

Submodular Maximization subject to a Knapsack Constraint: Combinatorial Algorithms with Near-optimal Adaptive Complexity*

Georgios Amanatidis[†] Federico Fusco[‡] Philip Lazos[‡] Stefano Leonardi[‡]
 Alberto Marchetti Spaccamela[‡] Rebecca Reiffenhäuser[‡]

February 17, 2021

Abstract

The growing need to deal with massive instances motivates the design of algorithms balancing the quality of the solution with applicability. For the latter, an important measure is the *adaptive complexity*, capturing the number of sequential rounds of parallel computation needed. In this work we obtain the first *constant factor* approximation algorithm for non-monotone submodular maximization subject to a knapsack constraint with *near-optimal* $O(\log n)$ adaptive complexity. Low adaptivity by itself, however, is not enough: one needs to account for the total number of function evaluations (or value queries) as well. Our algorithm asks $\tilde{O}(n^2)$ value queries, but can be modified to run with only $\tilde{O}(n)$ instead, while retaining a low adaptive complexity of $O(\log^2 n)$. Besides the above improvement in adaptivity, this is also the first *combinatorial* approach with sublinear adaptive complexity for the problem and yields algorithms comparable to the state-of-the-art even for the special cases of cardinality constraints or monotone objectives. Finally, we showcase our algorithms’ applicability on real-world datasets.

1 Introduction

Submodular optimization is a very popular topic relevant to various research areas as it captures the natural notion of *diminishing returns*. Its numerous applications include viral marketing [27, 29], data summarization [39, 35], feature selection [13, 14, 36] and clustering [34]. Prominent examples from combinatorial optimization are cut functions in graphs and coverage functions.

Submodularity is often implicitly associated with monotonicity, and many results rely on that assumption. However, non-monotone submodular functions do naturally arise in applications, either directly or from combining monotone submodular objectives and modular *penalization* or *regularization* terms [27, 39, 35, 6, 1]. Additional constraints, like cardinality, matroid, knapsack, covering, and packing constraints, are prevalent in applications and have been extensively studied. In this list, *knapsack* constraints are among the most natural, as they capture limitations on budget, time, or size of the elements. Like matroid constraints, they generalize cardinality constraints, yet they are not captured by the former.

*This work was supported by the ERC Advanced Grant 788893 AMDROMA “Algorithmic and Mechanism Design Research in Online Markets”, the MIUR PRIN project ALGADIMAR “Algorithms, Games, and Digital Markets”, and the NWO Veni project No. VI.Veni.192.153.

[†]Department of Mathematical Sciences, University of Essex, UK, and Institute for Logic, Language and Computation, University of Amsterdam, The Netherlands. Email: georgios.amanatidis@essex.ac.uk

[‡]Department of Computer, Control, and Management Engineering “Antonio Ruberti”, Sapienza University of Rome, Italy. Email: {fuscof, lazos, leonardi, alberto, rebeccar}@diag.uniroma1.it

The main computational bottleneck in submodular optimization comes from the necessity to repeatedly evaluate the objective function for various candidate sets. These so-called *value queries* are often notoriously heavy to compute, e.g., for exemplar-based clustering [15], log-determinant of submatrices [28], and accuracy of ML models [13, 30]. With real-world instances of these problems growing to enormous sizes, simply reducing the number of queries is not always sufficient and parallelisation has become an increasingly central paradigm. However, classic results in the area, often based on the greedy method, are inherently sequential: the intuitive approach of building a solution element-by-element contradicts the requirement of running *independent* computations on many machines in parallel. The degree to which an algorithm can be parallelized is measured by the notion of *adaptive complexity*, or adaptivity, introduced in Balkanski et al. [3]. It is defined as the number of sequential rounds of parallel computation needed to terminate. In each of these rounds, polynomially many value queries may be asked, but they can only depend on the answers to queries issued in past rounds.

Contribution. We design the first combinatorial randomized algorithms for maximizing a (possibly) non-monotone submodular function subject to a knapsack constraint that combine constant approximation, low adaptive complexity, and a small number of queries. In particular, we obtain

- a 9.465-approximation algorithm, PARKKNAPSACK, that has $O(\log n)$ adaptivity and uses $O(n^2 \log^2 n)$ value queries. This is the *first* constant factor approximation to the problem with optimal adaptive complexity up to a $O(\log \log n)$ factor (Theorem 1).
- a variant of our algorithm with the same approximation, near-optimal $O(n \log^3 n)$ query complexity, and $O(\log^2 n)$ adaptivity (Theorem 2). This is the first constant factor approximation algorithm that uses only $\tilde{O}(n)$ queries and has *sublinear* adaptivity.
- 3-approximation algorithms for *monotone* objectives that combine $O(\log n)$ adaptivity with $O(n^2 \log^2 n)$ total queries, and $O(\log^2 n)$ adaptivity with $O(n \log^3 n)$ queries, respectively (Theorem 3). Even in the monotone setting, the latter is the first $O(1)$ -approximation algorithm combining $\tilde{O}(n)$ queries and sublinear adaptivity.
- 5.83-approximation algorithms for *cardinality* constraints that match or surpass the state-of-the-art when it comes to the combination of approximation, adaptivity and total queries (Theorem 4).

See Table 1 for an overview of our results.

Technical Challenges. Like existing work for cardinality or matroid constraints, e.g., [5, 6], in order to reduce the adaptive complexity we iteratively sample sequences of feasible elements and add large chunks of them to our solution. However, knapsack constraints do not allow for the elegant counting arguments used in the case of cardinality or matroid constraints. The reason is that while the latter can be interpreted as a 1-independence system, a knapsack constraint induces a $\Theta(n)$ -independence system, leading to poor results when naively adjusting existing approaches. A natural and very successful way of circumventing the resulting difficulties is to turn towards a *continuous* version of the problem. This, however requires evaluating the objective function also for *fractional* sets, i.e., such algorithms require access to an oracle for the multilinear relaxation and its gradient. Typically, these values are estimated by sampling, requiring $\tilde{\Theta}(n^2)$ samples (see Chekuri and Quanrud [9]). Our choice to avoid the resulting increase in query complexity and deal directly with the discreteness of the problem calls for specifically tailored algorithmic approaches. Most

Reference	Objective	Constraint	Approx.	Adaptivity	Queries
Ene et al. [19]	General	Knapsack	$e + \varepsilon$	$O(\log^2 n)$	$\tilde{O}(n^2)$
Theorem 1 (this work)	General	Knapsack	$9.465 + \varepsilon$	$O(\log n)$	$\tilde{O}(n^2)$
Theorem 2 (this work)	General	Knapsack	$9.465 + \varepsilon$	$O(\log^2 n)$	$\tilde{O}(n)$
Ene et al. [19]	Monotone	Knapsack	$\frac{e}{e-1} + \varepsilon$	$O(\log n)$	$\tilde{O}(n^2)$
Chekuri and Quanrud [9]	Monotone	Knapsack	$\frac{e}{e-1} + \varepsilon$	$O(\log n)$	$\tilde{O}(n^2)$
Theorem 3 (this work)	Monotone	Knapsack	$3 + \varepsilon$	$O(\log n)$	$\tilde{O}(n^2)$
Theorem 3 (this work)	Monotone	Knapsack	$3 + \varepsilon$	$O(\log^2 n)$	$\tilde{O}(n)$
Ene and Nguyen [17]	General	k -Cardinality	$e + \varepsilon$	$O(\log n)$	$\tilde{O}(nk^2)$
Kuhnle [31]	General	k -Cardinality	$6 + \varepsilon$	$O(\log n)$	$\tilde{O}(n)$
Kuhnle [31]	General	k -Cardinality	$5.18 + \varepsilon$	$O(\log^2 n)$	$\tilde{O}(n)$
Theorem 4 (this work)	General	k -Cardinality	$5.83 + \varepsilon$	$O(\log n)$	$\tilde{O}(nk)$
Theorem 4 (this work)	General	k -Cardinality	$5.83 + \varepsilon$	$O(\log n \log k)$	$\tilde{O}(n)$

Table 1: Our results—main result highlighted—compared to the state-of-the-art for low-adaptivity. Bold indicates the best result(s) in each setting. In the last two columns the dependence on ε is omitted; in the last column only the leading terms are stated.

crucially, our main subroutine THRESHSEQ needs to balance a suitable definition of *good quality* candidates with a way to also reduce the size (not simply by cardinality, but a combination of overall cost and absolute marginal values) of the candidate set by a constant factor in each adaptive round.

Both these goals are further hindered by our second main challenge, non-monotonicity. In presence of elements with negative marginals, not only is it harder to maintain a good quality of our solution, but size measures like the overall absolute marginals of our candidate sets are no longer inclusion-monotone. In fact, even one such element can arbitrarily deteriorate intuitive quality measures like the overall marginal density of the candidate set, causing a new adaptive round. Our approach combines carefully designed stopping times in THRESHSEQ with a separate handling of the elements responsible for most of the above mentioned *discreteness issues*, i.e., elements with cost less than $1/n$ of the budget and elements of maximum value.

Related Work. Submodular maximization has been studied extensively since the seminal work of Nemhauser et al. [37]. For *monotone* submodular objectives subject to a knapsack constraint the $\frac{e}{e-1}$ -approximation algorithm of Sviridenko [38] is best-possible, unless $P = NP$ [22]. For the *non-monotone* case, a number of continuous greedy approaches [24, 32, 11] led to the current best factor of e when a knapsack, or any downward closed, constraint is involved. Combinatorial approaches [25, 1] achieve somewhat worse approximation, but are often significantly faster and thus relevant in practice.

Adaptive complexity for submodular maximization was introduced by Balkanski and Singer [2] for monotone objectives and a cardinality constraint, where they achieved an $O(1)$ approximation with $O(\log n)$ adaptivity, along with an almost matching lower bound: to get an $o(\log n)$ approximation, adaptivity must be $\Omega(\frac{\log n}{\log \log n})$. This result has been then improved [4, 16, 20] and recently Breuer et al. [6] achieved an optimal $\frac{e}{e-1}$ -approximation in $O(\log n \log^2 \log k)$ adaptive rounds and $O(n \log \log k)$ query complexity, where k is the cardinality constraint.

The study of adaptivity for non-monotone objectives was initiated by Balkanski et al. [3] again for a cardinality constraint, showing a constant approximation in $O(\log^2 n)$ adaptive rounds, later improved by [21, 17, 31]. Non-monotone maximization is also interesting in the unconstrained

scenario. Recently, Ene et al. [18] and Chen et al. [12] achieved a $2 + \varepsilon$ approximation with constant adaptivity depending only on ε . Note that the algorithm of Chen et al. [12] needs only $\tilde{O}(n)$ value queries, where the \tilde{O} hides terms poly-logarithmic in ε^{-1} and n .

Richer constraints, e.g., matroids and multiple packing constraints, have also been studied [5, 19, 10, 9]. For knapsack constraints (as a special case of packing constraints) Ene et al. [19] and Chekuri and Quanrud [9] provide low adaptivity results— $O(\log^2 n)$ for non-monotone and $O(\log n)$ for monotone—via continuous approaches (see Table 1; notice that the query complexity of these algorithms is stated with respect to queries to f and not to its multilinear extension). Chekuri and Quanrud [9] also provide two combinatorial algorithms for the monotone case: one with optimal approximation and adaptivity but $O(n^4)$ value queries, and one with linear query complexity, optimal adaptivity but an approximation factor parameterized by $\max_{x \in \mathcal{N}} c(x)$ which can be arbitrarily bad.

2 Preliminaries

Let $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}$ be a set function over a ground set \mathcal{N} of n elements. For $S, T \subseteq \mathcal{N}$, $f(S|T)$ denotes the *marginal value* of S with respect to T , i.e., $f(S \cup T) - f(T)$. To ease notation, we write $f(x|T)$ instead of $f(\{x\}|T)$. The function f is *non-negative* if $f(S) \geq 0$, $\forall S \subseteq \mathcal{N}$, *monotone* if $f(S) \leq f(T)$, $\forall S, T \subseteq \mathcal{N}$, and *submodular* if $f(x|T) \leq f(x|S)$, $\forall S, T \subseteq \mathcal{N}$ with $S \subseteq T$ and $x \notin T$.

We study non-negative, possibly *non-monotone*, submodular maximization under a *knapsack constraint*. Formally, we are given a budget $B > 0$, a non-negative submodular function f and a non-negative additive cost function $c : 2^{\mathcal{N}} \rightarrow \mathbb{R}_{>0}$. The goal is to find $O^* \in \arg \max_{T \subseteq \mathcal{N}: c(T) \leq B} f(T)$. Let $OPT = f(O^*)$ denote the value of such an optimal set. Given a (randomized) algorithm for the problem, let ALG denote the expected value of its output. We say that the algorithm is a β -approximation algorithm if $ALG \cdot \beta \geq OPT$. Throughout this work, we assume, without loss of generality, that $\max_{x \in \mathcal{N}} c(x) \leq B$.

We assume access to f through value queries, i.e., for each $S \subseteq \mathcal{N}$, an oracle returns $f(S)$ in constant time. Given such an oracle for f , the *adaptive complexity* or *adaptivity* of an algorithm is the minimum number of rounds in which the algorithm makes $O(\text{poly}(n))$ *independent* queries to the evaluation oracle. In each adaptive round the queries may *only* depend on the answers to queries from past rounds. With respect to the same oracle, the *query complexity* of an algorithm is the total number of value queries it makes.

We finally state some widely known properties of submodular functions that are extensively used in the rest of the paper. The first lemma summarizes two equivalent definitions of submodular functions shown by Nemhauser et al. [37].

Lemma 1. *Let $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}$ be a submodular function and S, T, U be any subsets of \mathcal{N} , with $S \subseteq T$. Then i) $f(U|T) \leq f(U|S)$, ii) $f(S|T) \leq \sum_{x \in S} f(x|T)$.*

The second lemma, Lemma 2.2 of Buchbinder et al. [7], is an important tool for tackling non-monotonicity.

Lemma 2 (Sampling Lemma). *Let $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}$ be a submodular function, $X \subseteq \mathcal{N}$ and X_p be a random subset of X , where each element of X is contained with probability at most p . Then $\mathbb{E}[f(X_p)] \geq (1 - p)f(\emptyset)$.*

Finally, we assume access to SUBMODMAX, an unconstrained submodular maximization oracle. For instance, this can be implemented via the combinatorial algorithm of Chen et al. [12], which outputs a $(2 + \varepsilon)$ -approximation of $\max_{T \subseteq \mathcal{N}} f(T)$ for a given precision ε in $O(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon})$ adaptive

rounds and linear query complexity. For our experiments, we use the much simpler 4-approximation of Feige et al. [23], which has adaptive complexity 1.

3 Non-Monotone Submodular Maximization

To achieve sublinear adaptivity we need to add large chunks of elements to the solution without using intermediate value queries. The sequence of elements that are candidates to be added to the current solution is randomly drawn in `SAMPLESEQ`. This subroutine receives as input a partial solution S , a set of feasible elements X and a budget, and outputs a sequence A each element of which is sequentially drawn uniformly at random among the remaining elements of X that do not cause $S \cup A$ to exceed the budget. We do, however, need to restrict ourselves to only adding a *suitable prefix* of A ; with each element added, the original “good” quality of the leftover candidates in X can quickly deteriorate.

The selection of the prefix of the sequence $A = [a_1, \dots, a_d]$ to be added to the current solution S is then done by `THRESHSEQ`. Given a threshold τ , we add to S a prefix $A_i = [a_1, \dots, a_i]$ such that for all $j < i$ the average contribution to $S \cup A_j$ of the elements in $X \setminus A_j$ is comparable to τ . Then the expected value of $f(A_i | S)$ should be comparable to $\tau \mathbb{E}[c(A_i)]$. In order to compute A_i in one single parallel round, one can *a posteriori* compute for each prefix A_j of A the *a priori* (with respect to the uniform samples) expected marginal value of a_{j+1} ; with a_{j+1} drawn uniformly at random from the elements in $X \setminus A_j$ still fitting the budget, this means simply averaging over their marginal densities. Since all the value queries depend only on S and A , finding the prefix needs only a single adaptive round.

The crucial difficulty lies in the fact that limiting the expected marginal density is insufficient to bound the number of adaptive steps. In the worst case, a single very negative element could trigger this condition. We circumvent the resulting adaptive complexity of up to n by imposing two different stopping conditions, corresponding to i^* and j^* in `THRESHSEQ`. The *cost condition* is triggered once an ε -fraction of all remaining candidates’ cost is due to elements that are no longer *good*, i.e., they now have marginal density below τ . The *value condition* is triggered at most ℓ times, which happens whenever the elements with negative marginal value make up an ε -fraction of the entire leftover marginal value. Now, in each adaptive step, either the overall cost or the summed-up marginal contributions of the candidate set decrease by a factor of $(1 - \varepsilon)$. These observations are formalized below.

Algorithm 1: `SAMPLESEQ`(S, X, B)

- 1: **Input:** current solution S , set X of remaining elements and budget $B > 0$
 - 2: $A \leftarrow []$, $i \leftarrow 1$
 - 3: **while** $X \neq \emptyset$ **do**
 - 4: Draw a_i uniformly at random from X
 - 5: $A \leftarrow [a_1, \dots, a_{i-1}, a_i]$
 - 6: $X \leftarrow \{x \in X \setminus \{a_i\} : c(x) + c(A) + c(S) \leq B\}$
 - 7: $i \leftarrow i + 1$
 - 8: **return** $A = [a_1, a_2, \dots, a_d]$
-

Lemma 3. *Let $\kappa(X) = \max_{x,y \in X} c(x)/c(y)$. Then `THRESHSEQ` runs in $O\left(\frac{1}{\varepsilon} \log(n\kappa(X)) + \ell\right)$ adaptive rounds and issues $O\left(n^2 \left(\frac{1}{\varepsilon} \log n\kappa(X) + \ell\right)\right)$ value queries.*

Proof. The adaptive rounds correspond to iterations of the while loop. In fact, once a new sequence is drawn by SAMPLESEQ, all the value queries needed are deterministically induced by it and hence can be assigned to different independent machines. Gathering this information we can determine k^* and start another iteration of the while loop. Bounding the number of such iterations where the value condition is triggered is easy, since it is forced to be at most ℓ . For the cost condition we use the geometric decrease in the total cost of X : every time it is triggered, the total cost of the feasible elements X is decreased by at least a $(1 - \varepsilon)$ factor. At the beginning of the algorithm, that cost is at most Cn , with $C = \max_{x \in X} c(x)$, and it needs to decrease below $c = \min_{x \in X} c(x)$ to ensure that $X = \emptyset$. Call r the number of such rounds. In the worst case we need $Cn(1 - \varepsilon)^r < c$, meaning that the adaptivity is upper bounded by $\frac{1}{\varepsilon} \log(n\kappa(X)) + \ell$. Finally, notice that the query complexity is just a n^2 factor greater than the adaptivity: each adaptive round contains $O(n^2)$ value queries, since the length of the sequence output by SAMPLESEQ may be linear in n and for each prefix the value of the marginals of all the remaining elements has to be considered. \square

Having settled the adaptive and query complexity of THRESHSEQ, we move to proving that our conditions ensure good expected marginal density.

Algorithm 2: THRESHSEQ($X, \tau, \varepsilon, \ell, B$)

```

1: Input: set  $X$  of elements, threshold  $\tau > 0$ , precision  $\varepsilon \in (0, 1)$ , parameter  $\ell$  and budget  $B$ 
2:  $S \leftarrow \emptyset$ ,  $\text{ctr} \leftarrow 0$ 
3:  $X \leftarrow \{x \in X : f(x) \geq \tau c(x)\}$ 
4: while  $X \neq \emptyset$  and  $\text{ctr} < \ell$  do
5:    $[a_1, a_2, \dots, a_d] \leftarrow \text{SAMPLESEQ}(S, X, B)$ ;
6:   for  $i = 1, \dots, d$  do
7:      $A_i \leftarrow \{a_1, a_2, \dots, a_i\}$ 
8:      $X_i \leftarrow \{a \in X \setminus A_i : c(a) + c(S \cup A_i) \leq B\}$ ;
9:      $G_i \leftarrow \{a \in X_i : f(a | S \cup A_i) \geq \tau \cdot c(a)\}$ 
10:     $E_i \leftarrow \{a \in X_i : f(a | S \cup A_i) < 0\}$ 
11:     $i^* \leftarrow \min\{i : c(G_i) \leq (1 - \varepsilon)c(X)\}$ 
12:     $j^* \leftarrow \min\left\{j : \sum_{x \in G_j} \varepsilon f(x | S \cup A_j) \leq \sum_{x \in E_j} |f(x | S \cup A_j)|\right\}$ 
13:     $k^* \leftarrow \min\{i^*, j^*\}$ 
14:     $S \leftarrow S \cup A_{k^*}$ 
15:     $X \leftarrow G_{k^*}$ 
16:    if  $j^* < i^*$  then
17:       $\text{ctr} \leftarrow \text{ctr} + 1$ 
18: return  $S$ 

```

Lemma 4. For any $X, \tau, \varepsilon \in (0, 1), \ell$ and b , the random set S output by THRESHSEQ is such that $\mathbb{E}[f(S)] \geq (1 - \varepsilon)^2 \tau \mathbb{E}[c(S)]$ and $c(S) \leq B$.

Proof. We first note that $c(S) \leq B$ with probability 1 since SAMPLESEQ always returns feasible sequences. The algorithm adds a chunk of elements to the solution in each iteration of the *while* loop. This, along with the fact that each of these chunks is an ordered prefix of a sequence output by SAMPLESEQ, induces a total ordering on the elements in S . To facilitate the presentation of this proof, we imagine that the elements of S are added one after the other, according to this total order.

Let us call the t -th such element s_t , and let \mathcal{F}_t denote the filtration capturing the randomness of the algorithm up to, but excluding, the adding of s_t to its chunk's random sequence. We show that whenever any s_t is added, its expected marginal density is at least $(1 - \varepsilon)^2\tau$.

Fix some s_t and consider the iteration of the while loop in which it is added to the solution. We denote with S_{old} the partial solution at the beginning of that while loop, with X the candidate set $\{x \in \mathcal{N} : f(x | S_{old}) \geq \tau \cdot c(x), c(x) + c(S_{old}) \leq B\}$ at that point, and with A the sequence drawn in that iteration by SAMPLESEQ. Let $A_{(t)}$ be the prefix of A up to, and excluding, s_t . Then $S_t = S_{old} \cup A_{(t)}$ is the set of all elements added to the solution before s_t . Note that, given \mathcal{F}_t , the sets X , S_{old} and $A_{(t)}$ are deterministic, while the rest of A is random. Recall that s_t is drawn uniformly at random from $X_{(t)} = \{x \in X \setminus A_{(t)} | c(S_t) + c(x) \leq B\}$. We need to show that $\mathbb{E}[f(s_t | S_t) | \mathcal{F}_t] \geq (1 - \varepsilon)^2\tau\mathbb{E}[c(s_t) | \mathcal{F}_t]$, where the randomness is with respect to the uniform sampling in $X_{(t)}$.

If s_t is the first element in A , this holds since all the elements in X exhibit a marginal density greater than τ . If s_t is not the first element, it means that the value and cost condition were not triggered for the previous one. Call G and E the sets of the good and negative elements with respect to S_t , i.e., $G = \{x \in X_{(t)} : f(x | S_t) \geq \tau c(x)\}$ and $E = \{x \in X_{(t)} : f(x | S_t) < 0\}$, which are also deterministically defined by \mathcal{F}_t . Finally, let p_x be $\mathbb{P}(s_t = x | \mathcal{F}_t)$ which is equal to $|X_{(t)}|^{-1}$ for all $x \in X_{(t)}$ and zero otherwise, then

$$\begin{aligned}
& \mathbb{E}[f(s_t | S_t) | \mathcal{F}_t] - (1 - \varepsilon)^2\tau\mathbb{E}[c(s_t) | \mathcal{F}_t] = \\
&= \sum_{x \in X} p_x f(x | S_t) - (1 - \varepsilon)^2\tau \sum_{x \in X} p_x c(x) \\
&\geq \sum_{x \in G \cup E} p_x f(x | S_t) - (1 - \varepsilon)^2\tau \sum_{x \in X} p_x c(x) \\
&= \varepsilon \sum_{x \in G} p_x f(x | S_t) - \sum_{x \in E} p_x |f(x | S_t)| \tag{1} \\
&\quad + (1 - \varepsilon)\tau \left[\sum_{x \in G} p_x c(x) - (1 - \varepsilon) \sum_{x \in X} p_x c(x) \right] \tag{2} \\
&\quad + (1 - \varepsilon) \sum_{x \in G} p_x \left[f(x | S_t) - \tau c(x) \right] \geq 0. \tag{3}
\end{aligned}$$

Expressions (1) and (2) are nonnegative since the value and cost conditions were not triggered before adding s_t . Expression (3) is nonnegative by the definition of G .

We have shown for all t and \mathcal{F}_t , the expected marginal density of the t -th element (if any) added by our algorithm is large enough. Next we carefully apply conditional expectations to get the desired bound. We have already argued how the algorithm induces an ordering on the elements it adds to the solution, so that they can be pictured as being added one after the other. To avoid complication in the analysis we suppose that after the algorithm stops it keeps on adding dummy elements of no cost and no value, so that in total it runs for n time steps. Consider the filtration $\{\mathcal{F}_t\}_{t=1}^n$ generated by the stochastic process associated to the algorithm, where \mathcal{F}_t narrates what happens up to the point when element s_t is considered. So that \mathcal{F}_1 is empty and \mathcal{F}_n contains all the story of the algorithm except for the last—possibly dummy—added element. From the above analysis we know that for each time $t \in \{1, \dots, n\}$ and any possible story \mathcal{F}_t of the algorithm it holds

$$\mathbb{E}[f(s_t | S_t) | \mathcal{F}_t] \geq (1 - \varepsilon)^2\tau\mathbb{E}[c(s_t) | \mathcal{F}_t]. \tag{4}$$

Note that this claim holds also if one considers the dummy elements after the actual termination of

the algorithm.

$$\begin{aligned}
\mathbb{E}[f(S)] &= \mathbb{E}\left[\sum_{t=1}^n f(s_t | S_t)\right] = \sum_{t=1}^n \mathbb{E}[f(s_t | S_t)] = \sum_{t=1}^n \mathbb{E}[\mathbb{E}[f(s_t | S_t) | \mathcal{F}_t]] \\
&= \mathbb{E}\left[\sum_{t=1}^n \mathbb{E}[f(s_t | S_t) | \mathcal{F}_t]\right] \geq (1 - \varepsilon)^2 \tau \mathbb{E}\left[\sum_{t=1}^n \mathbb{E}[c(s_t) | \mathcal{F}_t]\right] \\
&= (1 - \varepsilon)^2 \tau \mathbb{E}\left[\sum_{t=1}^n c(s_t)\right] = (1 - \varepsilon)^2 \tau \mathbb{E}[c(S)].
\end{aligned}$$

The second and fourth equalities hold by linearity of expectation, the third and fifth equalities hold by the law of total expectation. Finally, the inequality follows from monotonicity of the conditional expectation and inequality (4). \square

Having established that S has expected density comparable to our threshold τ , we move on to showing that when THRESHSEQ terminates, either a large portion of the budget is used up in expectation, or we can bound the value of good candidates that are left outside the solution.

Lemma 5. *When THRESHSEQ terminates we have $f(S) \geq \varepsilon \ell \sum_{x \in G} f(x | S)$, where $G = \{x \in X \setminus S : f(x | S) \geq \tau c(x), c(x) + c(S) \leq B\}$.*

Proof. THRESHSEQ terminates in one of two cases. Either X is empty, meaning that there are no elements still fitting in the budget whose marginal density is greater than τ —and in that case the inequality we want to prove trivially holds—or the value condition has been triggered ℓ times.

For the latter, suppose that the value condition was triggered for the i th time during iteration t_i of the while loop. Denote by S_{t_i} the solution at the end of that iteration. We are interested in the sets $X_{j^*}, G_{j^*}, E_{j^*}$ of that particular iteration of the while loop. In order to be consistent across iterations, we use $X_{(i)}, G_{(i)}$, and $E_{(i)}$ to denote these sets for iteration t_i . Since the value condition was triggered during t_i , we have $\varepsilon \sum_{x \in G_{(i)}} f(x | S_{t_i}) \leq \sum_{x \in E_{(i)}} |f(x | S_{t_i})|$. Clearly, $G_{(\ell)}$ is what we denoted by G in the statement and S_{t_ℓ} is S . Also notice that $E_{(j)} \cap E_{(k)} = \emptyset$ for $j \neq k$.

Now, by non-negativity of f and Lemma 1, we have

$$0 \leq f\left(S_{t_\ell} \cup \bigcup_{i=1}^{\ell} E_{(i)}\right) \leq f(S_{t_\ell}) + \sum_{i=1}^{\ell} \sum_{x \in E_{(i)}} f(x | S_{t_i}).$$

Rearranging the terms and using the value condition, we get

$$f(S_{t_\ell}) \geq \sum_{i=1}^{\ell} \sum_{x \in E_{(i)}} |f(x | S_{t_i})| \geq \sum_{i=1}^{\ell} \varepsilon \sum_{x \in G_{(i)}} f(x | S_{t_i}) \geq \ell \varepsilon \sum_{x \in G_{(\ell)}} f(x | S_{t_\ell}).$$

The last inequality follows from the submodularity of f and the fact that that $G_{(1)} \supseteq G_{(2)} \supseteq \dots \supseteq G_{(\ell)}$. \square

Lemma 5 still leaves a gap: how can we account for the elements which have marginal density greater than τ but are not considered due to the budget constraint? It can be the case that due to some poor random choices we initially filled the solution with low quality elements, preventing the algorithm at later stages to consider good elements with large costs. To handle this, we need the following simple lemma.

Lemma 6. *Suppose that $\mathbb{E}[c(S)] < \frac{B}{2}(1 - \varepsilon)$. Then, running THRESHSEQ $\frac{1}{\varepsilon} \log(\frac{1}{\varepsilon})$ times, with probability at least $(1 - \varepsilon)$, there is at least one run where $c(S) < \frac{B}{2}$.*

Proof. Let \mathcal{E} be the event $\{c(S) \geq \frac{B}{2}\}$, then

$$(1 - \varepsilon) \frac{B}{2} > \mathbb{E}[c(S)] \geq \mathbb{E}[c(S) | \mathcal{E}] \mathbb{P}(\mathcal{E}) \geq \mathbb{P}(\mathcal{E}) \frac{B}{2}.$$

Hence $\mathbb{P}(\mathcal{E}^C) > \varepsilon > 0$, so repeating the experiment $\frac{1}{\varepsilon} \log(\frac{1}{\varepsilon})$ independent times is enough to have a probability at least $1 - \varepsilon$ of observing at least once \mathcal{E}^C . \square

We can now present the full *parallel* algorithm PARKNAPSACK for the non-monotone case. It considers separately the set \mathcal{N}_- of “small” elements, each with cost smaller than B/n , and the set of “large” elements $\mathcal{N}_+ = \mathcal{N} \setminus \mathcal{N}_-$. The set \mathcal{N}_- is fed to the low adaptive complexity unconstrained maximization routine SUBMODMAX as discussed in Section 2. For the large elements, PARKNAPSACK samples each element of \mathcal{N}_+ with probability p to get a random subset H , and then it runs THRESHSEQ a logarithmic number of times on H , in parallel, for different guesses of the “right” threshold. The partition between \mathcal{N}_+ and \mathcal{N}_- is critical in bounding the adaptivity of THRESHSEQ, as $\kappa(\mathcal{N}_+) \leq n$.

Algorithm 3: PARKNAPSACK($\mathcal{N}, f, \varepsilon, \alpha, p, B$)

- 1: **Input:** Ground set \mathcal{N} , submodular function f , budget B , precision ε , parameter α and sampling probability p
 - 2: $\mathcal{N}_- \leftarrow \{x \in \mathcal{N} : c(x) < \frac{B}{n}\}$; $\mathcal{N}_+ \leftarrow \mathcal{N} \setminus \mathcal{N}_-$
 - 3: $x^* \leftarrow \max_{x \in \mathcal{N}_+} f(x)$; $\hat{\tau} \leftarrow \alpha n \frac{f(x^*)}{B}$
 - 4: $\hat{\varepsilon} \leftarrow \varepsilon/125$; $\ell \leftarrow \hat{\varepsilon}^{-2}$; $k \leftarrow \hat{\varepsilon}^{-1} \log(n)$
 - 5: $S_- \leftarrow \text{SUBMODMAX}(\mathcal{N}_-, \hat{\varepsilon})$
 - 6: $H \leftarrow$ sample each element in \mathcal{N}_+ independently at random with probability p
 - 7: **for** $i = 0, 1, \dots, k$ **in parallel do**
 - 8: $\tau_i \leftarrow \hat{\tau} \cdot (1 - \hat{\varepsilon})^i$
 - 9: **for** $j = 1, 2, \dots, \hat{\varepsilon}^{-1} \log(\hat{\varepsilon}^{-1})$ **in parallel do**
 - 10: $S_j^i \leftarrow \text{THRESHSEQ}(H, \tau_i, \hat{\varepsilon}, \ell, B)$
 - 11: **return** $T \in \arg \max_{i,j} \{f(S_j^i), f(x^*), f(S_-)\}$
-

Theorem 1. *For $\alpha = 2 - \sqrt{3}$, $p = \frac{1 - \alpha}{2}$ and $\varepsilon < \frac{1}{3}$, PARKNAPSACK is a $(9.465 + \varepsilon)$ -approximation algorithm with $O(\frac{1}{\varepsilon} \log n)$ adaptivity and $O(\frac{n^2}{\varepsilon^3} \log^2 n \log \frac{1}{\varepsilon})$ total queries.*

Proof. Excluding the call to SUBMODMAX, the claim on the adaptivity follows directly from Lemma 3 with $\ell = O(\varepsilon^{-2})$, and the observation that $\kappa(\mathcal{N}_+) \leq n$. The adaptivity is indeed only due to THRESHSEQ, since the guessing of the threshold, as well as the multiple runs of THRESHSEQ, happen independently in parallel. Relative to the query complexity, we have the bound in Lemma 3 multiplied by an extra $O(\frac{\log n}{\varepsilon^2} \log \frac{1}{\varepsilon})$ factor caused by the two *for* loops. SUBMODMAX does not affect these asymptotics since it has adaptivity bounded by $O(\frac{1}{\varepsilon})$ and linear query complexity.

Consider now the approximation guarantee. Call O^* the optimal solution, and O^+, O^- its intersections with \mathcal{N}_+ and \mathcal{N}_- respectively. We can upper bound $f(O^-)$ with the unconstrained max

on \mathcal{N}_- , since there are at most n elements in \mathcal{N}_- whose cost is at most $\frac{B}{n}$. Using the combinatorial algorithm of Chen et al. [12], we get

$$f(O^-) \leq (2 + \hat{\varepsilon}) \cdot f(S_-) \leq (2 + \hat{\varepsilon}) \cdot ALG \quad (5)$$

Let $O \in \arg \max\{f(T) : T \subseteq \mathcal{N}_+, c(T) \leq B\}$, i.e., O is an optimal solution in \mathcal{N}_+ . Clearly $f(O^+) \leq f(O)$, so we will upper bound the latter. Let $O_H = O \cap H$. By submodularity and monotonicity of $f(\cdot \cap O)$, we have $pf(O) \leq \mathbb{E}[f(O_H)]$. Outside O , the function may be non-monotone, so we need Lemma 2. In particular, we apply it on the submodular function $g(\cdot) = f(\cdot \cup O)$. Since elements belong to H with probability p , for $S \subseteq H$ we get

$$p(1-p)f(O) \leq (1-p)\mathbb{E}[f(O_H)] \leq \mathbb{E}[\mathbb{E}[f(S \cup O_H) | O_H]] = \mathbb{E}[f(S \cup O_H)]. \quad (6)$$

Let $\tau^* = \alpha f(O)/B$. By the parallel guesses we have that there exists $\tau = \tau_i$ such that $(1-\hat{\varepsilon})\tau^* \leq \tau < \tau^*$. This directly follows from the definitions of τ^* and $\hat{\tau}$ and the fact that $nf(x^*) \geq f(O) \geq f(x^*)$. We focus only on this particular τ and consider two cases, depending on $\mathbb{E}[c(S)]$, where S is the set outputted by THRESHSEQ for τ . If $\mathbb{E}[c(S)] \geq (1-\hat{\varepsilon})\frac{B}{2}$, then, from Lemma 4 we have

$$ALG \geq \mathbb{E}[f(S)] \geq (1-\hat{\varepsilon})^2 \tau \mathbb{E}[c(S)] \geq (1-\hat{\varepsilon})^3 \alpha f(O) \frac{\mathbb{E}[c(S)]}{B} \geq (1-\hat{\varepsilon})^4 \frac{\alpha}{2} f(O). \quad (7)$$

If $\mathbb{E}[c(S)] < (1-\hat{\varepsilon})\frac{B}{2}$ we need a more careful analysis, via Lemmata 5 and 6. Consider the multiple runs of THRESHSEQ corresponding to τ . Let \mathcal{G} be the event that at least one of those runs outputs S with $c(S) < \frac{B}{2}$ and consider that solution; recall that $\mathbb{P}(\mathcal{G}) \geq (1-\hat{\varepsilon})$ from Lemma 6. What we want to bound is the total value of the elements of O_H which are not in S . The ones retaining a good marginal density with respect to S can be divided into two categories, depending on the reason why they were not added to S :

$$\begin{aligned} G &= \{x \in H : f(x | S) \geq \tau c(x), c(x) + c(S) \leq B\}, \\ \tilde{G} &= \{x \in H : f(x | S) \geq \tau c(x), c(x) + c(S) > B\}. \end{aligned}$$

The total contribution of the elements in G can be bounded applying Lemma 5. For $\tilde{G} \cap O_H$ we know that it contains at most one element \tilde{x} , since we are conditioning on \mathcal{G} and thus, if such \tilde{x} exists, $c(\tilde{x}) > \frac{B}{2}$. Moreover, $f(\tilde{x}) \leq f(x^*)$. Finally, we know that the marginal density of all the other elements in $O_H \setminus S$ is at most τ . Let \mathcal{E} be the event that $\tilde{G} \cap O_H \neq \emptyset$ given \mathcal{G} and q its probability. We have

$$\begin{aligned} f(S \cup O_H) &\leq f(S) + \mathbf{1}_{\mathcal{E}} \cdot f(\tilde{x} | S) + \sum_{x \in G} f(x | S) + \sum_{x \in O_H \setminus (G \cup \tilde{G})} f(x | S) \\ &\leq f(S)(1 + \hat{\varepsilon}) + \mathbf{1}_{\mathcal{E}} \cdot (f(\tilde{x} | S) - \tau c(\tilde{x})) + \tau c(O_H) \\ &\leq f(S)(1 + \hat{\varepsilon}) + \mathbf{1}_{\mathcal{E}} \cdot (f(x^*) - \tau \frac{B}{2}) + \tau c(O_H). \end{aligned}$$

Keeping fixed H , let's apply the expectation on the randomness in THRESHSEQ, conditioning on \mathcal{G} and recalling that both $f(S)$ and $f(x^*)$ are upper bounded by ALG :

$$\mathbb{E}[f(S \cup O_H) | \mathcal{G}] \leq (1 + \hat{\varepsilon} + q)ALG + \tau c(O_H) - q\tau \frac{B}{2},$$

Now move on to the expectation with respect to H . Note that by submodularity $f(S \cup O_H) \leq 2f(O)$. We have

$$\mathbb{E}[f(S \cup O_H)] = \mathbb{E}[f(S \cup O_H) | \mathcal{G}] \mathbb{P}(\mathcal{G}) + \mathbb{E}[f(S \cup O_H) | \mathcal{G}^C] \mathbb{P}(\mathcal{G}^C)$$

$$\leq \mathbb{E}[f(S \cup O_H) | \mathcal{G}] + 2\hat{\varepsilon}f(O).$$

Putting together the last two inequalities and recalling that $\mathbb{E}[c(O_H)] = pc(O) \leq pB$, we have

$$\mathbb{E}[f(S \cup O_H)] \leq (2\hat{\varepsilon} + \alpha - q\frac{\alpha}{2})f(O) + (1 + \hat{\varepsilon} + q)ALG.$$

Combining that with (6), we finally obtain

$$f(O) \leq \frac{(1 + q + \hat{\varepsilon})}{[p(1 - p) - \alpha p + \frac{\alpha q}{2} - 2\hat{\varepsilon}]}ALG \quad (8)$$

At this point we need to optimize the constants in (5), (6), (8), also using that

$$OPT \leq f(O^+) + f(O^-) \leq f(O) + f(O^-).$$

Setting $p = \frac{1}{2}(\sqrt{3} - 1)$, $\alpha = 2 - \sqrt{3}$ and rescaling $\hat{\varepsilon} = \frac{\varepsilon}{125}$ we get, for small enough $\hat{\varepsilon}$ and for any value of $q \in (0, 1)$ the desired bound: $OPT \leq (2(3 + \sqrt{3}) + \varepsilon)ALG$. \square

3.1 Variants and Implications

As mentioned already, an interesting feature of our approach is that—with few modifications—yields a number of algorithms that match or improve the state-of-the-art. We only sketch these modifications here and we defer the details to the appendix.

We begin with a discussion on the possible trade-offs between adaptivity and query complexity. THRESHSEQ can be adapted to spare $\Theta(\frac{n}{\log n})$ value queries at the cost of $O(\log n)$ extra adaptive rounds. The idea is to use *binary search* to locate k^* in the while loop of THRESHSEQ. Only a logarithmic number of prefixes needs to be sequentially considered, instead of all of them in parallel. To be able to binary search k^* , though, a carefully modified version of the value condition is implemented, since the one used in THRESHSEQ exhibits a multi-modal behaviour.

Theorem 2. *For $\varepsilon \in (0, 1/3)$, it is possible to achieve a $(9.465 + \varepsilon)$ -approximation in $O(\frac{1}{\varepsilon} \log^2 n)$ adaptive rounds and $O(\frac{n}{\varepsilon^3} \log^3 n \log \frac{1}{\varepsilon})$ queries.*

3.1.1 Monotone Submodular Functions

For monotone objectives, the approximation ratio of PARKNAPSACK can be significantly improved. In particular, in THRESHSEQ, we do not need to address the value condition any more. Moreover, the small elements can be accounted for without any extra loss in the approximation. As in the case of Theorems 1 and 2, it is possible to trade a logarithmic loss in adaptivity for an almost linear gain in query complexity.

Theorem 3. *For $\varepsilon \in (0, 1)$ it is possible to achieve a $3 + \varepsilon$ approximation in $O(\frac{1}{\varepsilon} \log n)$ adaptive rounds and $O(\frac{n^2}{\varepsilon^3} \log^2 n \log \frac{1}{\varepsilon})$ queries or in $O(\frac{1}{\varepsilon} \log^2 n)$ adaptive rounds and $O(\frac{n}{\varepsilon^3} \log^3 n \log \frac{1}{\varepsilon})$ queries.*

Note that the variant using $\tilde{O}(n)$ queries is the first $O(1)$ -approximation algorithm for the problem combining this few queries with sublinear adaptivity.

3.1.2 Cardinality constraints

PARKNAPSACK can be directly applied to cardinality constraints for (possibly) non-monotone objectives. Again with some simple modifications, it is possible to achieve a much better approximation.

Theorem 4. *For $\varepsilon \in (0, 2/5)$ it is possible to achieve a $5.83 + \varepsilon$ approximation, in $O(\frac{1}{\varepsilon} \log n)$ adaptive rounds and $O(\frac{nk}{\varepsilon^3} \log n \log k \log \frac{1}{\varepsilon})$ queries, or in $O(\frac{1}{\varepsilon} \log n \log k)$ adaptive rounds and $O(\frac{n}{\varepsilon^3} \log n \log^2 k \log(\frac{1}{\varepsilon}))$ queries.*


Although we do not heavily adjust our algorithms to cardinality constraints, [Theorem 4](#) is directly comparable to the very recent results of Ene and Nguyen [17] and Kuhnle [31] which are tailored for the problem.

4 Experiments


We evaluate the performance of PARKNAPSACK on real datasets and real-world applications, as is often the case in the related literature [35, 21, 1, 6, 31]. All three objectives we use are non-monotone submodular. We compare against the state-of-the-art of *fast* algorithms for non-monotone submodular maximization subject to a knapsack constraint, in order to demonstrate that PARKNAPSACK produces almost equally good solutions with an exponential improvement on the adaptivity. We provide two kinds of figures: objective versus budget (or instance size) and objective versus adaptive steps, for a given instance, as in Balkanski et al. [3] and Fahrback et al. [21]. Note that we use the version of PARKNAPSACK from [Theorem 2](#) to ensure $\tilde{O}(n)$ query complexity. The benchmarks we use are plain GREEDY, FANTOM of Mirzasoleiman et al. [35] and SAMPLEGREEDY of Amanatidis et al. [1]. The last two have the state-of-the-art performance in terms of objective value for knapsack constraints among algorithms with practical running times, i.e., among algorithms with subquadratic query complexity. On the other hand, these algorithms are not designed for low adaptivity but, the only alternative, i.e. continuous methods, are impractical for the instance sizes we consider. The GREEDY algorithm builds a solution step by step by adding the element with the highest marginal value, until the budget is exhausted. While this naive approach has no theoretical guarantees, it is very fast and often has acceptable performance in practice. FANTOM builds on GREEDY and is robust for intersecting p -systems and knapsack constraints providing a $10(1 + \varepsilon)$ -approximation for our setting. Finally, SAMPLEGREEDY greedily selects elements according to their marginal value per cost ratio, but only adds them to the solution with some probability. This leads to a 5.83-approximation. These algorithms need $O(n \log n)$ queries and adaptive steps, when implemented using with *lazy evaluations* [33].

Apart from GREEDY, all other algorithms need some constant parameters as part of the input, in addition to the submodular instance. For SAMPLEGREEDY and PARKNAPSACK, we set $p = 0.9$, without further optimizing it for each problem. Although this is not the theoretical optimal, the instances in question come from real-world datasets, which typically are not arbitrarily non-monotone. As a result the algorithm can afford to discard elements at a lower rate, as it is less likely that these will hurt the objective in the long run. Moreover, we set $\varepsilon = 1/8$ for FANTOM and $\varepsilon = 1/8, \alpha = 2 - \sqrt{3}$ for PARKNAPSACK balancing performance and running time. Each experiment on the top row was repeated 3 times, to gain an estimate of the variance. In total, the whole array of tests ran on four `t2.micro` instances on the Amazon Elastic Compute Cloud (EC2), which is part of Amazon Web Services (AWS).


Movie Recommendation. Given a set of movies A , a list of genres C_i such that $C_1 \cup C_2 \cup \dots \cup C_k = A$ and a list of user generated keyword tags t_{iu} and ratings r_{iu} , where $i \in A$ and u is the id of a




(a) Movie Recommendation




(b) Movie Recommendation




(c) Revenue Maximization



(d) Revenue Maximization



(e) Graph Cut on $G(n, 0.1)$



(f) Graph Cut on $G(n, 0.1)$

Figure 1: Each row contains two plots, corresponding to the same submodular problem. The left column contains the objective function value for different budget or instance sizes. The right column focuses on one vertical slice of the top one: for a specific budget and instance size, the objective value is presented as a function of the number of adaptive rounds, indicating that if we required to stop after a small number of rounds all other algorithms would perform extremely poorly. In all cases, the results are consistent: PARKNAPSACK has comparable performance, with drastically improved adaptivity.

user, a movie recommendation system aims to use this information to provide a short list of diverse options that match certain preferences. The [MovieLens dataset](#) [26] provides a very large set of movies that include user generated tags and ratings. We calculate the similarity between two movies (following the procedure of Amanatidis et al. [1]; see below) and produce a weighted complete graph, where each vertex is a movie. For $i, j \in A$ the weight w_{ij} represents their similarity. In addition,

we use χ_{ij} to indicate if the two movies share a genre. Putting everything together, the objective function is: $v(S) = \alpha \sum_{i \in S} r_i + \beta (\sum_{i \in S} \sum_{j \in A} w_{ij} - \sum_{i \in S} \sum_{j \in S} (\lambda + \chi_{ij} \mu) w_{ij})$ for $\lambda, \mu, \alpha, \beta \geq 0$ where r_i represents the average rating of movie i . This is a weighted average of the ratings of the movies in S and a modified *maximal marginal relevance* [8]. The second part is similar to a max cut (in fact it is a max cut for $\lambda = 1$ and $\mu = 0$), but allows the internal edges to be penalized differently, depending on whether the movies are similar or belong to the same genre. For the experiments we consider a subset of 5000 movies and set $\alpha = \beta = 0.5, \lambda = 3$ and $\mu = 7$. Each movie is assigned a cost sampled uniformly from $[0, 1]$ and the total budget ranges from 0.01 to 0.1 of the total cost.

We outline how the weights w_{ij} are created, given the MovieLens dataset. Each movie i is associated with a tag vector $t^i \in [0, 1]^{1128}$, which encodes how much each tag applies to it. The tags are user generated. For example, a movie like “Titanic” could have a score of 0.9 for “ships”, 0.8 for “romance” and 0 for “talking animals”. These tag vectors are not normalized. Given those, we compute the similarities as:

$$w_{ij} = \sqrt{\sum_{k=1}^{1128} (\min\{t_k^i, t_k^j\})^2}. \quad (9)$$

This approach ensures that movies with tag vectors that are close appear more similar. There are various trade-offs in the selection of the similarity metric. For instance, if one defined $w_{ij} = t^i \cdot t^j$, then a movie with all tags set to 1 would appear more similar to any other movie, even when comparing a movie with itself! Going to the other extreme, if $w_{ij} = t^i \cdot t^j / (|t^i| |t^j|)$ this issue would be avoided, but some information would be lost as every movie would have a normalised tag vector, even though having a high score on one tag should not impact the score on other tags. Ultimately, using (9) appears to be a reasonable compromise between the two. We stress that our experimental findings about the performance of each algorithm remain qualitatively the same for any sensible choice of w_{ij} .

Revenue Maximization. Representing a social network as a weighted graph, where each edge signifies how much one user is influenced by another, our goal is to select a subset S of users who are given a product to advertise, in order to maximize the revenue from sales. We use the [YouTube Network](#) [40], and consider the subgraph induced by selecting its Top 5000 communities, which has 39841 vertices and 224235 edges. We assign edge weights w_{ij} sampled uniformly in $[0, 1]$ and each user $i \in V$ has a suggestibility parameter α_i drawn from a Pareto Type II distribution with $\lambda = 1, \alpha = 2$. The objective to maximize is: $v(S) = \sum_{i \in V \setminus S} \alpha_i \sqrt{\sum_{j \in S} w_{ij}}$. Each user is assigned a cost proportional to their incident edges, with the budget ranging from 0.01 to 0.1 of the total cost.

Maximum Weighted Cut. Given an Erdős–Rényi graph $G(n, p)$ where n is the number of vertices and p the probability of including each edge, the objective is to find a *cut of maximum weight*. Fixing $p = 0.1$, we let $n \in \{30, \dots, 5000\}$ (with an exponential step) and assign random edge weights and costs sampled uniformly from $[0, 1]$, while the budget is fixed at 15% of the total cost.

5 Conclusions

In this paper we close the gap for the adaptive complexity of non-monotone submodular maximization subject to a knapsack constraint, up to a $O(\log \log n)$ factor. Our algorithm, PARKNAPSACK, is combinatorial and can be modified to achieve trade-offs between adaptivity and query complexity. In particular, it may use nearly linear queries, while achieving an exponential improvement on adaptivity compared to existing algorithms with subquadratic query complexity.

A Adapting PARKNAPSACK to Use Binary Search

This appendix addresses the adaptivity versus queries trade-off mentioned in [Section 3.1](#) which leads to [Theorem 2](#) and, implicitly, to [Theorems 3](#) and [4](#). The algorithms and the proofs of these two theorems are presented in [Appendices B](#) and [C](#), respectively.

As already pointed out in the main text, the value condition in THRESHSEQ may exhibit a multi-modal behaviour along a single iteration of the while loop. In order to enable binary search for k^* , we want to tweak the value condition so that if it is triggered for a certain prefix A_i , it remains activated for all A_j for $j \geq i$ in the specific while loop iteration.

So, fix an arbitrary iteration of the while loop. Call S the initial solution and A the sequence drawn by SAMPLESEQ, $\{X_i\}_i$ the sequence of the sets of elements still fitting into the budget relative to each prefix A_i and with G_i and E_i the subsets of X_i containing the *good*, i.e., marginal density greater than τ , and *bad*, i.e., negative marginal density, elements, respectively. First, note that the cost condition is clearly unimodal: the $\{X_i\}_i$ is a decreasing sequence of sets and hence $c(X_i)$ is a non-increasing sequence of costs, while $c(X)$ stays fixed: as soon as the cost of X_i drops below $(1 - \varepsilon)c(X)$ it stays there for all the prefixes longer than A_i .

For the value condition we need a bit more work; if for some j it holds that $\varepsilon \sum_{x \in G_j} f(x | S \cup A_j) \leq \sum_{x \in E_j} |f(x | S \cup A_j)|$, it may be the case that the inequality switches direction later in the same iteration of the while loop. Notice that it can happen for one of two reasons: either elements with negative marginals are added to the solution or they are thrown away due to budget constraint. We want a modification which is robust to these corner cases. To this end, we add to the value condition the absolute contribution of two sets of items.

First, we redefine the set E_i to contain also all the bad elements considered in that while loop, regardless of the budget condition, i.e., $E_i \leftarrow \{a \in X : f(a | S \cup A_i) < 0\}$. Moreover for each prefix A_i , we define \mathcal{E}_i as the set of all the items in the prefix A_i which added negative marginal when inserted in the solution. i.e., $\mathcal{E}_i = \{a_t \in A_i : f(a_t | S \cup A_{t-1}) < 0\}$.

The new value condition then reads:

$$\varepsilon \sum_{x \in G_i} f(x | S \cup A_i) \leq \sum_{x \in E_i} |f(x | S \cup A_i)| + \sum_{a_j \in \mathcal{E}_i} |f(a_j | S \cup A_j)|.$$

Notice that now everything works out just fine: the left hand side of the condition is monotonically decreasing in i , while the right hand side is monotonically increasing, by submodularity and the fact that now $\{E_t \cup \mathcal{E}_t\}_t$ is an increasing set sequence.

Given the new algorithm, THRESHBIN, we need to show that it retains the right properties of THRESHSEQ and argue about its adaptive and query complexity.

Lemma 7. *Consider a run of THRESHBIN and denote with S and \bar{S} the preliminary and final solution as in the algorithm. Then the following properties hold:*

- $c(\bar{S}) \leq c(S) \leq B$
- $f(\bar{S}) \geq f(S)$
- $\mathbb{E}[f(S)] \geq \tau(1 - \varepsilon)^2 \mathbb{E}[c(S)]$
- *Call G the set of elements still fitting in the budget after S , whose marginal density with respect to S is greater than τ . Then $f(\bar{S}) \geq \varepsilon \ell \sum_{x \in G} f(x | S)$.*

THRESHBIN needs $O\left(\log n \left(\frac{\log n \kappa(X)}{\varepsilon} + \ell\right)\right)$ adaptive rounds and $O\left(n \log n \left(\frac{\log n \kappa(X)}{\varepsilon} + \ell\right)\right)$ value queries.

Algorithm 4: THRESHBIN($X, \tau, \varepsilon, \ell, B$)Variant of THRESHSEQ that utilises binary search

```
1: Input: set  $X$  of elements, threshold  $\tau > 0$ , precision  $\varepsilon \in (0, 1)$ , parameter  $\ell$  and budget  $B$ 
2:  $S \leftarrow \emptyset$ ; ctr  $\leftarrow 0$ ; flag  $\leftarrow 0$ 
3:  $X \leftarrow \{x \in X : f(x) \geq \tau c(x)\}$ 
4: while  $X \neq \emptyset$  and ctr  $< \ell$  do
5:    $[a_1, a_2, \dots, a_d] \leftarrow \text{SAMPLESEQ}(S, X, B)$ ;
6:    $b_l \leftarrow 1, b_r \leftarrow d$ ;
7:   while  $b_l < b_r$  do
8:      $i = \lfloor (b_r + b_l)/2 \rfloor$ ;
9:      $A_i \leftarrow \{a_1, a_2, \dots, a_i\}$ 
10:     $X_i \leftarrow \{a \in X \setminus A_i : c(a) + c(S \cup A_i) \leq B\}$ ;
11:     $G_i \leftarrow \{a \in X_i : f(a | S \cup A_i) \geq \tau \cdot c(a)\}$ 
12:     $E_i \leftarrow \{a \in X : f(a | S \cup A_i) < 0\}$ 
13:     $\mathcal{E}_i \leftarrow \{a_s \in A_i : f(a_s | S \cup A_s) < 0\}$ 
14:     $c_1 \leftarrow c(G_i) \leq (1 - \varepsilon)c(X)$ ;
15:     $c_2 \leftarrow \sum_{x \in G_i} \varepsilon f(x | S \cup A_i) \leq \sum_{x \in E_i} |f(x | S \cup A_i)| + \sum_{a_j \in \mathcal{E}_i} |f(a_j | S \cup A_j)|$ ;
16:    if  $c_1$  or  $c_2$  then
17:       $b_r \leftarrow i$ 
18:    else
19:       $b_l \leftarrow i + 1$ 
20:    if  $c_2$  and not  $c_1$  then
21:      flag  $\leftarrow 1$ ;
22:    else
23:      flag  $\leftarrow 0$ ;
24:     $k^* = b_r$ ;
25:     $S \leftarrow S \cup A_{k^*}$ ;
26:     $X \leftarrow G_{k^*}$ 
27:    ctr  $\leftarrow$  ctr + flag
28: Suppose  $S = \{s_1, s_2, \dots, s_{|S|}\}$ , where the indices imply the total ordering from the proof of
    Lemma 4
29:  $\bar{S} \leftarrow \emptyset$ 
30: for  $t = 1, \dots, |S|$  do
31:   if  $f(s_t | \{s_1, \dots, s_{t-1}\}) > 0$  then
32:      $\bar{S} \leftarrow \bar{S} \cup \{s_t\}$ 
33: return  $\bar{S}$ 
```

Proof. The proof of this Lemma is quite similar to the one for THRESHSEQ, so we just highlight the differences.

First, the new value condition is stricter than the old one, so the $\mathbb{E}[f(s_t | S_t) | \mathcal{F}_t] \geq \tau(1 - \varepsilon)^2 \mathbb{E}[c(s_t) | \mathcal{F}_t]$ inequality holds as well, where S_t and \mathcal{F}_t are as in the proof of Lemma 4. This implies that

$$\mathbb{E}[f(S)] \geq \tau(1 - \varepsilon)^2 \mathbb{E}[c(S)] .$$

The bounds on adaptivity and query complexity follow easily from the binary search and the analysis of Lemma 3. Further, the first and second bullets follow directly from $\bar{S} \subseteq S$ and the fact that we only filter out from S elements with negative contribution.

Consider now the last remaining statement to prove (the analog of Lemma 5). For all real numbers a we denote with $a_+ = \max\{a, 0\}$ its positive part and with $a_- = \max\{-a, 0\}$ the negative one. Clearly $a = a_+ - a_-$.

$$f(S) = \sum_{t=1}^T (f(s_t | S_t)_+ - f(s_t | S_t)_-) \leq \sum_{t=1}^T f(s_t | S_t)_+ \leq \sum_{s_t \in \bar{S}} f(s_t | \bar{S}) = f(\bar{S}) ,$$

where in the last inequality we used submodularity.

Similarly to the proof for THRESHSEQ, consider $t_1, \dots, t_\ell, E_{(1)}, \dots, E_{(\ell)}, G_{(1)}, \dots, G_{(\ell)},$ and $\mathcal{E}_{(1)}, \dots, \mathcal{E}_{(\ell)}$. Notice that they are all disjoint. Like before, for $s_i \in S$, S_i denotes $\{s_1, \dots, s_{i-1}\}$, but we slightly abuse the notation and have S_{t_j} denote the set S at the end of the iteration of the outer while loop where ctr is increased for the ℓ th time. We have

$$\begin{aligned} 0 &\leq f(S_{t_\ell} \cup \bigcup_{j=1}^{\ell} E_{(j)}) \leq f(S_{t_\ell}) + f\left(\bigcup_{j=1}^{\ell} E_{(j)} \mid S_{t_\ell}\right) \\ &\leq \sum_{s_i \in S_{t_\ell}} (f(s_i | S_i)_+ - f(s_i | S_i)_-) + \sum_{j=1}^{\ell} f(E_{(j)} \mid S_{t_j}) \\ &\leq \sum_{s_i \in S_{t_\ell}} (f(s_i | S_i)_+ - f(s_i | S_i)_-) + \sum_{j=1}^{\ell} \sum_{x \in E_{(j)}} f(x \mid S_{t_j}) . \end{aligned}$$

Rearranging terms, and using the value condition, we get

$$\begin{aligned} f(\bar{S}) &\geq \sum_{s_i \in S_{t_\ell}} f(s_i | S_i)_+ \geq \sum_{s_i \in S_{t_\ell}} f(s_i | S_i)_- + \sum_{j=1}^{\ell} \sum_{x \in E_{(j)}} |f(x \mid S_{t_j})| \\ &\geq \sum_{j=1}^{\ell} \left[\sum_{x \in E_{(j)}} |f(x \mid S_{t_j})| + \sum_{s_i \in \mathcal{E}_{(j)}} |f(s_i \mid S_i)| \right] \geq \\ &\geq \varepsilon \sum_{j=1}^{\ell} \left[\sum_{x \in G_{(j)}} f(x \mid S_{t_j}) \right] \geq \ell \varepsilon \sum_{x \in G_{(\ell)}} f(x \mid S_{t_\ell}) . \end{aligned}$$

Observing that $S_{t_\ell} = S$ concludes the proof. \square

Theorem 5. For $\varepsilon \in (0, 1/3)$, it is possible to achieve a $(9.465 + \varepsilon)$ -approximation in $O(\frac{1}{\varepsilon} \log^2 n)$ adaptive rounds and $O(\frac{n}{\varepsilon^3} \log^3 n \log \frac{1}{\varepsilon})$ queries.

Proof. A large part of this proof is very similar to the proof of Theorem 1. Hence, we only highlight the differences, while retaining the same notation. Note that now S is no more the output of the

algorithm, but the non-filtered output of THRESHBIN (the filtered version being \bar{S}). The two cases in the analysis are similar.

If $\mathbb{E}[c(S)] \geq (1 - \varepsilon)\frac{B}{2}$, then we have a result analogous to (7) in the main text:

$$ALG \geq \mathbb{E}[f(\bar{S})] \geq \mathbb{E}[f(S)] \geq \tau(1 - \varepsilon)^2 \mathbb{E}[c(S)] \geq \frac{1}{2}\alpha(1 - \varepsilon)^4 f(O).$$

Otherwise, we can repeat the algorithm $\frac{1}{\varepsilon} \log(\frac{1}{\varepsilon})$ times to be sure that, with probability at least $1 - \varepsilon$, we observe $\frac{B}{2} > c(S) > c(\bar{S})$. From this we can infer that at most one element is contained in $\tilde{G} \cap O_H$. Notice however a difference here: G and \tilde{G} are the elements with good marginal with respect to S , not with respect to \bar{S} .

$$\begin{aligned} f(S \cup O_H) &\leq f(S) + \mathbf{1}_{\mathcal{E}} \cdot f(\tilde{x} | S) + \sum_{x \in G} f(x | S) + \sum_{x \in O_H \setminus (G \cup \tilde{G})} f(x | S) \\ &\leq f(\bar{S})(1 + \hat{\varepsilon}) + \mathbf{1}_{\mathcal{E}} \cdot (f(\tilde{x} | S) - \tau c(\tilde{x})) + \tau c(O_H) \\ &\leq f(\bar{S})(1 + \hat{\varepsilon}) + \mathbf{1}_{\mathcal{E}} \cdot (f(x^*) - \tau \frac{B}{2}) + \tau c(O_H), \end{aligned}$$

where \mathcal{E} is the event that $\tilde{G} \cap O_H$ is not empty given that $c(S) < \frac{B}{2}$. Proceeding as in the proof of [Theorem 1](#) and noting that $f(S) \leq f(\bar{S}) \leq ALG$ we arrive at the same inequality as in (8) in the main text:

$$f(O) \leq \frac{(1 + q + \hat{\varepsilon})}{[p(1 - p) - \alpha p + \frac{\alpha q}{2} - 2\hat{\varepsilon}]} ALG.$$

The rest of the proof is essentially the same with the proof of [Theorem 1](#). \square

B Monotone Objectives and a Knapsack Constraint

For monotone objectives we can improve the approximation factor by slightly modifying the main algorithm. Notice, moreover, that in THRESHSEQ the only relevant condition is the cost condition since no element can have a negative marginal value.

Lemma 8. *For any set X , threshold τ , precision $\varepsilon \in (0, 1)$, parameter ℓ and budget B , the random set S output by THRESHSEQ is such that*

$$\mathbb{E}[f(S)] \geq \tau(1 - \varepsilon)\mathbb{E}[c(S)].$$

The random set S is always a feasible solution and if $c(S) < B$, then all the elements in $X \setminus S$ have either marginal density with respect to S smaller than τ or there is no room for them in the budget. Finally, the adaptivity is upperbounded by $\frac{1}{\varepsilon} \log(n\kappa(X))$, while the query complexity by $\frac{n^2}{\varepsilon} \log(n\kappa(X))$.

Proof. Again the proof is similar to the one for the non-monotone case in the main text. There are three main differences. First, the adaptive complexity is given only by the number of times the cost condition is triggered, hence an upper bound is given by $\frac{1}{\varepsilon} \log(n\kappa(X))$. The query complexity is simply obtained multiplying that by a n^2 factor as in the proof of [Theorem 1](#).

Second, the algorithm can now only stop if the budget is exhausted or there are no good elements fitting within the budget; the while loop terminates only in those two cases. Finally, the main chain of inequalities is simply

$$\mathbb{E}[f(s_t | S_t) | \mathcal{F}_t] = \sum_{x \in X} p_x f(x | S_t) \geq \tau \sum_{x \in G} p_x c(x) \geq \tau(1 - \varepsilon) \sum_{x \in X} p_x c(x) = (1 - \varepsilon)\tau \mathbb{E}[c(s_t) | \mathcal{F}_t],$$

where in the second inequality we used the fact that the cost condition is not triggered. Taking the expectation on the whole process and reasoning as in the proof of [Lemma 4](#), we conclude the proof. \square

We are ready to present the full algorithm for the monotone case. There are two main differences to the non-monotone case. First, there is no need to sample a subset H and use the Sampling Lemma, since $f(S) \leq f(S \cup O)$. Second, if one defines the small elements to be the ones with cost smaller than $\varepsilon \frac{B}{n}$ it is possible to account for them by simply adding all of them to the solution at the cost of filling an ε fraction of the budget. Notice that this can be done while keeping $\kappa(\mathcal{N}_+)$ linear in n . The remaining $(1 - \varepsilon)$ fraction of the budget is then filled via THRESHSEQ on the large elements. The pseudocode is given in [Algorithm 5](#) below.

Algorithm 5: PARKNAPSACKMONOTONE($\mathcal{N}, f, \varepsilon, \alpha, B$)

Full algorithm for monotone objectives and a knapsack constraint

- 1: **Input:** Ground set \mathcal{N} , monotone submodular function f , budget B , precision $\varepsilon \in (0, 1)$ and parameter $\alpha \in (0, 1)$
 - 2: $\mathcal{N}_- \leftarrow \{x \in \mathcal{N} : c(x) < \varepsilon \frac{B}{n}\}$
 - 3: $\mathcal{N}_+ \leftarrow \mathcal{N} \setminus \mathcal{N}_-$
 - 4: $x^* \leftarrow \max_{x \in \mathcal{N}} f(x)$, $\hat{\tau} \leftarrow \alpha n \frac{f(x^*)}{B}$
 - 5: $\hat{\varepsilon} \leftarrow \frac{1}{10}\varepsilon$, $k \leftarrow \frac{1}{\hat{\varepsilon}} \log(n)$
 - 6: **for** $i = 0, \dots, k$ **in parallel do**
 - 7: $\tau_i \leftarrow \hat{\tau} \cdot (1 - \hat{\varepsilon})^i$
 - 8: **for** $j = 1, \dots, \frac{1}{\hat{\varepsilon}} \log(\frac{1}{\hat{\varepsilon}})$ **in parallel do**
 - 9: $S_j^i \leftarrow \text{THRESHSEQ}(\mathcal{N}_+, \tau_i, \hat{\varepsilon}, (1 - \hat{\varepsilon})B)$
 - 10: $T_j^i \leftarrow S_j^i \cup \mathcal{N}_-$
 - 11: $T \leftarrow \arg \max\{f(T_i^j), f(x^*)\}$
 - 12: **Return** T
-

Theorem 6. For $\varepsilon \in (0, 1)$ it is possible to achieve a $3 + \varepsilon$ approximation in $O(\frac{1}{\varepsilon} \log n)$ adaptive rounds and $O(\frac{n^2}{\varepsilon^3} \log^2 n \log \frac{1}{\varepsilon})$ queries or in $O(\frac{1}{\varepsilon} \log^2 n)$ adaptive rounds and $O(\frac{n}{\varepsilon^3} \log^3 n \log \frac{1}{\varepsilon})$ queries.

Proof. We show that PARKNAPSACKMONOTONE with parameters $\alpha = \frac{2}{3}$ and any $\varepsilon \in (0, 1)$ satisfies the statement of the theorem. We start by noting that the adaptivity bound is given by combining [Lemma 8](#), the fact that the thresholds are guessed in parallel, and the fact that $\kappa(\mathcal{N}_+) \in O(\frac{n}{\varepsilon})$. We remark that now, since the cost condition is unimodal, binary search in THRESHSEQ works without any major adjustment.

Let O^* be the optimal solution and let $\tau^* = \alpha \frac{f(O^*)}{B}$. By the parallel guesses we have that there exists a $\tau = \tau_i$ such that $(1 - \hat{\varepsilon})\tau^* \leq \tau < \tau^*$. As in the non-monotone case, this is because $f(x^*) \geq f(O^*) \geq f(x^*)$. Let S be the random set outputted by THRESHSEQ for that τ . Also, let $T = S \cup \mathcal{N}_-$ and notice that $c(S \cup \mathcal{N}_-) \leq B$.

We can distinguish two cases. First, if $\mathbb{E}[c(S)] \geq \frac{B}{2}(1 - 2\hat{\varepsilon})(1 - \hat{\varepsilon})$, then we apply [Lemma 8](#) and we have

$$f(O^*) \leq \frac{2}{\alpha(1 - \hat{\varepsilon})^3(1 - 2\hat{\varepsilon})} f(S) \leq \frac{2}{\alpha(1 - \hat{\varepsilon})^3(1 - 2\hat{\varepsilon})} f(T). \quad (10)$$

Let's now address the other case. We can argue as we did in the monotone case: if we run $\frac{1}{\hat{\varepsilon}} \log(\frac{1}{\hat{\varepsilon}})$ independent times the algorithm, at least one of them respects $c(S) < B(\frac{1}{2} - \hat{\varepsilon})$, with

probability at least $(1 - \hat{\varepsilon})$ (the proof of this fact is practically the same as the one of [Lemma 6](#)). Let's call \mathcal{G} that event, similarly to what we have done in the main text. Focus on that run and call \mathcal{E} the event that in that run there is a good element with respect to the solution not fitting in the budget. Clearly there may be at most one such item which belongs to the optimal solution O^* ; we call such element \tilde{x} . If \tilde{x} exists, then $c(\tilde{x}) \geq \frac{B}{2}$. This is because in the budget of THRESHSEQ, i.e., $B(1 - \hat{\varepsilon})$, at least $\frac{B}{2}$ budget is empty, under \mathcal{G} .

$$\begin{aligned} f(O^*) &\leq f(S \cup O^*) = f(T \cup (O^* \setminus \mathcal{N}_-)) \leq f(T) + \mathbf{1}_{\mathcal{E}} f(\tilde{x} | T) + \sum_{x \in O^* \setminus \{\tilde{x}\}} f(x | T) \\ &\leq f(T) + \mathbf{1}_{\mathcal{E}} f(x^*) + \sum_{x \in O^* \setminus \{\tilde{x}\}} f(x | S) \leq f(T) + \mathbf{1}_{\mathcal{E}} (f(x^*) - \alpha \frac{f(O^*)}{2}) + \alpha f(O^*). \end{aligned}$$

Passing to the expectation and calling q the conditioned probability of the event \mathcal{E} given \mathcal{G} we have, similarly to the main text:

$$\begin{aligned} f(O^*) &\leq \mathbb{E}[f(O^* \cup S)] = \mathbb{E}[f(O^* \cup S) | \mathcal{G}] \mathbb{P}(\mathcal{G}) + \mathbb{E}[f(O^* \cup S) | \mathcal{G}^C] \mathbb{P}(\mathcal{G}^C) \\ &\leq \mathbb{E}[f(O^* \cup S) | \mathcal{G}] (1 - \hat{\varepsilon}) + \hat{\varepsilon} 2f(O^*) \\ &\leq (1 + q)(1 - \hat{\varepsilon})ALG + (2\hat{\varepsilon} + \alpha(1 - \hat{\varepsilon})(1 - \frac{q}{2}))f(O). \end{aligned}$$

Notice we used the bound $f(S \cup O^*) \leq 2f(O^*)$ which is universal as long as $c(S) \leq B$. Rearranging the terms we have

$$f(O^*) \leq \frac{(1 + q)(1 - \hat{\varepsilon})}{1 - 2\hat{\varepsilon} - \alpha(1 - \hat{\varepsilon})(1 - \frac{q}{2})} ALG. \quad (11)$$

Putting together [\(10\)](#) and [\(11\)](#) and setting $\alpha = \frac{2}{3}$, we have

$$OPT \leq (3 + 10\hat{\varepsilon})ALG,$$

for any value of q . Rescaling $\hat{\varepsilon}$ by a factor of 10 one yields the desired result. \square

C Non-Monotone Objectives and a Cardinality Constraint

In presence of cardinality constraints, there is no need to address separately small and large elements. Moreover, when bounding the elements of the solution whose marginal density is greater than τ but do not fit in the budget, we just need to consider the case $\mathbb{E}[c(S)] > (1 - \hat{\varepsilon})k$ instead of considering half of the ‘‘budget’’.

If $\mathbb{E}[|S|] \geq (1 - \hat{\varepsilon})k$ we have immediately a good bound in expectation. Otherwise, if we run it at least $\frac{1}{\hat{\varepsilon}} \log(\frac{1}{\hat{\varepsilon}})$ independent times, we have that, with probability at least $(1 - \hat{\varepsilon})$ the cardinality constraint k is not met, meaning that all the good elements belong to the set G , as defined in the main text. The full algorithm, PARCARDINAL, is presented in [Algorithm 6](#) below.

Theorem 7. *For $\varepsilon \in (0, 2/5)$ it is possible to achieve a $5.83 + \varepsilon$ approximation, in $O(\frac{1}{\varepsilon} \log n)$ adaptive rounds and $O(\frac{nk}{\varepsilon^3} \log n \log k \log \frac{1}{\varepsilon})$ queries, or in $O(\frac{1}{\varepsilon} \log n \log k)$ adaptive rounds and $O(\frac{n}{\varepsilon^3} \log n \log^2 k \log(\frac{1}{\varepsilon}))$ queries.*

Proof. PARCARDINAL with parameters $\alpha = 3 - 2\sqrt{2}$, $p = (1 - \alpha)/2$ and $\varepsilon \in (0, \frac{2}{5})$ does the job. The adaptive and query complexity are as in the knapsack case. The only difference is that now each sequence drawn from SAMPLESEQ has at most length k . For the approximation guarantees, we consider two cases, this time depending on the relative ordering of $\mathbb{E}[|S|]$ and $(1 - \hat{\varepsilon})k$.

Algorithm 6: PARCARDINAL($\mathcal{N}, f, \varepsilon, \alpha, B$)Full algorithm for non-monotone and a cardinality constraint

- 1: **Input:** Ground set \mathcal{N} , submodular function f , cardinality k , precision $\varepsilon \in (0, 1)$, parameter $\alpha \in (0, 1)$ and sampling probability p
 - 2: $x^* \leftarrow \max_{x \in \mathcal{N}} f(x)$, $\hat{\tau} \leftarrow \alpha n \frac{f(x^*)}{k}$
 - 3: $\hat{\varepsilon} \leftarrow \frac{1}{70} \varepsilon$, $\ell \leftarrow \frac{1}{\hat{\varepsilon}^2}$, $k \leftarrow \frac{1}{\hat{\varepsilon}} \log(n)$
 - 4: $H \leftarrow$ sample each element in \mathcal{N} independently at random with probability p
 - 5: **for** $i = 0, \dots, k$ in parallel **do**
 - 6: $\tau_i \leftarrow \hat{\tau} \cdot (1 - \hat{\varepsilon})^i$
 - 7: **for** $j = 1, \dots, \frac{1}{\hat{\varepsilon}} \log(\frac{1}{\hat{\varepsilon}})$ in parallel **do**
 - 8: $S_j^i \leftarrow \text{THRESHSEQ}(H, \tau_i, \hat{\varepsilon}, \ell, k)$
 - 9: $T \leftarrow \arg \max \{f(S_i^j), f(x^*)\}$
 - 10: **Return** T
-

If $\mathbb{E}[|S|] \geq (1 - \hat{\varepsilon})k$, then, by [Lemma 4](#), we have

$$f(O^*) \leq \frac{1}{\alpha(1 - \hat{\varepsilon})^4} \mathbb{E}[f(S)]. \quad (12)$$

Otherwise, with probability at least ε , each run of the algorithm does not fill all the cardinality constraint, meaning that with $\frac{1}{\hat{\varepsilon}} \log(\frac{1}{\hat{\varepsilon}})$ independent rounds one has that $|S| < k$ in at least one round with probability at least $1 - \hat{\varepsilon}$; as usual we call \mathcal{G} that event. Focus on that round and recall the definition of G from [Lemma 5](#): G contains the good elements still fitting in the cardinality constraint. We have

$$f(S \cup O_H) \leq f(S) + \sum_{x \in G} f(x|S) + \sum_{x \in O_H \setminus (S \cup G)} f(x|S) \leq f(S)(1 + \hat{\varepsilon}) + \tau c(O_H).$$

Passing to the conditional expectation with respect to \mathcal{G} and using $f(S \cup O_H) \leq 2f(O^*)$ we have

$$\begin{aligned} p(1 - p)f(O^*) &\leq \mathbb{E}[f(S \cup O_H)] \leq \mathbb{E}[f(S \cup O_H) | \mathcal{G}] (1 - \hat{\varepsilon}) + 2\hat{\varepsilon}f(O^*) \\ &\leq \mathbb{E}[f(S)] (1 + \hat{\varepsilon})(1 - \hat{\varepsilon}) + \alpha p(1 - \hat{\varepsilon})f(O^*) + 2\hat{\varepsilon}f(O^*). \end{aligned}$$

By rearranging the terms, we get

$$f(O^*) \leq \frac{(1 + \hat{\varepsilon})(1 - \hat{\varepsilon})}{p(1 - p - (1 - \hat{\varepsilon})\alpha - 2\frac{\hat{\varepsilon}}{p})} \mathbb{E}[f(S)], \quad (13)$$

Plugging $\alpha = 3 - 2\sqrt{2}$ and $p = \frac{1 - \alpha}{2}$ in [\(12\)](#) and [\(13\)](#) one gets

$$OPT \leq (3 + 2\sqrt{2} + 70\hat{\varepsilon})ALG.$$

If we rescale $\hat{\varepsilon} = \frac{\varepsilon}{70}$, we have the desired result. \square

References

- [1] G. Amanatidis, F. Fusco, P. Lazos, S. Leonardi, and R. Reiffenhäuser. Fast adaptive non-monotone submodular maximization subject to a knapsack constraint. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [2] E. Balkanski and Y. Singer. The adaptive complexity of maximizing a submodular function. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 1138–1151. ACM, 2018.
- [3] E. Balkanski, A. Breuer, and Y. Singer. Non-monotone submodular maximization in exponentially fewer iterations. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 2359–2370, 2018.
- [4] E. Balkanski, A. Rubinstein, and Y. Singer. An exponential speedup in parallel running time for submodular maximization without loss in approximation. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 283–302. SIAM, 2019.
- [5] E. Balkanski, A. Rubinstein, and Y. Singer. An optimal approximation for submodular maximization under a matroid constraint in the adaptive complexity model. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*, pages 66–77. ACM, 2019.
- [6] A. Breuer, E. Balkanski, and Y. Singer. The FAST algorithm for submodular maximization. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 1134–1143. PMLR, 2020.
- [7] N. Buchbinder, M. Feldman, J. Naor, and R. Schwartz. Submodular maximization with cardinality constraints. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 1433–1452. SIAM, 2014.
- [8] J. Carbinell and J. Goldstein. The use of MMR, diversity-based reranking for reordering documents and producing summaries. *SIGIR Forum*, 51(2):209–210, 2017.
- [9] C. Chekuri and K. Quanrud. Submodular function maximization in parallel via the multilinear relaxation. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 303–322. SIAM, 2019.
- [10] C. Chekuri and K. Quanrud. Parallelizing greedy for submodular set function maximization in matroids and beyond. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*, pages 78–89. ACM, 2019.
- [11] C. Chekuri, J. Vondrák, and R. Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. *SIAM J. Comput.*, 43(6):1831–1879, 2014.

- [12] L. Chen, M. Feldman, and A. Karbasi. Unconstrained submodular maximization with constant adaptive complexity. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*, pages 102–113. ACM, 2019.
- [13] A. Das and D. Kempe. Algorithms for subset selection in linear regression. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 45–54. ACM, 2008.
- [14] A. Das and D. Kempe. Approximate submodularity and its applications: Subset selection, sparse approximation and dictionary selection. *J. Mach. Learn. Res.*, 19:3:1–3:34, 2018.
- [15] D. Dueck and B. J. Frey. Non-metric affinity propagation for unsupervised image categorization. In *IEEE 11th International Conference on Computer Vision, ICCV 2007, Rio de Janeiro, Brazil, October 14-20, 2007*, pages 1–8. IEEE Computer Society, 2007.
- [16] A. Ene and H. L. Nguyen. Submodular maximization with nearly-optimal approximation and adaptivity in nearly-linear time. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 274–282. SIAM, 2019.
- [17] A. Ene and H. L. Nguyen. Parallel algorithm for non-monotone DR-submodular maximization. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 2902–2911. PMLR, 2020.
- [18] A. Ene, H. L. Nguyen, and A. Vladu. A parallel double greedy algorithm for submodular maximization. *CoRR*, abs/1812.01591, 2018.
- [19] A. Ene, H. L. Nguyen, and A. Vladu. Submodular maximization with matroid and packing constraints in parallel. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*, pages 90–101. ACM, 2019.
- [20] M. Fahrbach, V. S. Mirrokni, and M. Zadimoghaddam. Submodular maximization with nearly optimal approximation, adaptivity and query complexity. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 255–273. SIAM, 2019.
- [21] M. Fahrbach, V. S. Mirrokni, and M. Zadimoghaddam. Non-monotone submodular maximization with nearly optimal adaptivity and query complexity. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 1833–1842. PMLR, 2019.
- [22] U. Feige. A threshold of $\ln n$ for approximating set cover. *J. ACM*, 45(4):634–652, 1998.
- [23] U. Feige, V. S. Mirrokni, and J. Vondrák. Maximizing non-monotone submodular functions. *SIAM J. Comput.*, 40(4):1133–1153, 2011.
- [24] M. Feldman, J. Naor, and R. Schwartz. A unified continuous greedy algorithm for submodular maximization. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pages 570–579. IEEE Computer Society, 2011.

- [25] A. Gupta, A. Roth, G. Schoenebeck, and K. Talwar. Constrained non-monotone submodular maximization: Offline and secretary algorithms. In *Internet and Network Economics - 6th International Workshop, WINE 2010, Stanford, CA, USA, December 13-17, 2010. Proceedings*, volume 6484 of *Lecture Notes in Computer Science*, pages 246–257. Springer, 2010.
- [26] F. M. Harper and J. A. Konstan. The movielens datasets: History and context. *ACM Trans. Interact. Intell. Syst.*, 5(4):19:1–19:19, 2016.
- [27] J. D. Hartline, V. S. Mirrokni, and M. Sundararajan. Optimal marketing strategies over social networks. In *Proceedings of the 17th International Conference on World Wide Web, WWW 2008, Beijing, China, April 21-25, 2008*, pages 189–198. ACM, 2008.
- [28] E. Kazemi, M. Zadimoghaddam, and A. Karbasi. Scalable deletion-robust submodular maximization: Data summarization with privacy and fairness constraints. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 2549–2558. PMLR, 2018.
- [29] D. Kempe, J. M. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. *Theory Comput.*, 11:105–147, 2015.
- [30] R. Khanna, E. R. Elenberg, A. G. Dimakis, S. N. Negahban, and J. Ghosh. Scalable greedy feature selection via weak submodularity. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, AISTATS 2017, 20-22 April 2017, Fort Lauderdale, FL, USA*, volume 54 of *Proceedings of Machine Learning Research*, pages 1560–1568. PMLR, 2017.
- [31] A. Kuhnle. Nearly linear-time, parallelizable algorithms for non-monotone submodular maximization. *To appear in AAAI*, abs/2009.01947, 2021.
- [32] A. Kulik, H. Shachnai, and T. Tamir. Approximations for monotone and nonmonotone submodular maximization with knapsack constraints. *Math. Oper. Res.*, 38(4):729–739, 2013.
- [33] M. Minoux. Accelerated greedy algorithms for maximizing submodular set functions. In *Optimization Techniques*, pages 234–243, Berlin, Heidelberg, 1978. Springer Berlin Heidelberg.
- [34] B. Mirzasoleiman, A. Karbasi, R. Sarkar, and A. Krause. Distributed submodular maximization: Identifying representative elements in massive data. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pages 2049–2057, 2013.
- [35] B. Mirzasoleiman, A. Badanidiyuru, and A. Karbasi. Fast constrained submodular maximization: Personalized data summarization. In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 1358–1367. JMLR.org, 2016.
- [36] B. Mirzasoleiman, J. A. Bilmes, and J. Leskovec. Coresets for data-efficient training of machine learning models. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 6950–6960. PMLR, 2020.
- [37] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions - I. *Math. Program.*, 14(1):265–294, 1978.

- [38] M. Sviridenko. A note on maximizing a submodular set function subject to a knapsack constraint. *Oper. Res. Lett.*, 32(1):41–43, 2004.
- [39] S. Tschitschek, R. K. Iyer, H. Wei, and J. A. Bilmes. Learning mixtures of submodular functions for image collection summarization. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 1413–1421, 2014.
- [40] J. Yang and J. Leskovec. Defining and evaluating network communities based on ground-truth. *Knowledge and Information Systems*, 42(1):181–213, 2015.