f(R\cap(B\cup C))]-\frac{OPT}{2n}.
\end{aligned}
$$

Again, we estimate $\mathbf{E}[f(R\cap(B\cup C))]=\mathbf{E}[f(C(1/2)\cup(B\setminus C)(1/2))]$ using Lemma 2.3. We get

$$\mathbf{E}[f(R\cap(B\cup C))]\geq\frac{1}{4}f(\emptyset)+\frac{1}{4}f(C)+\frac{1}{4}f(B\setminus C)+\frac{1}{4}f(B\cup C)\geq\frac{OPT}{4}+\frac{\gamma}{4}.$$

Now we combine our estimates for $\mathbf{E}[f(R)]$:

$$\mathbf{E}[f(R)]+\frac{OPT}{2n}\geq\frac{1}{2}\mathbf{E}[f(R\cup(B\cap C))]+\frac{1}{2}\mathbf{E}[f(R\cap(B\cup C))]\geq\frac{OPT}{4}+\frac{\beta}{8}+\frac{\gamma}{8}.$$

Finally, the expected value obtained by the algorithm is

$$\frac{8}{9}\mathbf{E}[f(R)]+\frac{1}{9}f(A)\geq\frac{2}{9}OPT-\frac{4}{9n}OPT+\frac{\beta}{9}+\frac{\gamma}{9}+\frac{\alpha}{9}\geq\left(\frac{1}{3}-\frac{4}{9n}\right)OPT$$

since $\alpha+\beta+\gamma\geq OPT$. $\qquad\square$

3 Adaptive algorithms

We turn to adaptive algorithms for the problem of maximizing a general nonnegative submodular function. We propose several algorithms improving the $\frac{1}{4}$-approximation achieved by a random set.

3.1 Deterministic local search

Here, we present a deterministic $\frac{1}{3}$-approximation algorithm. Our deterministic algorithm is based on a simple local-search technique. We try to increase the value of our solution S by either including a new element in S or discarding one of the elements of S. We call S a *local optimum* if no such operation increases the value of S. Local optima have the following property which was first observed in [7, 23].

Lemma 3.1. *Given a submodular function f, if S is a local optimum of f, and $T \subseteq S$ or $T \supseteq S$, then $f(T) \le f(S)$.*

This property turns out to be very useful in comparing a local optimum to the global optimum. However, it is known that finding a local optimum for the Max Cut problem is PLS-complete [44]. Therefore, we relax our local search and find an *approximate local optimal solution.*

Local Search Algorithm: LS.

1. Let $S := \{v\}$ where $f(\{v\})$ is the maximum over all singletons $v \in X$.

2. If there exists an element $a \in X \backslash S$ such that $f(S \cup \{a\}) > (1 + \frac{\epsilon}{n^2})f(S)$, then let $S := S \cup \{a\}$, and go back to Step 2.

3. If there exists an element $a \in S$ such that $f(S \backslash \{a\}) > (1 + \frac{\epsilon}{n^2})f(S)$, then let $S := S \backslash \{a\}$, and go back to Step 2.

4. Return the maximum of $f(S)$ and $f(X \backslash S)$.

It is easy to see that if the algorithm terminates, the set S is a $(1 + \frac{\epsilon}{n^2})$-approximate local optimum, in the following sense.

Definition 3.2. *Given $f : 2^X \to \mathbb{R}$, a set S is called a $(1+\alpha)$-approximate local optimum, if $(1+\alpha)f(S) \ge f(S \setminus \{v\})$ for any $v \in S$, and $(1 + \alpha)f(S) \ge f(S \cup \{v\})$ for any $v \notin S$.*

We prove the following analogue of Lemma 3.1.

Lemma 3.3. *If S is an $(1+\alpha)$-approximate local optimum for a submodular function f then for any set T such that $T \subseteq S$ or $T \supseteq S$, we have $f(T) \le (1 + n\alpha)f(S)$.*

Proof. Let $T = T_1 \subseteq T_2 \subseteq \ldots \subseteq T_k = S$ be a chain of sets where $T_i \backslash T_{i-1} = \{a_i\}$. For each $2 \le i \le k$, we know that $f(T_i) - f(T_{i-1}) \ge f(S) - f(S \setminus \{a_i\}) \ge -\alpha f(S)$ using the submodularity and approximate local optimality of S. Summing up these inequalities, we get $f(S) - f(T) \ge -k\alpha f(S)$. Thus $f(T) \le (1 + k\alpha)f(S) \le (1 + n\alpha)f(S)$. This completes the proof for set $T \subseteq S$. The proof for $T \supseteq S$ is very similar. $\square$

Theorem 3.4. *Algorithm LS is a $\left(\frac{1}{3} - \frac{\epsilon}{n}\right)$-approximation algorithm for maximizing nonnegative submodular functions, and a $\left(\frac{1}{2} - \frac{\epsilon}{n}\right)$-approximation algorithm for maximizing nonnegative symmetric submodular functions. The algorithm uses at most $O(\frac{1}{\epsilon}n^3 \log n)$ oracle calls.*

Proof. Consider an optimal solution C and let $\alpha = \frac{\epsilon}{n^2}$. If the algorithm terminates, the set S obtained at the end is a $(1 + \alpha)$-approximate local optimum. By Lemma 3.3, $(1 + n\alpha)f(S) \ge f(S \cap C)$ and $(1 + n\alpha)f(S) \ge f(S \cup C)$. Using submodularity, $f(S \cup C) + f(X \backslash S) \ge f(C \backslash S) + f(X) \ge f(C \backslash S)$, and $f(S \cap C) + f(C \backslash S) \ge f(C) + f(\emptyset) \ge f(C)$. Putting these inequalities together, we get

$$
\begin{aligned}
2(1 + n\alpha)f(S) + f(X \backslash S) &\ge f(S \cap C) + f(S \cup C) + f(X \backslash S) \\
&\ge f(S \cap C) + f(C \backslash S) \ge f(C).
\end{aligned}
$$

For $\alpha = \frac{\epsilon}{n^2}$, this implies that either $f(S) \geq (\frac{1}{3} - \frac{\epsilon}{n})OPT$ or $f(X\backslash S) \geq (\frac{1}{3} - \frac{\epsilon}{n})OPT$.

For symmetric submodular functions, we use Lemma 3.3 to obtain the inequality $(1 + n\alpha)f(S) \geq f(S \cup \bar{C})$, and therefore

$$2(1 + n\alpha)f(S) \geq f(S \cap C) + f(S \cup \bar{C}) = f(S \cap C) + f(\bar{S} \cap C) \geq f(C)$$

and hence $f(S)$ is a $(\frac{1}{2} - \frac{\epsilon}{n})$-approximation.

To bound the running time of the algorithm, let v be a singleton of maximum value $f(\{v\})$. It is simple to see that $OPT \leq nf(\{v\})$. After each iteration, the value of the function increases by a factor of at least $(1 + \frac{\epsilon}{n^2})$. Therefore, if we iterate k times, then $f(S) \geq (1 + \frac{\epsilon}{n^2})^k f(\{v\})$ and yet $f(S) \leq nf(\{v\})$, hence $k = O(\frac{1}{\epsilon}n^2 \log n)$. The number of queries is $O(\frac{1}{\epsilon}n^3 \log n)$. $\square$

Tight example. Consider a submodular function defined as a cut function in a directed graph $D = (V, A)$. The graph has 4 vertices, $V = \{a, b, c, d\}$ and 4 arcs, $A = \{(a, b), (b, c), (c, b), (c, d)\}$. The cut function $f(S)$ denotes the number of arcs *leaving* S. The optimum solution is $S^* = \{a, c\}$ which gives a cut of size $f(S^*) = 3$. However, the algorithm LS could terminate with the set $S = \{a, b\}$ which is a local optimum of value $f(S) = 1$ (switching any element keeps the objective value the same).

3.2 Randomized local search

Next, we present a randomized algorithm which improves the approximation ratio of $1/3$ to $2/5$. The main idea behind this algorithm is to find a "smoothed" local optimum, where elements are sampled randomly but with different probabilities, based on some underlying set A. The general approach of local search, based on a function derived from the one we are interested in, has been referred to as "non-oblivious local search" in the literature [2]. The auxiliary function we are trying to optimize is the *multilinear extension* of $f(S)$ (see [5, 6]):

$$F(x) = \sum_{S \subseteq X} f(S) \prod_{i \in S} x_i \prod_{j \notin S} (1 - x_j).$$

Definition 3.5. *We say that a set is sampled with bias δ based on A, if elements in A are sampled independently with probability $p = \frac{1+\delta}{2}$ and elements outside of A are sampled independently with probability $q = \frac{1-\delta}{2}$. We denote this random set by $\mathcal{R}(A, \delta)$.*

Observe that the expected value $\mathbf{E}[f(\mathcal{R}(A, \delta))]$ is equal to $F(x)$ evaluated at a point $x \in [0, 1]^X$ whose coordinates have only two distinct values, p and q. The set A describes exactly those coordinates whose value is p. Our algorithm effectively performs a local search over such points.

The Smooth Local Search algorithm: SLS.

1. Fix parameters $\delta, \delta' \in [-1, 1]$. Start with $A = \emptyset$. Let $n = |X|$ denote the total number of elements. In the following, use an estimate for OPT, for example from Algorithm LS.

2. For each element x, define

$$\omega_{A,\delta}(x) = \mathbf{E}[f(\mathcal{R}(A, \delta) \cup \{x\})] - \mathbf{E}[f(\mathcal{R}(A, \delta) \setminus \{x\})].$$

   By repeated sampling, we compute $\tilde{\omega}_{A,\delta}(x)$, an estimate of $\omega_{A,\delta}(x)$ within $\pm\frac{1}{n^2}OPT$ w.h.p.

3. If there is $x \in X \setminus A$ such that $\tilde{\omega}_{A,\delta}(x) > \frac{2}{n^2}OPT$, include x in A and go to Step 2.

4. If there is $x \in A$ s.t. $\tilde{\omega}_{A,\delta}(x) < -\frac{2}{n^2}OPT$, remove x from A and go to Step 2.

5. Return a random set $\mathcal{R}(A, \delta')$, for a suitably chosen δ'.

In effect, we find an approximate local optimum of a derived function $\Phi(A) = \mathbf{E}[f(\mathcal{R}(A,\delta))]$. Then we return a set sampled according to $\mathcal{R}(A,\delta')$; possibly for $\delta' \neq \delta$. One can run Algorithm SLS with $\delta = \delta'$ and prove that the best approximation for such parameters is achieved by setting $\delta = \delta' = \frac{\sqrt{5}-1}{2}$, the golden ratio. Then, we get an approximation factor of $\frac{3-\sqrt{5}}{2} - o(1) \simeq 0.38$. (We omit the proof here.) However, our best result is achieved by taking a combination of two choices as follows.

Theorem 3.6. *Algorithm SLS runs in polynomial time. If we run SLS with $\delta = 1/3$ and δ' is chosen randomly to be $1/3$ with probability 0.9, or -1 with probability 0.1, the expected value of the solution is at least $(\frac{2}{5} - o(1))OPT$.*

Proof. Let $\Phi(A) = \mathbf{E}[f(\mathcal{R}(A,\delta))]$. We set $B = X \setminus A$. Recall that in $\mathcal{R}(A,\delta)$, elements from A are sampled with probability $p = \frac{1+\delta}{2}$, while elements from B are sampled with probability $q = \frac{1-\delta}{2}$. Consider Step 3 where an element x is added to A. Let $A' = A \cup \{x\}$ and $B' = B \setminus \{x\}$. The reason why x is added to A is that $\tilde{\omega}_{A,\delta}(x) > \frac{2}{n^2}OPT$; i.e. $\omega_{A,\delta}(x) > \frac{1}{n^2}OPT$. During this step, $\Phi(A)$ increases by

$$
\begin{aligned}
\Phi(A') - \Phi(A) &= \mathbf{E}[f(A'(p) \cup B'(q)) - f(A(p) \cup B(q))] \\
&= (p-q)\,\mathbf{E}[f(A(p) \cup B'(q) \cup \{x\}) - f(A(p) \cup B'(q))] \\
&= \delta\,\mathbf{E}[f(\mathcal{R}(A,\delta) \cup \{x\}) - f(\mathcal{R}(A,\delta) \setminus \{x\})] \\
&= \delta\,\omega_{A,\delta}(x) > \frac{\delta}{n^2}OPT.
\end{aligned}
$$

Similarly, executing Step 4 increases $\Phi(A)$ by at least $\frac{\delta}{n^2}OPT$. Since the value of $\Phi(A)$ is always between 0 and OPT, the algorithm cannot iterate more than n^2/δ times and thus it runs in polynomial time.

From now on, let A be the set at the end of the algorithm and $B = X \setminus A$. We also use $R = A(p) \cup B(q)$ to denote a random set from the distribution $\mathcal{R}(A,\delta)$. We denote by C the optimal solution, while our algorithm returns either R (for $\delta' = \delta$) or B (for $\delta' = -1$). When the algorithm terminates, we have $\omega_{A,\delta}(x) \geq -\frac{3}{n^2}OPT$ for any $x \in A$, and $\omega_{A,\delta}(x) \leq \frac{3}{n^2}OPT$ for any $x \in B$. Consequently, for any $x \in B$ we have $\mathbf{E}[f(R \cup \{x\}) - f(R)] = \Pr[x \notin R]\mathbf{E}[f(R \cup \{x\}) - f(R \setminus \{x\})] = \frac{2}{3}\omega_{A,\delta}(x) \leq \frac{2}{n^2}OPT$, using $\Pr[x \notin R] = p = 2/3$. By submodularity, we get

$$
\begin{aligned}
\mathbf{E}[f(R \cup (B \cap C))] &\leq \mathbf{E}[f(R)] + \sum_{x \in B \cap C} \mathbf{E}[f(R \cup \{x\}) - f(R))] \\
&\leq \mathbf{E}[f(R)] + |B \cap C|\frac{2}{n^2}OPT \\
&\leq \mathbf{E}[f(R)] + \frac{2}{n}OPT.
\end{aligned}
$$

Similarly, we can obtain $\mathbf{E}[f(R \cap (B \cup C))] \leq \mathbf{E}[f(R)] + \frac{2}{n}OPT$. This means that instead of R, we can analyze $R \cup (B \cap C)$ and $R \cap (B \cup C)$. In order to estimate $\mathbf{E}[f(R \cup (B \cap C))]$ and $\mathbf{E}[f(R \cap (B \cup C))]$, we use a further extension of Lemma 2.3 which can be proved by another iteration of the same proof:

$$
(*) \quad \mathbf{E}[f(A_1(p_1) \cup A_2(p_2) \cup A_3(p_3))] \geq \sum_{I \subseteq \{1,2,3\}} \prod_{i \in I} p_i \prod_{i \notin I} (1-p_i)\, f\left(\bigcup_{i \in I} A_i\right).
$$

First, we deal with $R \cap (B \cup C) = (A \cap C)(p) \cup (B \cap C)(q) \cup (B \setminus C)(q)$. We plug in $\delta = 1/3$, i.e. $p = 2/3$ and $q = 1/3$. Then $(*)$ yields

$$
\mathbf{E}[f(R \cap (B \cup C))] \geq \frac{8}{27}f(A \cap C) + \frac{2}{27}f(B \cup C) + \frac{2}{27}f(B \cap C) + \frac{4}{27}f(C) + \frac{4}{27}f(F) + \frac{1}{27}f(B)
$$

where we denote $F = (A \cap C) \cup (B \setminus C)$ and we discarded the terms $f(\emptyset) \geq 0$ and $f(B \setminus C) \geq 0$. Similarly, we estimate $\mathbf{E}[f(R \cup (B \cap C))]$, applying $(*)$ to a submodular function $h(R) = f(R \cup (B \cap C))$ and writing $\mathbf{E}[f(R \cup (B \cap C))] = \mathbf{E}[h(R)] = \mathbf{E}[h((A \cap C)(p) \cup (A \setminus C)(p) \cup B(q))]$:

$$
\mathbf{E}[f(R \cup (B \cap C))] \geq \frac{8}{27}f(A \cup C) + \frac{2}{27}f(B \cup C) + \frac{2}{27}f(B \cap C) + \frac{4}{27}f(C) + \frac{4}{27}f(\bar{F}) + \frac{1}{27}f(B).
$$

Here, $\bar{F} = (A \setminus C) \cup (B \cap C)$. We use $\mathbf{E}[f(R)] + \frac{2}{n} OPT \geq \frac{1}{2}(\mathbf{E}[f(R \cap (B \cup C))] + \mathbf{E}[f(R \cup (B \cap C))])$ and combine the two estimates.

$$
\begin{aligned}
\mathbf{E}[f(R)] + \frac{2}{n} OPT \geq \quad & \frac{4}{27} f(A \cap C) + \frac{4}{27} f(A \cup C) + \frac{2}{27} f(B \cap C) + \frac{2}{27} f(B \cup C) \\
+ \quad & \frac{4}{27} f(C) + \frac{2}{27} f(F) + \frac{2}{27} f(\bar{F}) + \frac{1}{27} f(B).
\end{aligned}
$$

Now we add $\frac{3}{27} f(B)$ on both sides and apply submodularity: $f(B) + f(F) \geq f(B \cup C) + f(B \setminus C) \geq f(B \cup C)$ and $f(B) + f(\bar{F}) \geq f(B \cup (A \setminus C)) + f(B \cap C) \geq f(B \cap C)$. This leads to

$$
\mathbf{E}[f(R)] + \frac{1}{9} f(B) + \frac{2}{n} OPT \geq \frac{4}{27} f(A \cap C) + \frac{4}{27} f(A \cup C) + \frac{4}{27} f(B \cap C) + \frac{4}{27} f(B \cup C) + \frac{4}{27} f(C)
$$

and since $f(A \cap C) + f(B \cap C) \geq f(C)$ and $f(A \cup C) + f(B \cup C) \geq f(C)$, we get

$$
\mathbf{E}[f(R)] + \frac{1}{9} f(B) + \frac{2}{n} OPT \geq \frac{12}{27} f(C) = \frac{4}{9} OPT.
$$

This implies that $\frac{9}{10} \mathbf{E}[f(R)] + \frac{1}{10} f(B) \geq (\frac{2}{5} - \frac{9}{5n}) OPT$. $\qquad\square$

The analysis of this algorithm is tight, as can be seen from the following example.

Tight example. Consider $V = \{a, b, c, d\}$ and a directed graph $D = (V, A)$, consisting of 3 edges, $A = \{(a, b), (b, c), (c, d)\}$. Let us define a submodular function $f : 2^V \to \mathbb{R}_+$ corresponding to directed cuts in D, by defining $f(S)$ as the number of edges going from S to $V \setminus S$.

Observe that $A = \{a, b\}$ could be the set found by our algorithm. In fact, this is a local optimum of $\mathbf{E}[f(\mathcal{R}(A, \delta))]$ for any value of δ, since we have $\mathbf{E}[f(\mathcal{R}(\{a, b\}, \delta))] = 2pq + p^2 = 1 - q^2$ (where $p = (1 + \delta)/2$ and $q = (1 - \delta)/2$). It can be verified that the value of $\mathbf{E}[f(\mathcal{R}(A, \delta))]$ for any A obtained by switching one element is at most $1 - q^2$.

Algorithm SLS returns either $\mathcal{R}(A, 1/3)$ of expected value $1 - (1/3)^2 = 8/9$, with probability 0.9, or B of value 0, with probability 0.1. Thus the expected value returned by SLS is 0.8, while the optimum is $f(\{a, c\}) = 2$.

This example can be circumvented if we define SLS to take the maximum of $\mathbf{E}[f(\mathcal{R}(A, 1/3))]$ and $f(B)$ rather than a weighted average. However, it is not clear how to use this in the analysis.

4 Inapproximability Results

In this section, we give hardness results for submodular maximization. Our results are of two flavors. First, we consider submodular functions in the form of a sum of "building blocks" of constant size, more precisely nonnegative submodular functions depending only on a constant number of elements. We refer to this as *succint representation*. Note that all the special cases such as Max Cut are of this type. For algorithms in this model, we prove complexity-theoretic inapproximability results, the strongest one stating that in the general case, a $(\frac{3}{4} + \epsilon)$-approximation for any fixed $\epsilon > 0$ would imply $P = NP$.

In the value oracle model, we show a much tighter result. Namely, any algorithm achieving a $(\frac{1}{2} + \epsilon)$-approximation for a fixed $\epsilon > 0$ would require an exponential number of queries to the value oracle. This holds even in the case of symmetric submodular functions, i.e. our $\frac{1}{2}$-approximation algorithm is optimal in this case.

4.1 NP-hardness results

Our reductions are based on Håstad's 3-bit and 4-bit PCP verifiers [27]. Some inapproximability results can be obtained immediately from [27], by considering the known special cases of submodular maximization, e.g. Max Cut in 4-uniform hypergraphs which is NP-hard to approximate within a factor better than $\frac{7}{8}$.

We obtain stronger hardness results by reductions from systems of parity equations. The parity function is not submodular, but we can obtain hardness results by a careful construction of a "submodular gadget" for each equation.

Theorem 4.1. *There is no polynomial-time $(\frac{5}{6} + \epsilon)$-approximation algorithm to maximize a nonnegative symmetric submodular function in succint representation, unless $P = NP$.*

Proof. Consider an instance of Max E4-Lin-2, a system $\mathcal{E}$ of m parity equations on 4 boolean variables each. Let us define two elements for each variable, T_i and F_i, corresponding to variable x_i being either true or false. For each equation e on variables (x_i, x_j, x_k, x_ℓ), we define a function $g_e(S)$. (This is our "submodular gadget".) Let $S' = S \cap \{T_i, F_i, T_j, F_j, T_k, F_k, T_\ell, F_\ell\}$. We say that S' is a *valid quadruple*, if it defines a boolean assignment to x_i, x_j, x_k, x_ℓ, i.e. contains exactly one element from each pair $\{T_i, F_i\}$. The function value is determined by S':

- If $|S'| < 4$, let $g_e(S) = |S'|$. If $|S'| > 4$, let $g_e(S) = 8 - |S'|$.

- If S' is a valid quadruple satisfying e, let $g_e(S) = 4$ (a *true quadruple*).

- If S' is a valid quadruple not satisfying e, let $g_e(S) = 8/3$ (a *false quadruple*).

- If $|S'| = 4$ but S' is not a valid quadruple, let $g_e(S) = 10/3$ (an *invalid quadruple*).

It can be verified that this is a submodular function, using the structure of the parity constraint. We define $f(S) = \sum_{e \in \mathcal{E}} g_e(S)$. This is again a nonnegative submodular function. Observe that for each equation, it is more profitable to choose an invalid assignment than a valid assignment which does not satisfy the equation. Nevertheless, we claim that WLOG the maximum is obtained by selecting exactly one of T_i, F_i for each variable: Consider a set S and call a variable *undecided*, if S contains both or neither of T_i, F_i. For each equation with an undecided variable, we get value at most $10/3$. Now, modify S into $\tilde{S}$ by randomly selecting exactly one of T_i, F_i for each undecided variable. The new set $\tilde{S}$ induces a valid assignment to all variables. For equations which had a valid assignment already in S, the value does not change. Each equation which had an undecided varible is satisfied by $\tilde{S}$ with probability $1/2$. Therefore, the expected value for each such equation is $\frac{1}{2}(\frac{8}{3} + 4) = \frac{10}{3}$, at least as before, and $\mathbf{E}[f(\tilde{S})] \geq f(S)$. Hence there must exist a set $\tilde{S}$ such that $f(\tilde{S}) \geq f(S)$ and $\tilde{S}$ induces a valid assignment.

Consequently, we have $OPT = \max f(S) = \frac{8}{3}m + \frac{4}{3}\#SAT$ where $\#SAT$ is the maximum number of satisfiable equations. Since it is NP-hard to distinguish whether $\#SAT \geq (1-\epsilon)m$ or $\#SAT \leq (\frac{1}{2} + \epsilon)m$, it is also NP-hard to distinguish between $OPT \geq (4 - \epsilon)m$ and $OPT \leq (\frac{10}{3} + \epsilon)m$. □

In the case of general nonnegative submodular functions, we improve the hardness threshold to $3/4$. This hardness result is slightly more involved. It requires certain properties of Håstad's 3-bit PCP verifier, implying that Max E3-Lin-2 (a system of parity equations with 3 boolean variables each) is NP-hard to approximate even for linear systems of a special structure.

Lemma 4.2. *Fix any $\epsilon > 0$ and consider systems of weighted linear equations (of total weight 1) over boolean variables, partitioned into $\mathcal{X}$ and $\mathcal{Y}$, so that each equation contains 1 variable $x_i \in \mathcal{X}$ and 2 variables $y_j, y_k \in \mathcal{Y}$. Define a matrix $P \in [0,1]^{\mathcal{Y} \times \mathcal{Y}}$ where P_{jk} is the weight of all equations where the first variable from $\mathcal{Y}$ is y_j and the second variable is y_k. Then it is NP-hard to decide whether there is a solution satisfying equations of weight at least $1 - \epsilon$ or whether any solution satisfies equations of weight at most $1/2 + \epsilon$, even in the special case where P is positive semidefinite.*

Proof. We show that the system of equations arising from Håstad's 3-bit PCP (see [27], pages 24-25) satisfies the properties that we need. In his notation, the equations are generated by choosing $f \in \mathcal{F}_U$ and $g_1, g_2 \in \mathcal{F}_W$ where U and W, $U \subset W$, are randomly chosen and $\mathcal{F}_U, \mathcal{F}_W$ are the spaces of all ± 1 functions on $\{-1, +1\}^U$ and $\{-1, +1\}^W$, respectively. An equation corresponds to a 3-bit test on f, g_1, g_2 and its weight is the probability that the verifier performs this particular test. One variable is associated with $f \in \mathcal{F}_U$, indexing a bit in the Long Code of the first prover, and two variables are associated with

$g_1, g_2 \in \mathcal{F}_W$, indexing bits in the Long Code of the second prover. This defines a natural partition of variables into $\mathcal{X}$ and $\mathcal{Y}$.

The actual variables appearing in the equations are determined by the folding convention; for the second prover, let us denote them by y_j, y_k where $j = \phi(g_1), k = \phi(g_2)$. The particular function ϕ will not matter to us, as long as it is the same for both g_1 and g_2 (which is the case in [27]). P_{jk} is the probability that the selected variables corresponding to the second prover are y_j and y_k. Let $P_{jk}^{U,W}$ be the same probability, conditioned on a particular choice of U, W. Since P is a positive linear combination of $P^{U,W}$, it suffices to prove that each $P^{U,W}$ is positive semidefinite. The way that g_1, g_2 are generated (for given U, W) is that $g_1 : \{-1, +1\}^W \rightarrow \{-1, +1\}$ is uniformly random and $g_2(y) = g_1(y)f(y|_U)\mu(y)$, where $f : \{-1, +1\}^U \rightarrow \{-1, +1\}$ uniformly random and $\mu : \{-1, +1\}^W \rightarrow \{-1, +1\}$ is a "random noise", where $\mu(x) = 1$ with probability $1 - \epsilon$ and -1 with probability ϵ. The value of ϵ will be very small, certainly $\epsilon < 1/2$.

Both g_1 and g_2 are distributed uniformly (but not independently) in $\mathcal{F}_W$. The probability of sampling (g_1, g_2) is the same as the probability of sampling (g_2, g_1), hence $P^{U,W}$ is a symmetric matrix. It remains to prove positive semidefiniteness. Let us choose an arbitrary function $A : \mathcal{Y} \rightarrow \mathbb{R}$ and analyze

$$\sum_{j,k} P_{jk}^{U,W} A(j)A(k) = \mathbf{E}_{g_1,g_2}[A(\phi(g_1))A(\phi(g_2))] = \mathbf{E}_{g_1,f,\mu}[A(\phi(g_1))A(\phi(g_1 f\mu))]$$

where g_1, f, μ are sampled as described above. If we prove that this quantity is always nonnegative, then $P^{U,W}$ is positive semidefinite. Let $B : \mathcal{F}_W \rightarrow \mathbb{R}$, $B = A \circ \phi$; i.e., we want to prove $\mathbf{E}[B(g_1)B(g_1 f\mu)] \geq 0$. We can expand B using its Fourier transform,

$$B(g) = \sum_{\alpha \subseteq \{-1,+1\}^W} \hat{B}(\alpha)\chi_\alpha(g).$$

Here, $\chi_\alpha(g) = \prod_{x \in \alpha} g(x)$ are the Fourier basis functions. We obtain

$$\begin{aligned}
\mathbf{E}[B(g_1)B(g_1 f\mu)] &= \sum_{\alpha,\beta \subseteq \{-1,+1\}^W} \mathbf{E}[\hat{B}(\alpha)\chi_\alpha(g_1)\hat{B}(\beta)\chi_\beta(g_1 f\mu)] \\
&= \sum_{\alpha,\beta \subseteq \{-1,+1\}^W} \hat{B}(\alpha)\hat{B}(\beta) \prod_{x \in \alpha\Delta\beta} \mathbf{E}_{g_1}[g_1(x)] \cdot \mathbf{E}_f[\prod_{y \in \beta} f(y|_U)] \prod_{z \in \beta} \mathbf{E}_\mu[\mu(z)].
\end{aligned}$$

The terms for $\alpha \neq \beta$ are zero, since then $\mathbf{E}_{g_1}[g_1(x)] = 0$ for each $x \in \alpha\Delta\beta$. Therefore,

$$\mathbf{E}[B(g_1)B(g_1 f\mu)] = \sum_{\beta \subseteq \{-1,+1\}^W} \hat{B}^2(\beta)\mathbf{E}_f[\prod_{y \in \beta} f(y|_U)] \prod_{z \in \beta} \mathbf{E}_\mu[\mu(z)].$$

Now all the terms are nonnegative, since $\mathbf{E}_\mu[\mu(z)] = 1 - 2\epsilon > 0$ for every z and $\mathbf{E}_f[\prod_{y \in \beta} f(y|_U)] = 1$ or 0, depending on whether every string in $\{-1, +1\}^U$ is the projection of an even number of strings in β (in which case the product is 1) or not (in which case the expectation gives 0 by symmetry). To conclude,

$$\sum_{j,k} P_{jk}^{U,W} A(j)A(k) = \mathbf{E}[B(g_1)B(g_1 f\mu)] \geq 0$$

for any $A : \mathcal{Y} \rightarrow \mathbb{R}$, which means that each $P^{U,W}$ and consequently also P is positive semidefinite. $\qquad\square$

Now we are ready to show the following.

Theorem 4.3. *There is no polynomial-time $(\frac{3}{4} + \epsilon)$-approximation algorithm to maximize a nonnegative submodular function in succint representation, unless $P = NP$.*

Proof. We define a reduction from the system of linear equations provided by Lemma 4.2. For each variable $x_i \in \mathcal{X}$, we have two elements T_i, F_i and for each variable $y_j \in \mathcal{Y}$, we have two elements $\tilde{T}_j, \tilde{F}_j$.

Denote the set of equations by $\mathcal{E}$. Each equation e contains one variable from $\mathcal{X}$ and two variables from $\mathcal{Y}$. For each $e \in \mathcal{E}$, we define a submodular function $g_e(S)$ tailored to this structure. Assume that $S \subseteq \{T_i, F_i, \tilde{T}_j, \tilde{F}_j, \tilde{T}_k, \tilde{F}_k\}$, the elements corresponding to this equation; g_e does not depend on other than these 6 elements. We say that S is a valid triple, if it contains exactly one of each $\{T_i, F_i\}$.

- The value of each singleton T_i, F_i corresponding to a variable in $\mathcal{X}$ is 1.

- The value of each singleton $\tilde{T}_j, \tilde{F}_j$ corresponding to a variable in $\mathcal{Y}$ is 1/2.

- For $|S| < 3$, $g_e(S)$ is the sum of its singletons, except $g_e(\{T_i, F_i\}) = 1$ (*a weak pair*).

- For $|S| > 3$, $g_e(S) = g_e(\bar{S})$.

- If S is a valid triple satisfying e, let $g_e(S) = 2$ (*true triple*).

- If S is a valid triple not satisfying e, let $g_e(S) = 1$ (*false triple*).

- If S is an invalid triple containing exactly one of $\{T_i, F_i\}$ then $g_e(S) = 2$ (*invalid triple of type I*).

- If S is an invalid triple containing both/neither of $\{T_i, F_i\}$ then $g_e(S) = 3/2$ (*invalid triple of type II*).

Verifying that g_e is submodular is more tedious here; we omit the details. Let us move on to the important properties of g_e. A true triple gives value 2, while a false triple gives value 1. For invalid assignments of value 3/2, we can argue as before that a random valid assignment achieves expected value 3/2 as well, so we might as well choose a valid assignment. However, in this gadget we also have invalid triples of value 2 (type I - we cannot avoid this due to submodularity.) Still, we prove that the optimum is attained for a valid boolean assignment. The main argument is, roughly, that if there are many invalid triples of type I, there must be also many equations where we get value only 1 (a weak pair or its complement). For this, we use the positive semidefinite property from Lemma 4.2.

We define $f(S) = \sum_{e \in \mathcal{E}} w(e) g_e(S)$ where $w(e)$ is the weight of equation e. We claim that $\max f(S) = 1 + \max w_{SAT}$, where w_{SAT} is the weight of satisfied equations. First, for a given boolean assignment, the corresponding set S selecting T_i or F_i for each variable achieves value $f(S) = w_{SAT} \cdot 2 + (1 - w_{SAT}) \cdot 1 = 1 + w_{SAT}$. The non-trivial part is proving that the optimum $f(S)$ is attained for a set inducing a valid boolean assignment.

Consider any set S and define $V : \mathcal{E} \to \{-1, 0, +1\}$ where $V(e) = +1$ if S induces a satisfying assignment to equation e, $V(e) = -1$ if S induces a non-satisfying assignment to e and $V(e) = 0$ if S induces an invalid assignment to e. Also, define $A : \mathcal{Y} \to \{-1, 0, +1\}$, where $A(j) = |S \cap \{\tilde{T}_j, \tilde{F}_j\}| - 1$, i.e. $A(j) = 0$ if S induces a valid assignment to y_j, and $A(j) = \pm 1$ if S contains both/neither of $\tilde{T}_j, \tilde{F}_j$. Observe that for an equation e whose $\mathcal{Y}$-variables are y_j, y_k, only one of $V(e)$ and $A(j)A(k)$ can be nonzero. The gadget $g_e(S)$ is designed in such a way that

$$g_e(S) \leq \frac{1}{2}(3 - A(j)A(k) + V(e)).$$

This can be checked case by case: for valid assignments, $A(j)A(k) = 0$ and we get value 2 or 1 depending on $V(e) = \pm 1$. For invalid assignments, $V(e) = 0$; if at least one of the variables y_j, y_k has a valid assignment, then $A(j)A(k) = 0$ and we can get at most 3/2 (an invalid triple of type II). If both y_j, y_k are invalid and $A(j)A(k) = 1$, then we can get only 1 (a weak pair or its complement) and if $A(j)A(k) = -1$, we can get 2 (an invalid triple of type I). The total value is

$$f(S) = \sum_{e \in \mathcal{E}} w(e) g_e(S) \leq \sum_{e = (x_i, y_j, y_k)} w(e) \cdot \frac{1}{2}(3 - A(j)A(k) + V(e)).$$

Now we use the positive semidefinite property of our linear system, which means that

$$\sum_{e = (x, y_j, y_k)} w(e) A(j)A(k) = \sum_{j,k} P_{jk} A(j)A(k) \geq 0.$$

for any function A. Hence, $f(S) \le \frac{1}{2} \sum_{e \in \mathcal{E}} w(e)(3+V(e))$. Let us modify S into a valid boolean assignment by choosing randomly one of T_i, F_i for all variables such that S contains both/neither of T_i, F_i. Denote the new set by $\tilde{S}$ and the equations containing any randomly chosen variable by $\mathcal{R}$. We satisfy each equation in $\mathcal{R}$ with probability $1/2$, which gives expected value $3/2$ for each such equation, while the value for other equations remains unchanged.

$$\mathbf{E}[f(\tilde{S})] \quad = \quad \frac{3}{2} \sum_{e \in \mathcal{R}} w(e) + \frac{1}{2} \sum_{e \in \mathcal{E} \setminus \mathcal{R}} w(e)(3 + V(e)) = \frac{1}{2} \sum_{e \in \mathcal{E}} w(e)(3 + V(e)) \ge f(S).$$

This means that there is a set $\tilde{S}$ of optimal value, inducing a valid boolean assignment. $\square$

4.2 Value query complexity results

Finally, we prove that our $\frac{1}{2}$-approximation for symmetric submodular functions is optimal in the value oracle model. First, we present a similar result for the "random set" model, which illustrates some of the ideas needed for the more general result.

Proposition 4.4. *For any $\delta > 0$, there is $\epsilon > 0$ such that for any (random) sequence of queries $\mathcal{Q} \subseteq 2^X$, $|\mathcal{Q}| \le 2^{\epsilon n}$, there is a nonnegative submodular function f such that (with high probability) for all queries $Q \in \mathcal{Q}$,*

$$f(Q) \le \left(\frac{1}{4} + \delta\right) OPT.$$

Proof. Let $\epsilon = \frac{1}{32}\delta^2$ and fix a sequence $\mathcal{Q} \subseteq 2^X$ of $2^{\epsilon n}$ queries. We prove the existence of f by the probabilistic method. Consider functions corresponding to cuts in a complete bipartite directed graph on (C, D), $f_C(S) = |S \cap C| \cdot |\bar{S} \cap D|$. We choose a uniformly random $C \subseteq X$ and $D = X \setminus C$. The idea is that for any query, a typical C bisects both Q and its complement, which means that $f_C(Q)$ is roughly $\frac{1}{4}OPT$. We call a query $Q \in \mathcal{Q}$ "successful", if $f_C(Q) > (\frac{1}{4} + \delta)OPT$. Our goal is to prove that with high probability, C avoids any successful query.

We use the Chernoff bound: For any set $A \subseteq X$ of size a, $|A \cap C|$ is a sum of a independent random variables uniform in $\{0, 1\}$. Therefore, by Theorem 1.2,

$$\Pr[|A \cap C| > \frac{1}{2}(1 + \delta)|A|] + \Pr[|A \cap C| < \frac{1}{2}(1 - \delta)|A|] < 2e^{-\delta^2 a/2}.$$

Similarly, with probability at least $1 - 2e^{-2\delta^2 n}$, the size of C is in $[(\frac{1}{2} - \delta)n, (\frac{1}{2} + \delta)n]$, so we can assume this is the case. We have $OPT \ge (\frac{1}{4} - \delta^2)n^2 \ge \frac{1}{4}n^2/(1 + \delta)$ (for small $\delta > 0$). No query can achieve $f_C(Q) > (\frac{1}{4} + \delta)OPT \ge \frac{1}{16}n^2$ unless $|Q| \in [\frac{1}{16}n, \frac{15}{16}n]$, so we can assume this is the case for all queries. By Theorem 1.2, $\Pr[|Q \cap C| > \frac{1}{2}(1 + \delta)|Q|] < e^{-\delta^2 n/32}$ and $\Pr[|\bar{Q} \cap D| > \frac{1}{2}(1 + \delta)|\bar{Q}|] < e^{-\delta^2 n/32}$. If neither of these events occurs, the query is not successful, since $f_C(Q) = |Q \cap C| \cdot |\bar{Q} \cap D| < \frac{1}{4}(1 + \delta)^2|Q| \cdot |\bar{Q}| \le \frac{1}{16}(1 + \delta)^2 n^2 \le \frac{1}{4}(1 + \delta)^3 OPT \le (\frac{1}{4} + \delta)OPT$ (we can assume that $\delta > 0$ is sufficiently small).

For now, fix a sequence of queries. By the union bound, we get that the probability that any query is successful is at most $2^{\epsilon n} 2e^{-\delta^2 n/32} = 2(\frac{2}{e})^{\epsilon n}$. Thus with high probability, there is no successful query for C. Even for a random sequence, the probabilistic bound still holds by averaging over all possible sequences of queries. We can fix any C for which the bound is valid, and then the claim of the lemma holds for the submodular function f_C. $\square$

This means that in the model where an algorithm only samples a sequence of polynomially many sets and returns the one of maximal value, we cannot improve our $\frac{1}{4}$-approximation (Section 2). As we show next, this example can be modified for the model of adaptive algorithms with value queries, to show that our $\frac{1}{2}$-approximation for symmetric submodular functions is optimal, even among all adaptive algorithms.

14

Theorem 4.5. *For any $\epsilon > 0$, there are instances of nonnegative symmetric submodular maximization, such that there is no (adaptive, possibly randomized) algorithm using less than $e^{\epsilon^2 n/8}$ queries that always finds a solution of expected value at least $(\frac{1}{2} + \epsilon)OPT$.*

Proof. We construct a nonnegative symmetric submodular function on $[n] = C \cup D$, $|C| = |D| = n/2$, which has the following properties:

- $f(S)$ depends only on $k = |S \cap C|$ and $\ell = |S \cap D|$. Henceforth, we write $f(k, \ell)$ to denote the value of any such set.

- When $|k - \ell| \leq \epsilon n$, the function has the form
$$f(k, \ell) = (k + \ell)(n - k - \ell) = |S|(n - |S|),$$
i.e., the cut function of a complete graph. The value depends only on the size of S, and the maximum attained by such sets is $\frac{1}{4}n^2$.

- When $|k - \ell| > \epsilon n$, the function has the form
$$f(k, \ell) = k(n - 2\ell) + (n - 2k)\ell - O(\epsilon n^2),$$
close to the cut function of a complete bipartite graph on (C, D) with edge weights 2. The maximum in this range is $OPT = \frac{1}{2}n^2(1 - O(\epsilon))$, attained for $k = \frac{1}{2}n$ and $\ell = 0$ (or vice versa).

If we construct such a function, we can argue as follows. Consider any algorithm, for now deterministic. (For a randomized algorithm, let us condition on its random bits.) Let the partition (C, D) be uniformly random and unknown to the algorithm. The algorithm issues some queries Q to the value oracle. Call Q "unbalanced", if $|Q \cap C|$ differs from $|Q \cap D|$ by more than ϵn. For any query Q, the probability that Q is unbalanced is at most $2e^{-\epsilon^2 n/4}$, by Theorem 1.2. Therefore, for any fixed sequence of $e^{\epsilon^2 n/8}$ queries, the probability that any query is unbalanced is still at most $2e^{\epsilon^2 n/8} \cdot e^{-\epsilon^2 n/4} = 2e^{-\epsilon^2 n/8}$. As long as queries are balanced, the algorithm gets the same answer regardless of (C, D). Hence, it follows the same path of computation and issues the same queries. With probability at least $1 - 2e^{-\epsilon^2 n/8}$, all its queries will be balanced and it will never find out any information about the partition (C, D). For a randomized algorithm, we can now average over its random choices; still, with probability at least $1 - 2e^{-\epsilon^2 n/8}$ the algorithm will never query any unbalanced set.

Alternatively, consider a function $g(S)$ which is defined by $g(S) = |S|(n - |S|)$ for *all sets* S. We proved that with high probability, the algorithm will never query a set where $f(S) \neq g(S)$ and hence cannot distinguish between the two instances. However, $\max_S f(S) = \frac{1}{2}n^2(1 - O(\epsilon))$, while $\max_S g(S) = \frac{1}{4}n^2$. This means that there is no $(\frac{1}{2} + \epsilon)$-approximation algorithm with a subexponential number of queries, for any $\epsilon > 0$.

It remains to construct the function $f(k, \ell)$ and prove its submodularity. For convenience, assume that ϵn is an integer. In the range where $|k - \ell| \leq \epsilon n$, we already defined $f(k, \ell) = (k + \ell)(n - k - \ell)$. In the range where $|k - \ell| \geq \epsilon n$, let us define
$$f(k, \ell) = k(n - 2\ell) + (n - 2k)\ell + \epsilon^2 n^2 - 2\epsilon n|k - \ell|.$$

The ϵ-terms are chosen so that $f(k, \ell)$ is a smooth function on the boundary of the two regions. E.g., for $k - \ell = \epsilon n$, we get $f(k, \ell) = (2k - \epsilon n)(n - 2k + \epsilon n)$ for both expressions. Moreover, the marginal values also extend smoothly. Consider an element $i \in C$ (for $i \in D$ the situation is symmetric). The marginal value of i added to a set S is $f(S \cup \{i\}) - f(S) = f(k + 1, \ell) - f(k, \ell)$. We split into three cases:

- If $k - \ell < -\epsilon n$, we have $f(k + 1, \ell) - f(k, \ell) = (n - 2\ell) + (-2\ell) + 2\epsilon n = (1 + 2\epsilon)n - 4\ell$.

- If $-\epsilon n \leq k - \ell < \epsilon n$, we have $f(k + 1, \ell) - f(k, \ell) = (k + 1 + \ell)(n - k - 1 - \ell) - (k + \ell)(n - k - \ell) = (n - k - 1 - \ell) - (k + 1 + \ell) = n - 2k - 2\ell - 2$. In this range, this is between $(1 \pm 2\epsilon) - 4\ell$.

- If $k - \ell \geq \epsilon n$, we have $f(k + 1, \ell) - f(k, \ell) = (n - 2\ell) + (-2\ell) - 2\epsilon n = (1 - 2\epsilon)n - 4\ell$.

Now it is easy to see that the marginal value is decreasing in both k and ℓ, in each range and also across ranges. $\qquad \square$

5 Conclusions and open questions

Subsequently to this work, a new algorithm has been proposed for the maximization of a non-negative (non-monotone) submodular function, which achieves an improved 0.41-approximation [41]. However, the existence of a 1/2-approximation algorithm still remains an open question. For non-adaptive algorithms, a remaining question is whether our algorithms can be derandomized and implemented by querying only a predetermined collection of polynomially many sets.

Even considering our query complexity results, it is still conceivable that a better than 1/2 approximation might be achieved in a model of computation requiring an explicit representation of $f(S)$. Our NP-hardness results in this case are still quite far away from 1/2. Considering the known approximation results for Max Cut, such an improvement would most likely require semidefinite programming or spectral techniques. Currently, we do not know how to implement such an approach.

References

[1] A. Ageev and M. Sviridenko. An 0.828 approximation algorithm for uncapacitated facility location problem, *Discrete Applied Mathematics* 93:2–3 (1999), 149–156.

[2] P. Alimonti. Non-oblivious local search for MAX 2-CCSP with application to MAX DICUT, *Proc. of the 23rd International Workshop on Graph-theoretic Concepts in Computer Science* (1997).

[3] N. Alon and J. H. Spencer. The Probabilistic Method (2nd edition). Wiley Interscience, 2000.

[4] P. Austrin. Improved inapproximability for submodular maximization, *Proc. of 13th APPROX*, 2010.

[5] G. Calinescu, C. Chekuri, M. Pál and J. Vondrák. Maximizing a submodular set function subject to a matroid constraint, *Proc. of 12th IPCO* (2007), 182–196.

[6] G. Calinescu, C. Chekuri, M. Pál, J. Vondrák. Maximizing a submodular set function subject to a matroid constraint, to appear in *SIAM J. on Computing*.

[7] V. Cherenin. Solving some combinatorial problems of optimal planning by the method of successive calculations, *Proc. of the Conference of Experiences and Perspectives of the Applications of Mathematical Methods and Electronic Computers in Planning* (in Russian), Mimeograph, Novosibirsk (1962).

[8] G. Cornuejols, M. Fischer and G. Nemhauser. Location of bank accounts to optimize float: an analytic study of exact and approximation algorithms, *Management Science* 23 (1977), 789–810.

[9] G. Cornuejols, M. Fischer and G. Nemhauser. On the uncapacitated location problem, *Annals of Discrete Math* 1 (1977), 163–178.

[10] J. Edmonds. Matroids, submodular functions and certain polyhedra, *Combinatorial Structures and Their Applications* (1970), 69–87.

[11] U. Feige and M. X. Goemans. Approximating the value of two-prover systems, with applications to MAX-2SAT and MAX-DICUT, *Proc. of the 3rd Israel Symposium on Theory and Computing Systems*, Tel Aviv (1995), 182–189.

[12] U. Feige. A threshold of $\ln n$ for approximating Set Cover, *Journal of the ACM* 45 (1998), 634–652.

[13] U. Feige. Maximizing social welfare when utility functions are subadditive, *Proc. of 38th ACM STOC* (2006), 41–50.

[14] U. Feige, V. Mirrokni and J. Vondrák. Maximizing non-monotone submodular functions, *Proc. of 48th IEEE FOCS* (2007), 461–471.

[15] U. Feige and J. Vondrák. Approximation algorithms for combinatorial allocation problems: Improving the factor of $1 - 1/e$, *Proc. of 47th IEEE FOCS* (2006), 667–676.

[16] M. L. Fisher, G. L. Nemhauser and L. A. Wolsey. An analysis of approximations for maximizing submodular set functions II, *Mathematical Programming Study* 8 (1978), 73–87.

[17] S. Fujishige. Canonical decompositions of symmetric submodular systems, *Discrete Applied Mathematics* 5 (1983), 175–190.

[18] L. Fleischer, S. Fujishige and S. Iwata. A combinatorial, strongly polynomial-time algorithm for minimizing submodular functions, *Journal of the ACM* 48:4 (2001), 761–777.

[19] A. Frank. Matroids and submodular functions, *Annotated Biblographies in Combinatorial Optimization* (1997), 65–80.

[20] M. Goemans, N. Harvey, S. Iwata and V. Mirrokni. Approximating submodular functions everywhere, *Proc. of 20th ACM-SIAM SODA* (2009), 535–544.

[21] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming, *Journal of the ACM* 42 (1995), 1115–1145.

[22] B. Goldengorin, G. Sierksma, G. Tijssen and M. Tso. The data correcting algorithm for the minimization of supermodular functions, *Management Science*, 45:11 (1999), 1539–1551.

[23] B. Goldengorin, G. Tijsssen and M. Tso. The maximization of submodular functions: Old and new proofs for the correctness of the dichotomy algorithm, *SOM Report*, University of Groningen (1999).

[24] V. Guruswami. Inapproximability results for set splitting and satisfiability problems with no mixed clauses, *Algorithmica* 38 (2004), 451–469.

[25] V. Guruswami and S. Khot. Hardness of Max 3-SAT with no mixed clauses, *Proc. of 20th IEEE Conference on Computational Complexity* (2005), 154–162.

[26] E. Halperin and U. Zwick. Combinatorial approximation algorithms for the maximum directed cut problem, *Proc. of 12th ACM-SIAM SODA* (2001), 1–7.

[27] J. Håstad. Some optimal inapproximability results, *Journal of the ACM* 48 (2001), 798–869.

[28] V. R. Khachaturov. Mathematical methods of regional programming (in Russian), *Nauka, Moscow*, 1989.

[29] S. Khot, G. Kindler, E. Mossel and R. O'Donnell. Optimal inapproximability results for MAX-CUT and other two-variable CSPs? *Proc. of 45th IEEE FOCS* (2004), 146–154.

[30] A. Kulik, H. Shachnai and T. Tamir. Maximizing submodular functions subject to multiple linear constraints, *Proc. of 20th ACM-SIAM SODA* (2009).

[31] D. Livnat, M. Lewin and U. Zwick. Improved rounding techniques for the MAX 2-SAT and MAX DI-CUT problems. *Proc. of 9th IPCO* (2002), 67–82.

[32] J. Lee, V. Mirrokni, V. Nagarajan and M. Sviridenko. Non-monotone submodular maximization under matroid and knapsack constraints, *Proc. of 41th ACM STOC* (2009), 323-332.

[33] H. Lee, G. Nemhauser and Y. Wang. Maximizing a submodular function by integer programming: Polyhedral results for the quadratic case, *European Journal of Operational Research* 94 (1996), 154–166.

[34] J. Lee, M. Sviridenko and J. Vondrák. Submodular maximization over multiple matroids via generalized exchange properties, *Proc. of APPROX 2009*, 244–257.

[35] L. Lovász. Submodular functions and convexity. A. Bachem et al., editors, *Mathematical Programmming: The State of the Art*, 235–257.

[36] V. Mirrokni, M. Schapira and J. Vondrák. Tight information-theoretic lower bounds for welfare maximization in combinatorial auctions, *Proc. of ACM Conference on Electronic Commerce*, 2008, 70–77.

[37] M. Minoux. Accelerated greedy algorithms for maximizing submodular functions, J. Stoer, ed., *Actes Congress IFIP, Springer Verlag, Berlin* (1977), 234–243.

[38] E. Mossel, R. O'Donnell and K. Oleszkiewicz. Noise stability of functions with low influences: invariance and optimality, *Proc. of 46th IEEE FOCS* (2005), 21–30.

[39] G. L. Nemhauser, L. A. Wolsey and M. L. Fisher. An analysis of approximations for maximizing submodular set functions I, *Mathematical Programming* 14 (1978), 265–294.

[40] G. L. Nemhauser and L. A. Wolsey. Best algorithms for approximating the maximum of a submodular set function, *Math. Oper. Research*, 3(3):177–188, 1978.

[41] S. Oveis Gharan and J. Vondrák. Submodular maximization by simulated annealing, *Proc. of 22nd ACM-SIAM SODA* (2011), 1098–1117.

[42] M. Queyranne. A combinatorial algorithm for minimizing symmetric submodular functions, *Proc. of 6th ACM-SIAM SODA* (1995), 98–101.

[43] T. Robertazzi and S. Schwartz. An accelated sequential algorithm for producing D-optimal designs, *SIAM Journal on Scientific and Statistical Computing* 10 (1989), 341–359.

[44] A. Schäfer and M. Yannakakis. Simple local search problems that are hard to solve, *SIAM J. Comput.* 20:1 (1991), 56–87.

[45] A. Schrijver. A combinatorial algorithm minimizing submodular functions in strongly polynomial time, *Journal of Combinatorial Theory, Series B* 80 (2000), 346–355.

[46] M. Sviridenko. A note on maximizing a submodular set function subject to a knapsack constraint, *Operations Research Letters* 32 (2004), 41–43.

[47] Z. Svitkina and L. Fleischer. Submodular approximation: Sampling-based algorithms and lower bounds, *Proc. of 49th IEEE FOCS* (2008), 697–706.

[48] L. Trevisan. Max Cut and the smallest eigenvalue, *Proc. of 41th ACM STOC* (2009), 263–272.

[49] J. Vondrák. Optimal approximation for the submodular welfare problem in the value oracle model, *Proc. of 40th ACM STOC* (2008), 67–74.

[50] J. Vondrák. Symmetry and approximability of submodular maximization problems, *Proc. of 50th IEEE FOCS* (2009), 651–670.