

Fast Adaptive Non-Monotone Submodular Maximization Subject to a Knapsack Constraint

Georgios Amanatidis

*Department of Mathematical Sciences,
University of Essex, UK.*

GEORGIOS.AMANATIDIS@ESSEX.AC.UK

Federico Fusco

Philip Lazos

Stefano Leonardi

Rebecca Reiffenhäuser

*Department of Computer, Control,
and Management Engineering “Antonio Ruberti”,
Sapienza University of Rome, Italy.*

FUSCOF@DIAG.UNIROMA1.IT

LAZOS@DIAG.UNIROMA1.IT

LEONARDI@DIAG.UNIROMA1.IT

REBECCAR@DIAG.UNIROMA1.IT

Abstract

Constrained submodular maximization problems encompass a wide variety of applications, including personalized recommendation, team formation, and revenue maximization via viral marketing. The massive instances occurring in modern-day applications can render existing algorithms prohibitively slow. Moreover, frequently those instances are also inherently stochastic. Focusing on these challenges, we revisit the classic problem of maximizing a (possibly non-monotone) submodular function subject to a knapsack constraint. We present a simple randomized greedy algorithm that achieves a 5.83-approximation and runs in $O(n \log n)$ time, i.e., at least a factor n faster than other state-of-the-art algorithms. The versatility of our approach allows us to further transfer it to a stochastic version of the problem. There, we obtain a $(9 + \varepsilon)$ -approximation to the best adaptive policy, which is the first constant approximation for non-monotone objectives. Experimental evaluation of our algorithms showcases their improved performance on real and synthetic data.

1. Introduction

Constrained submodular maximization is a fundamental problem at the heart of discrete optimization. The reason for this is as simple as it is clear: submodular functions capture the notion of *diminishing returns* present in a wide variety of real-world settings.

Consequently to its striking importance and coinciding NP-hardness (Feige, 1998), extensive research has been conducted on submodular maximization since the seventies (e.g., Edmonds, 1971; Nemhauser et al., 1978), with focus lately shifting towards handling the massive datasets emerging in modern applications. With a wide variety of possible constraints, often regarding cardinality, independence in a matroid, or knapsack-type restrictions, the number of applications is vast. To name just a few, there are recent works on feature selection in machine learning (Das and Kempe, 2008, 2018; Khanna et al., 2017), influence maximization in viral marketing (Babaei et al., 2013; Kempe et al., 2015), data summarization (Sipos et al., 2012; Mirzasoleiman et al., 2013; Tschitschek et al., 2014; Dütting et al., 2022), and decision making under uncertainty (Shperberg and Shimony, 2017; Boodaghians et al., 2020). Many of these applications have *non-monotone* submod-

ular objectives, meaning that adding an element to an existing set might actually decrease its value. Two such examples are discussed in detail in Section 5.

Modern-day applications increasingly force us to face two distinct, but often entangled challenges. First, the massive size of occurring instances fuels a need for very fast algorithms. As the running time is dominated by the *objective function evaluations* (also known as *value oracle calls*), it is typically measured (as in this work) by their number. So, here the goal is to design algorithms requiring an almost linear number of such evaluations. There is extensive research focusing on this issue, be it in the standard algorithmic setting (Mirzasoleiman et al., 2016), in streaming (Badanidiyuru et al., 2014; Chekuri et al., 2015a; Feldman et al., 2018; Alaluf et al., 2020), in distributed submodular maximization (Mirzasoleiman et al., 2013; da Ponte Barbosa et al., 2015), or in the adaptive complexity framework (Kuhnle, 2021; Amanatidis et al., 2021). The second challenge is the inherent uncertainty in problems like sensor placement or revenue maximization, where one does not learn the exact marginal value of an element until it is added to the solution (and thus “paid for”). This, too, has motivated several works on adaptive submodular maximization (Golovin and Krause, 2011; Gotovos et al., 2015; Gupta et al., 2017; Mitrovic et al., 2019) or on adaptivity gap techniques (Bradac et al., 2019; Boodaghians et al., 2020). Note that even estimating the expected value to a partially unknown objective function can be very costly and this makes the reduction of the number of such calls all the more important.

Knapsack constraints are one of the most natural types of restriction that occurs in real-world problems and are often *hard* budget, time, or size constraints. Other combinatorial constraints like partition matroid constraints, on the other hand, model less stringent requirements, e.g., avoiding too many similar items in the solution. As the soft versions of such constraints can be often hardwired in the objective itself (see the *Video Recommendation* application in Section 5), we do not deal with them directly here.

The nearly-linear time requirement, without large constants involved, leaves little room for using sophisticated approaches like continuous greedy methods (Feldman et al., 2011) or enumeration of initial solutions (Sviridenko, 2004). To further highlight the delicate balance between function evaluations and approximation, it is worth mentioning that, even for the monotone case, the first result combining $O(n \log n)$ oracle calls with an approximation better than 2 is the very recent $\frac{e}{e-1}$ -approximation algorithm of Ene and Nguyen (2019). While this is a very elegant theoretical result, the huge constants involved render it unusable in practice.

At the same time, the strikingly simple 2.8-approximation *adapted density greedy* algorithm of Wolsey (1982) deals well with both issues in the *monotone* case: Sort the items in decreasing order according to their marginal value over cost ratio and pick as many items as possible in that order without violating the constraint. Finally, return the best among this solution and the best single item.¹ When combined with lazy evaluations (Minoux, 1978), this algorithm requires only $O(n \log n)$ value oracle calls and can be adjusted to work equally well for adaptive submodular maximization (Golovin and Krause, 2011). For *non-monotone* objectives, however, the only practical algorithm is the $(10 + \varepsilon)$ -approximation FANTOM algorithm of Mirzasoleiman et al. (2016) requiring $O(n^2 \log n)$ value oracle calls (see also

1. In terms of guarantees, this is equivalent to the algorithm described in Wolsey (1982) which considers the first item left out of the density greedy solution rather than the best single item. For simplicity, by *adapted density greedy* we will refer to the algorithm we state above.

Remark 1). Moreover, there is no known algorithm for the adaptive setting that can handle anything beyond a cardinality constraint (Gotovos et al., 2015).

We aim to tackle both aforementioned challenges for non-monotone submodular maximization under a knapsack constraint, by revisiting the simple algorithmic principle of Wolsey’s algorithm. Our approach is along the lines of recent results on *random greedy* combinatorial algorithms (Buchbinder et al., 2014; Feldman et al., 2017), which show that introducing randomness into greedy algorithms can extend their guarantees to the non-monotone case. Here we give the *first* such algorithm for a knapsack constraint.

To get some intuition regarding the need for (and use of) randomization, consider the following example which also demonstrates that the adapted density greedy algorithm may produce arbitrarily poor solutions when the objective is non-monotone. Suppose there are $n + 1$ items, x_1, \dots, x_n and y . The value of a set S is either its cardinality if $y \notin S$, or $1 + \epsilon$ otherwise (it is easy to check that this is a non-monotone submodular function). All items have weight $1/n$ and our knapsack size is 1. Note that the adapted density greedy and the standard greedy algorithms would both start by adding y , thus producing a final solution of value $1 + \epsilon$, whereas the optimal value is n . Now suppose that we allow adapted density greedy to avoid each item with a certain constant probability p : it avoids y with probability p in its first step, leading to a solution with expected value $(1 - p) \cdot (1 + \epsilon) + p^2 \cdot n = \Theta(n)$.

1.1 Contribution and Outline

In this work we show that introducing some randomization to the adapted density greedy algorithm leads to a simple algorithm, SAMPLEGREEDY, that outperforms existing algorithms both in theory and in practice. SAMPLEGREEDY flips a coin before greedily choosing any item in order to decide whether to include it to the solution or ignore it. The algorithmic simplicity of such an approach keeps SAMPLEGREEDY fast, easy to implement, and flexible enough to adjust to other related settings. At the same time the added randomness prevents it from getting trapped in solutions of poor quality.

In Section 3 we show that SAMPLEGREEDY is a $(3 + 2\sqrt{2} + \epsilon)$ -approximation algorithm using $O(n\epsilon^{-1} \log(n/\epsilon))$ value oracle calls; specifically, this implies a 5.83-approximation algorithm that makes only $O(n \log n)$ calls. When all singletons have small value compared to an optimal solution, the approximation factor improves to almost 4. This is the *first* constant-factor approximation algorithm for the non-monotone case using this few queries. The only other algorithm fast enough to be suitable for large instances is the aforementioned FANTOM (Mirzasoleiman et al., 2016) which, for a knapsack constraint,² achieves an approximation factor of $(10 + \epsilon)$ with $O(nr\epsilon^{-1} \log n)$ queries, where r is the size of the largest feasible set and can be as large as $\Theta(n)$. Even if we modify FANTOM to use lazy evaluations, we still improve the query complexity by a logarithmic factor (see also Remark 1).

Then we study the problem in the adaptive submodular maximization framework of Golovin and Krause (2011) and Gotovos et al. (2015), where the stochastic submodular objective is learned as we build the solution and its value depends only on the state of the elements in the evaluated set. For this adaptive variant, we show in Section 4 that a natural

2. FANTOM can handle more general constraints, like a p -system constraint and ℓ knapsack constraints. Here we refer to its performance and running time when restricted to a single knapsack constraint.

adaptation of our algorithm, ADAPTIVEGREEDY, still guarantees a $(9 + \varepsilon)$ -approximation to the *best adaptive policy*. This is not only a relatively small loss given the considerably stronger benchmark, but is in fact the *first* constant approximation known for the problem in this framework. Hence we fill a notable theoretical gap, given that models with incomplete prior information, or those capturing evolving settings, are becoming increasingly important in practice.

From a technical point of view, our algorithm combines the simple principle of always choosing a high-density item with maintaining a careful exploration-exploitation balance, as is the case in many stochastic learning problems. Adding greedily elements might result in local maxima of poor quality, so we need some randomization to use some budget on the exploration of other elements. Given these features, this paper is therefore directly related to the recent simple randomized greedy approaches for maximizing non-monotone submodular objectives subject to other (i.e., non-knapsack) constraints (Buchbinder et al., 2014; Chekuri et al., 2015a; Feldman et al., 2017). However, there are underlying technical difficulties that make the analysis for knapsack constraints significantly more challenging. Every single result in this line of work critically depends on making a random choice in each step, in a way so that “good progress” is consistently made. This is not possible under a knapsack constraint. Instead, we argue globally about the value of the SAMPLEGREEDY output via a comparison with a carefully maintained *almost integral* solution. When it comes to extending this approach to the adaptive non-monotone submodular maximization framework, we crucially use the fact that the algorithm builds the solution iteratively, committing in every step to all the past choices. This is the main technical reason why it is not possible to adjust algorithms with multiple “parallel” runs, like FANTOM, to the adaptive setting.

Our algorithms provably handle well the aforementioned emerging, modern-day challenges, i.e., stochastically evolving objectives and rapidly growing real-world instances. In Section 5 we showcase the fact that our theoretical results indeed translate into applied performance. We focus on three applications that fit within the framework of non-monotone submodular maximization subject to a knapsack constraint, namely *video recommendation*, *influence-and-exploit marketing*, and *influence maximization*. We run experiments on real and synthetic data that indicate that SAMPLEGREEDY consistently performs better than FANTOM while being much faster. For ADAPTIVEGREEDY we highlight the fact that its adaptive behavior results in a significant improvement over non-adaptive alternatives.

1.2 Related Work

There is an extensive literature on submodular maximization subject to knapsack or other constraints, going back several decades (see, e.g., Nemhauser et al., 1978; Wolsey, 1982). For a *monotone* submodular objective subject to a knapsack constraint there is a deterministic $\frac{e}{e-1}$ -approximation algorithm (Khuller et al., 1999; Sviridenko, 2004) which is tight, unless $P = NP$ (Feige, 1998). This algorithm has a running time of $O(n^5)$ but there are other, much faster, greedy approaches with weaker approximation guarantees, like Wolsey’s 2.8-approximation algorithm (Wolsey, 1982) used as a starting point here, and the recent 2-approximation algorithm of Yaroslavtsev et al. (2020).

On non-monotone submodular functions, Lee et al. (2010) provided a 5-approximation algorithm for k knapsack constraints, which was the first constant factor algorithm for the problem. Fadaei et al. (2011) building on the approach of Lee et al. (2010), reduced this factor to 4. One of the most interesting algorithms for a single knapsack constraint is the 6-approximation algorithm of Gupta et al. (2010). As this is a greedy combinatorial algorithm based on running Sviridenko’s algorithm twice, it is often used as a subroutine by other algorithms in the literature, e.g., (da Ponte Barbosa et al., 2015), despite its running time of $O(n^4)$. A number of continuous greedy approaches (Feldman et al., 2011; Kulik et al., 2013; Chekuri et al., 2014) led to the current best factor of e when a knapsack—or even a general downwards closed—constraint is involved. However, continuous greedy algorithms are impractical for most real-world applications. The fastest such algorithm for our setting is the $(e + \varepsilon)$ -approximation algorithm of Chekuri et al. (2015b) requiring $O(n^3\varepsilon^{-4}\text{polylog}(n))$ function evaluations. Possibly the only algorithm that is directly comparable to our SAMPLEGREEDY in terms of running time is FANTOM by Mirzasoleiman et al. (2016). FANTOM achieves a $(1 + \varepsilon)(p + 1)(2p + 2\ell + 1)/p$ -approximation for ℓ knapsack constraints and a p -system constraint in time $O(nr p \varepsilon^{-1} \log(n))$, where r is the size of the largest feasible solution.

As mentioned above, there is a number of recent results on randomizing simple greedy algorithms so that they work for non-monotone submodular objectives (Buchbinder et al., 2014; Chekuri et al., 2015a; Gotovos et al., 2015; Feldman et al., 2017; Feldman and Zenklusen, 2018). Our paper extends this line of work, as we are the first to successfully apply this approach for a knapsack constraint.

Golovin and Krause (2011) introduced the notions of adaptive monotonicity and submodularity and showed it is possible to achieve guarantees with respect to the optimal adaptive policy that are similar to the guarantees one gets in the standard algorithmic setting with respect to an optimal solution. Our Section 4 fits into this framework as it was generalized by Gotovos et al. (2015) for non-monotone objectives. Gotovos et al. (2015) showed that a variant of the random greedy algorithm of Buchbinder et al. (2014) achieves a $\frac{e}{e-1}$ -approximation in the case of a cardinality constraint. Tang (2021a) recently presented a different analysis of the algorithm of Gotovos et al. (2015) that yields the same approximation factor when the objective is adaptive submodular and possibly non-pointwise submodular. Finally, in a very recent unpublished manuscript, Tang (2021b) claims that (an equivalent version of) SampleGreedy achieves a constant factor approximation for the problem of adaptive submodular maximization subject to a knapsack constraint even without resorting to the assumption of pointwise adaptive submodularity that is typically made in the literature (Krause et al., 2008; Gotovos et al., 2015; Amanatidis et al., 2020).

Implicitly related to our quest for few value oracle calls is the recent line of work on the adaptive complexity of submodular maximization that measures the number of sequential rounds of independent value oracle calls needed to obtain a constant factor approximation (see Balkanski and Singer, 2018; Balkanski et al., 2019; Ene et al., 2019; Fahrback et al., 2019b,a; Kuhnle, 2021; Amanatidis et al., 2021, and references therein). For non-monotone functions and a knapsack constraint, in particular, Ene et al. (2019) and Amanatidis et al. (2021) give $O(1)$ -approximation algorithms that need $O(\log^2(n))$ and $O(\log(n))$ rounds of independent value oracle calls, respectively. While the elaborate continuous approach of Ene et al. (2019) requests $O(n^2)$ value queries, Amanatidis et al. (2021) uses only $O(n \log^3 n)$

value queries. While the latter result is posterior to our work, we are still faster by a $\log^2 n$ factor, while retaining a better approximation factor.

2. Preliminaries

In this section we formally introduce the problem of submodular maximization with a knapsack constraint in both the standard and the adaptive setting.

Submodular functions. Let $E = [n] = \{1, 2, \dots, n\}$ be a set of n items, a function $v : 2^E \rightarrow \mathbb{R}$ is *submodular* if and only if

$$v(S \cup \{i\}) - v(S) \geq v(T \cup \{i\}) - v(T)$$

for all $S \subseteq T \subseteq E$ and $i \notin T$.

A function $v : 2^E \rightarrow \mathbb{R}$ is *non-decreasing* (often referred to as *monotone*), if $v(S) \leq v(T)$ for any $S \subseteq T \subseteq E$. We consider general (i.e., not necessarily monotone), normalized (i.e., $v(\emptyset) = 0$), non-negative submodular valuation functions. Since *marginal* values are extensively used, we adopt the shortcut $v(T | S)$ for the marginal value of set T with respect to a set S , i.e. $v(T | S) = v(T \cup S) - v(S)$. If $T = \{i\}$ we write simply $v(i | S)$. While this is the most standard definition of submodularity in this setting, there are alternative equivalent definitions that will be useful later.

Theorem 1 (Nemhauser et al. (1978)). *Given a function $v : 2^E \rightarrow \mathbb{R}$, the following are equivalent:*

1. $v(i | S) \geq v(i | T)$ for all $S \subseteq T \subseteq E$ and $i \notin T$. (original definition)
2. $v(S) + v(T) \geq v(S \cup T) + v(S \cap T)$ for all $S, T \subseteq E$.
3. $v(T) \leq v(S) + \sum_{i \in T \setminus S} v(i | S) - \sum_{i \in S \setminus T} v(i | S \cup T \setminus \{i\})$ for all $S, T \subseteq E$.

Moreover, we restate a key result which connects random sampling and submodular maximization. The original version of the theorem was due to Feige et al. (2011), although here we use a variant from Buchbinder et al. (2014).

Lemma 1 (Lemma 2.2 of Buchbinder et al. (2014)). *Let $v : 2^E \rightarrow \mathbb{R}$ be a (possibly not normalized) submodular set function, let $X \subseteq E$ and let $X(p)$ be a sampled subset, where each element of X appears with probability at most p (not necessarily independent). Then*

$$\mathbb{E}[v(X(p))] \geq (1 - p)v(\emptyset).$$

We assume access to a *value oracle* that returns $v(S)$ when given as input a set S .

Knapsack Constraint. We associate a positive cost c_i with each element $i \in E$ and consider a given budget B . The goal of our problem is to find a subset of E of maximum value among the subsets whose total cost is at most B . Formally, we want some $S^* \in \arg \max\{v(S) \mid S \subseteq E, \sum_{i \in S} c_i \leq B\}$. In the following we denote with OPT the value of S^* . Without loss of generality, we may assume that $c_i \leq B$ for all $i \in E$, since any element with cost exceeding B is not contained in any feasible solution and can be discarded. Given

the hardness of this optimization problem, in this work we focus on finding solution that approximates the optimal solution; in particular, given an algorithm, we denote with ALG the expected value of the set output by it and we say that it gives an α approximation, for $\alpha \in [1, +\infty)$, if $OPT \leq \alpha \cdot \text{ALG}$.

Adaptive submodularity. We now present the adaptive optimization framework (Krause et al., 2008). On a high level, in many applications (e.g. sensor placing, traffic control and influence maximization) we do not know how the world works and what situations occur with which probability. However, which of those we will be actually dealing with is inferred over time by the bits of information we learn. To model these situations, along with set E , we introduce the *state space* Ω which is endowed with some probability measure. By $\omega = (\omega_i)_{i \in E} \in \Omega$ we specify the *state* of each element in E . The adaptive valuation function v is then defined over $E \times \Omega$; the value of a subset $S \subseteq E$ depends on both the subset and ω . Due to the probability measure over Ω , $v(S, \omega)$ is a random variable. We define $v(S) = \mathbb{E}[v(S, \omega)]$, the expectation being with respect to ω . Like before, the costs c_i are deterministic and known in advance.

For each $\omega \in \Omega$ and $S \subseteq E$, we define the partial realization of state ω on S as the couple $(S, \omega|_S)$, where $\omega|_S = (\omega_i)_{i \in S}$. It is natural to assume that the true value of a set S does not depend on the whole state, but only on $\omega|_S$, i.e., $v(S, \omega) = v(S, \psi)$, for all $\omega, \psi \in \Omega$ such that $\omega|_S = \psi|_S$. Therefore, sometimes we overload the notation and use $v(S, \omega|_S)$ instead of $v(S, \omega)$. There is a clear partial ordering on the set of all possible partial realizations: $(S, \omega|_S) \subseteq (T, \omega|_T)$ if $S \subseteq T$ and $\omega|_T$ coincides with $\omega|_S$ over all the elements of S . The marginal value of a set S' given a partial realization $(S, \omega|_S)$ is

$$v(S' \mid (S, \omega|_S)) = \mathbb{E}[v(S' \cup S, \omega) - v(S, \omega) \mid \omega|_S].$$

We are now ready to introduce the concepts of adaptive submodularity and monotonicity.

Definition 1. The function $v(\cdot, \cdot)$ is *adaptive submodular* if $v(i \mid (S, \omega|_S)) \geq v(i \mid (T, \omega|_T))$ for all partial realizations $(S, \omega|_S) \subseteq (T, \omega|_T)$ and for any $i \notin T$. Further, $v(\cdot, \cdot)$ is *adaptive monotone* if $v(i \mid (S, \omega|_S)) \geq 0$ for all partial realizations $(S, \omega|_S)$ and for all $i \notin S$.

In Section 4 we assume access to a value oracle that given an element i and a partial realization returns the expected marginal value of i . Using the properties of conditional expectation, it is straightforward to show that if $v(\cdot, \cdot)$ is adaptive submodular, then its expected value $v(\cdot)$ is submodular. In analogy with Gotovos et al. (2015), we assume v to be pointwise submodular, i.e., $v(\cdot, \omega)$ is a submodular set function for each $\omega \in \Omega$.

In this framework it is possible to define *adaptive policies* to maximize v . An adaptive policy π is a function which associates with every partial realization a distribution on the next element to be added to the solution. The optimal solution to the adaptive submodular maximization problem is to find an adaptive policy that maximizes the expected value while respecting the knapsack constraint (the expectation being taken over Ω and the randomness of the policy itself). Notice that the knapsack constraint has to be respected point-wise, i.e. for each realization of the randomization of the algorithm and of the state. The definition of approximation is defined similarly as in the non-adaptive case.

The following result of Gotovos et al. (2015) that holds even for non-monotone objectives will come handy later.

Lemma 2 (Lemma 1 of Gotovos et al. (2015)). *If v is adaptive submodular, then, for any $k \in \mathbb{N}$, any policy π that terminates after k steps, and any partial realization $(S, \omega|_S)$, it holds that*

$$\mathbb{E} [v(S'(\pi) \mid (S, \omega|_S))] \leq \sum_{i \in M_k} v(i \mid (S, \omega|_S)),$$

where $S'(\pi)$ is the (random) set selected according to the policy π and M_k the set containing k elements with the highest marginal values given $(S, \omega|_S)$. The expectation on the left side is with respect to the randomness of $S'(\pi)$.

We present now a brief example to help the reader familiarizing with the notion of adaptive submodularity. Consider the task of finding a max cut in a graph, where edge weights are randomly extracted from some known distribution. In this problem a *state* corresponds to an actual realization of the edge weights on the whole graph, while partial realizations of a state refer to the restriction of the realized edge weights on a specific subset of the edges. The distributions on different edges may be correlated, so every time the actual weight of an edge is revealed all the prior relative to unknown edge weights are updated. Finally, the adaptive submodular objective function depends on the topology of the graph (that is known up front) *and* on the realized weights. More on this example can be found in the experimental section.

3. A Simple Algorithm for the Standard Algorithmic Setting

We present and analyze SAMPLEGREEDY, a randomized 5.83-approximation algorithm for maximizing a submodular function subject to a knapsack constraint. As we mentioned already, SAMPLEGREEDY is based on the adapted density greedy algorithm of Wolsey (1982). Since the latter may perform arbitrarily bad for non-monotone objectives, we add a sampling phase, similar to the sampling phase of the Sample Greedy algorithm of Feldman et al. (2017). First we present in detail the algorithm and its analysis, then we show how a lazy implementation is enough to achieve $O(n \log n)$ sample complexity and finally how better approximation guarantees are achievable in the large instance scenario.

The algorithm. SAMPLEGREEDY first selects a random subset E' of E by independently picking each element with probability p . Then it runs Wolsey’s algorithm only on E' . To formalize this second step, using $v(i)$ as a shorthand for $v(\{i\})$, let

$$j_1 \in \arg \max_{i \in E'} v(i)/c_i,$$

while, for $k \geq 1$,

$$j_{k+1} \in \arg \max_{i \in E' \setminus \{j_1, \dots, j_k\}} v(i \mid \{j_1, \dots, j_k\})/c_i.$$

If ℓ is the largest integer such that $\sum_{i=1}^{\ell} c_{j_i} \leq B$, then $S = \{j_1, \dots, j_{\ell}\}$. In the end, the output is the one yielding the largest value between S and an element from $\arg \max_{i \in E} v(i)$.

We formally present this algorithm in pseudocode below. Notice that to simplify the analysis, instead of selecting the entire set E' immediately at the start of the algorithm, we defer this decision and toss a coin with probability of success p each time an item is considered to be added to the solution. Both versions of the algorithm behave identically.

```

SAMPLEGREEDY( $E, v, \mathbf{c}, B$ )
1  $i^* \in \arg \max_{j \in E} v(j)$  /* best single item */
2  $S = \emptyset$  /* greedy solution */
3  $F = \{j \in E \mid v(j \mid S) > 0\}$  /* initial set of feasible items */
4  $R = B$  /* remaining knapsack capacity */
5 while  $F \neq \emptyset$  do
6   Let  $i \in \arg \max_{j \in F} \frac{v(j \mid S)}{c_j}$ 
7   Let  $r_i \sim \text{Bernoulli}(p)$  /* independent random bit */
8   if  $r_i = 1$  then
9      $S = S \cup \{i\}$ 
10     $R = R - c_i$ 
11     $E = E \setminus \{i\}$ 
12     $F = \{j \in E \mid v(j \mid S) > 0 \text{ and } c_j \leq R\}$ 
13 return  $\arg \max\{v(i^*), v(S)\}$ 

```

Theorem 2. For $p = \sqrt{2} - 1$, SAMPLEGREEDY is a $(3 + 2\sqrt{2})$ -approximation algorithm.

Proof. For the analysis of the algorithm we are going to use the auxiliary set O , an extension of the set S that respects the knapsack constraint and uses feasible items from an optimal solution. In particular, let S^* be an optimal solution and let s_1, s_2, \dots, s_r be its elements, sorted in increasing order of cost.

Then, O is a *fuzzy* set that is initially equal to S^* and during each iteration of the *while* loop it is updated as follows:

- If $r_i = 1$, then $O = O \cup \{i\}$. In case this addition violates the knapsack constraint, i.e., $c(O) > B$, then we repetitively remove items from $O \setminus S$ in increasing order with respect to their cost until the cost of O becomes exactly B . Note that this means that the last item removed may be removed only partially. More precisely, if $c(O) > B$ and $c(O \setminus \{s_j\}) \leq B$, where s_j is the item of S^* of maximum index in $O \setminus S$, then we keep a $(B - c(O) + c_j) / c_j$ fraction of s_j in O and stop its update for the current iteration.
- Else (i.e., if $r_i = 0$), $O = O \setminus \{i\}$.

If an item j was considered (in line 6) in some iteration of the *while* loop, then let S_j and O_j denote the sets S and O , respectively, at the beginning of that iteration. Moreover, let O'_j denote O at the end of that iteration. If j was never considered, then S_j and O_j (or O'_j) denote the final versions of S and O , respectively. In fact, in what follows we exclusively use S and O for their final versions.

It should be noted that, for all $j \in E$, $S_j \subseteq O_j$ and also no item in $O_j \setminus S_j$ has been considered in any of the previous iterations of the *while* loop.

Before stating the next lemma, let us introduce some notation for the sake of readability. Note that, by construction, $O \setminus S$ is either empty or consists of a single fractional item \hat{i} . In case $O \setminus S = \emptyset$, by \hat{i} we denote the last item removed from O . For every $j \in E$, we define

$Q_j = O_j \setminus (O'_j \cup S \cup \{\hat{i}\})$. Note that if j was never considered during the execution of the algorithm, then $Q_j = \emptyset$.

Lemma 3. *For every realization of the Bernoulli random variables, it holds that*

$$v(S \cup S^*) \leq v(S) + v(\hat{i}) + \sum_{j \in E} c(Q_j) \frac{v(j | S_j)}{c_j}.$$

Proof of Lemma 3. Assume that the random bits r_1, r_2, \dots are fixed. Also, without loss of generality, assume the items are numbered according to the order in which they are considered by SAMPLEGREEDY, with the ones not considered by the algorithm numbered arbitrarily (but after the ones considered). That is, item j —if considered—is the item considered during the j^{th} iteration.

Consider now any round j of the while loop of SAMPLEGREEDY. An item is removed from O_j in two cases. First, it could be item j itself that was originally in S^* but $r_j = 0$ (and hence it will never get back in O_k for any $k > j$). Second, it could be some other item that was in S^* and is taken out to make room for the new item j . In the latter case the only possibility for the removed item to return in O_k for some $k > j$ is to be selected by the algorithm and inserted in S . We can hence conclude that $Q_j \cap Q_k = \emptyset$ for all $j \neq k$. In addition to that, it is clear that $S \cup S^* = S \cup \{\hat{i}\} \cup_{j=1}^{\ell} Q_j$.

Therefore, if items $1, 2, \dots, \ell$ where all the items ever considered, using submodularity and the fact that $S_j \subseteq S \subseteq S \cup_{r=j+1}^{\ell} Q_r$, we have

$$\begin{aligned} v(S \cup S^*) - v(S) - v(\hat{i}) &\leq v((S \cup S^*) \setminus \hat{i}) - v(S) = \sum_{j=1}^{\ell} v(Q_j | S \cup_{r=j+1}^{\ell} Q_r) \\ &\leq \sum_{j=1}^{\ell} v(Q_j | S_j) \leq \sum_{j=1}^{\ell} \sum_{x \in Q_j} \frac{v(x | S_j)}{c_x} \cdot c_x \\ &\leq \sum_{j=1}^{\ell} \sum_{x \in Q_j} \frac{v(j | S_j)}{c_j} \cdot c_x = \sum_{j=1}^{\ell} \frac{v(j | S_j)}{c_j} \cdot \sum_{x \in Q_j} c_x \\ &= \sum_{j=1}^{\ell} \frac{v(j | S_j)}{c_j} \cdot c(Q_j), \end{aligned} \tag{1}$$

where in a slight abuse of notation we consider c_x to be the fractional (linear) cost if $x \in Q_j$ is a fractional item. While the first three inequalities directly follow from the submodularity of v , for the last inequality we need to combine the optimality of $v(j | S_j)/c_j$ at the step j was selected with the fact that every single item x appearing in the sum $\sum_{j=1}^{\ell} \sum_{x \in Q_j} v(x | S_j)$ was feasible (as a whole item) at that step. The latter is true because of the way we remove items from O . If x is removed, it is removed before (any part of) \hat{i} is removed. Thus, x is removed when the available budget is still at least $c_{\hat{i}}$. Given that $c_x \leq c_{\hat{i}}$, we get that x is feasible until removed.

To conclude the proof of the Lemma it is sufficient to note that $c(Q_j) = 0$ for all items that were not considered. \square

While the previous Lemma holds for *each* realization of the random coin tosses in the algorithm, we next consider inequalities holding in expectation over the randomness of the $\{r_i\}_{i=1}^{|E|}$ in SAMPLEGREEDY. The indexing of the elements is hence to be considered deterministic and fixed in advance, not as in the proof of Lemma 3.

Lemma 4.
$$\mathbb{E} \left[\sum_{j \in E} c(Q_j) \frac{v(j | S_j)}{c_j} \right] \leq \frac{\max\{p, 1 - p\}}{p} \mathbb{E}[v(S)] .$$

Proof of Lemma 4. For all $i \in E$, we define G_i to be the random gain because of i at the time i is added to the solution ($G_i = v(i | S_i)$ if i is added and 0 otherwise)

Since $v(S) = \sum_{i \in E} G_i$, by linearity, it suffices to show that the following inequalities hold in expectation over the coin tosses:

$$c(Q_i) \frac{v(i | S_i)}{c_i} \leq \frac{\max\{p, 1 - p\}}{p} G_i, \quad \text{for all } i \in E. \tag{2}$$

In order to achieve that, following Feldman et al. (2017), let \mathcal{E}_i be any event specifying the random choices of the algorithm up to the point i is considered (if i is never considered, \mathcal{E}_i captures all the randomness). If \mathcal{E}_i is an event that implies i is not considered, then the Eq. (2) is trivially true, due to $G_i = 0$ and $Q_i = \emptyset$. We focus now on the case \mathcal{E}_i implies that i is considered. Analyzing the algorithm, it is clear that

$$\mathbb{E}[c(Q_i) | \mathcal{E}_i] \leq \begin{cases} 0 \cdot \mathbb{P}(r_i = 1) + c_i \cdot \mathbb{P}(r_i = 0) = (1 - p) \cdot c_i, & \text{if } i \in O_i, \\ c_i \cdot \mathbb{P}(r_i = 1) + 0 \cdot \mathbb{P}(r_i = 0) = p \cdot c_i, & \text{otherwise.} \end{cases}$$

In short, $\mathbb{E}[c(Q_i) | \mathcal{E}_i] \leq \max\{p, 1 - p\} \cdot c_i$. It is here that we use the fuzziness of O : without the fractional items it would be hopeless to bound $c(Q_t)$ with c_t .

At this point, we exploit the fact that \mathcal{E}_i contains the information on S_i , i.e., $S_i = S_i(\mathcal{E}_i)$ deterministically. Recall that S_i is the solution set at the time item i is considered by the algorithm.

$$\begin{aligned} \mathbb{E}[G_i | \mathcal{E}_i] &= \mathbb{P}(i \in S | \mathcal{E}_i) v(i | S_i) = \mathbb{P}(r_i = 1) v(i | S_i) = p \cdot c_i \frac{v(i | S_i)}{c_i} \\ &\geq \frac{p}{\max\{p, 1 - p\}} \mathbb{E}[c(Q_i) | \mathcal{E}_i] \frac{v(i | S_i)}{c_i} \\ &= \frac{p}{\max\{p, 1 - p\}} \mathbb{E} \left[c(Q_i) \frac{v(i | S_i)}{c_i} \mid \mathcal{E}_i \right]. \end{aligned}$$

We can hence conclude the proof by using the law of total probability over \mathcal{E}_i and the monotonicity of the conditional expectation:

$$\begin{aligned} \mathbb{E}[G_i] &= \mathbb{E}[\mathbb{E}[G_i | \mathcal{E}_i]] \geq \mathbb{E} \left[\frac{p}{\max\{p, 1 - p\}} \mathbb{E} \left[c(Q_i) \frac{v(i | S_i)}{c_i} \mid \mathcal{E}_i \right] \right] = \\ &= \frac{p}{\max\{p, 1 - p\}} \mathbb{E} \left[c(Q_i) \frac{v(i | S_i)}{c_i} \right]. \quad \square \end{aligned}$$

Lemma 5. $v(S^*) \leq \frac{1}{1 - p} \mathbb{E}[v(S \cup S^*)].$

Proof of Lemma 5. Let S^* be an optimal set for the constrained submodular maximization problem. We define $g : 2^E \rightarrow \mathbb{R}_+$ as follows: $g(T) = v(T \cup S^*)$. It is a simple exercise to see that such function is indeed submodular, moreover $g(\emptyset) = v(S^*)$. If we now apply Lemma 1 to g , observing that the elements in the set S output by the algorithm are chosen with probability at most p , we conclude that:

$$\mathbb{E}[v(S \cup S^*)] = \mathbb{E}[g(S)] \geq (1 - p)g(\emptyset) = (1 - p)v(S^*). \quad \square$$

Combining Lemmata 3, 4 and 5 we get

$$\begin{aligned} (1 - p)v(S^*) &\leq \mathbb{E}[v(S \cup S^*)] \\ &\leq \mathbb{E}\left[v(S) + v(\hat{i}) + \sum_{j \in E} c(Q_j) \frac{v(j | S_j)}{c_j}\right] \\ &\leq \mathbb{E}[v(S)] + v(i^*) + \frac{\max\{p, 1 - p\}}{p} \mathbb{E}[v(S)] \\ &= \max\left\{2, \frac{1}{p}\right\} \cdot \mathbb{E}[v(S)] + v(i^*). \end{aligned} \quad (3)$$

By substituting $\sqrt{2} - 1$ for p , this yields

$$\begin{aligned} OPT = v(S^*) &\leq (3 + 2\sqrt{2}) \max\{\mathbb{E}[v(S)], v(i^*)\} \\ &\leq (3 + 2\sqrt{2}) \mathbb{E}[\max\{v(S), v(i^*)\}] = (3 + 2\sqrt{2}) \text{ALG}, \end{aligned}$$

where the second inequality follows from Jensen's inequality. □

Lazy implementation of SampleGreedy. A naive implementation of SAMPLEGREEDY needs $\Theta(n^2)$ value oracle calls in the worst case. Indeed, in each iteration all the remaining elements have their marginals updated and for large enough B the greedy solution may contain a constant fraction of E . Applying lazy evaluations (Minoux, 1978), however, we can cut the number of queries down to $O(n\varepsilon^{-1} \log(n/\varepsilon))$ losing only an additive ε in the approximation factor (Ene and Nguyen, 2019). To achieve this, instead of recomputing all the marginals at every step, we maintain an ordered queue of the elements sorted by their last known *densities* (i.e., their marginal value per cost ratios) and use it to get a *sufficiently good* element to add.

More formally, the lazy implementation of SAMPLEGREEDY maintains the elements in a priority queue in decreasing order of density, which is initialised using the ratios $v(i)/c_i$. At each step we pop the element on top of the queue. If its density with respect to the current solution is within a $1 + \varepsilon$ factor of its old one, then it is picked by the algorithm, otherwise it is reinserted in the queue according to its new density and we pop the next element. Submodularity guarantees that the density of a picked element is at least $1/(1 + \varepsilon)$ of the best density for that step. As soon as an element has been updated $\log(n/\varepsilon)/\varepsilon$ times, we discard it.

Theorem 3. *The lazy version of SAMPLEGREEDY achieves an approximation factor of $3 + 2\sqrt{2} + \varepsilon$ using $O(n\varepsilon^{-1} \log(n/\varepsilon))$ value oracle calls.*

Proof. For a given $\varepsilon \in (0, 1)$ let $\varepsilon' = \varepsilon/6$. We perform lazy evaluations using ε' , with \log denoting the binary logarithm.

It is straightforward to argue about the number of value oracle calls. Since the marginal value of each element i has been updated at most $\frac{\log(n/\varepsilon')}{\varepsilon'}$ times, we have a total of at most $n \frac{\log(n/\varepsilon')}{\varepsilon'} = O\left(\frac{n \log(n/\varepsilon)}{\varepsilon}\right)$ function evaluations.

The approximation ratio is also easy to show. There are two distinct sources of loss in approximation. We first bound the total value of the discarded elements due to too many updates. This value appears as the upper bound of an extra additive term in the first line of (1). Indeed, now besides $\sum_{j=1}^{\ell} v(Q_j | S \cup_{r=j+1}^{\ell} Q_r)$ we need to account for the elements of O that were ignored because of too many updates. Such elements, once they become “inactive” do not contribute to the cost of the current O and are never pushed out as new elements come into S . The definition of the Q_j s in the proof of Theorem 2 should be adjusted accordingly. That is, if W_j are the elements of O that become inactive because they were updated too many times during iteration j , we have

$$v((S \cup S^*) \setminus \hat{i}) - v(S) \leq \sum_{j=1}^{\ell} v(Q_j | S_j) + \sum_{j=1}^{\ell} v(W_j | S_j).$$

However, by noticing that for $x \in (0, 1)$ it holds that $x \leq \log(1 + x)$, we have

$$\begin{aligned} \sum_{j=1}^{\ell} v(W_j | S_j) &\leq \sum_{i \in \bigcup_j W_j} (1 + \varepsilon')^{-\frac{\log(n/\varepsilon')}{\varepsilon'}} v(i) \\ &\leq \sum_{i \in \bigcup_j W_j} (1 + \varepsilon')^{-\frac{\log(n/\varepsilon')}{\log(1+\varepsilon')}} \max_{k \in E} v(k) \\ &\leq \sum_{i \in E} (1 + \varepsilon')^{-\log_{1+\varepsilon'}(n/\varepsilon')} v(S^*) \\ &= \sum_{i \in E} \frac{\varepsilon}{6n} v(S^*) = \frac{\varepsilon}{6} v(S^*). \end{aligned}$$

For the second source of loss in approximation, recall that the marginals only decrease due to submodularity. So, we know that if some item j is considered during iteration j (following the renaming of Lemma 3), then $(1 + \varepsilon')v(j | S_j)/c_j \geq \arg \max_{k \in F} v(k | S_j)/c_k$. The only difference this makes (compared to the proof of Theorem 2) is that in the last inequality of (1) we have an extra factor of $1 + \varepsilon'$.

Combining the above, we get the following analog of Lemma 3:

$$v(S \cup S^*) \leq v(S) + v(\hat{i}) + \frac{\varepsilon}{6} v(S^*) + \sum_{j \in E} \left(1 + \frac{\varepsilon}{6}\right) c(Q_j) \frac{v(j | S_j)}{c_j},$$

which carries over to inequality (3), while Lemmata 4 and 5 are not affected at all. It is then a matter of simple calculations to see that for $p = \sqrt{2} - 1$, we still get $v(S^*) \leq (3 + 2\sqrt{2} + \varepsilon) \max\{\mathbb{E}[v(S)], v(i^*)\}$. The final passage is as before. \square

Large instance scenario. Additionally, our analysis implies that SAMPLEGREEDY performs significantly better in the *large instance* scenario, i.e., when the value of the optimal solution is much larger than the value of any single element. While it is not expected to have exact knowledge of the factor δ in the following proposition, often some estimate is accessible. Especially for massive instances, it is reasonable to assume that δ is bounded by a very small constant.

Theorem 4. *If $\max_{i \in E} v(i) \leq \delta \cdot \text{OPT}$ for $\delta \in (0, 1/2)$, then SAMPLEGREEDY with $p = \frac{1-\delta}{2}$ is a $(4 + \varepsilon_\delta)$ -approximation algorithm, where $\varepsilon_\delta = \frac{4\delta(2-\delta)}{(1-\delta)^2}$.*

Proof. Starting from inequality (3) and exploiting the large instance property, we get:

$$(1-p)v(S^*) \leq \max \left\{ 2, \frac{1}{p} \right\} \cdot \mathbb{E}[v(S)] + v(i^*) \leq \max \left\{ 2, \frac{1}{p} \right\} \cdot \mathbb{E}[v(S)] + \delta \cdot v(S^*).$$

Rearranging the terms and assuming $p + \delta < 1$, we have:

$$v(S^*) \leq \frac{\max \left\{ 2, \frac{1}{p} \right\}}{1-p-\delta} \cdot \mathbb{E}[v(S)].$$

Optimizing for $p \in (0, 1 - \delta)$ we get the desired statement. □

4. Adaptive Submodular Maximization

In this section we modify SAMPLEGREEDY to achieve a good approximation guarantee in the adaptive framework. Recall that the adaptive valuation function $v(\cdot, \cdot)$ depends on the state of the system which is discovered a bit at a time, in an adaptive fashion. Indeed, SAMPLEGREEDY is compatible with this framework and can be applied nearly as it is. We stick to the interpretation of SAMPLEGREEDY discussed right before Theorem 2. That is, there is no initial sampling phase. Instead, we directly begin to choose greedily with respect to the density (marginal value with respect to the current solution over cost). Each time we are about to pick an element of E , we throw a p -biased coin that determines whether we keep or discard the element.

Here the main difference with the greedy part of SAMPLEGREEDY is that the marginals are to be considered with respect to the *partial realization* relative to the current solution. Moreover, since it is not possible to return the largest between $\max_{i \in E} v(i)$ and the result of the greedy exploration, the choice between these two quantities has to be settled before starting the exploration. Formally, at the beginning of the algorithm a p_0 -biased coin is tossed to decide between the two. The pseudo-code for the resulting algorithm, ADAPTIVEGREEDY, is given below.

Before proving that ADAPTIVEGREEDY works as promised, we need some observations. Let us denote by S the output of a run of our algorithm, and S^* the output of a run of the optimal adaptive strategy. Fix a realization $\omega \in \Omega$. Now, using pointwise submodularity and directly applying Lemma 1 we have

$$\mathbb{E}[v(S \cup S^*, \omega) | \omega] \geq (1-p) \cdot v(S^*, \omega).$$

```

ADAPTIVEGREEDY
1 Let  $r_0 \sim \text{Bernoulli}(p_0)$ 
2 if  $r_0 = 1$  then
3    $i^* \in \arg \max_{j \in E} v(j)$                                 /* best single item in expectation */
4   Observe  $\omega_{i^*}$  and return  $(i^*, \omega_{i^*})$ 
5  $S = \emptyset, R = B$                                         /* greedy solution and remaining knapsack capacity */
6  $F = \{j \in E \mid v(j) > 0\}$                             /* initial set of candidate items */
7 while  $F \neq \emptyset$  do
8   Let  $i \in \arg \max_{j \in F} \frac{v(j \mid (S, \omega_{|S}))}{c_j}$ 
9   Let  $r_i \sim \text{Bernoulli}(p)$                                 /* independent random bit */
10  if  $r_i = 1$  then
11    Observe  $\omega_i$ :  $S = S \cup \{i\}, R = R - c_i$ 
12     $E = E \setminus \{i\}, F = \{j \in E \mid v(j \mid (S, \omega_{|S})) > 0 \text{ and } c_k \leq R\}$ 
13 return  $(S, \omega_{|S})$ 

```

Since ω (and therefore, S^*) is fixed, the only randomness is due to the coin flips in our algorithm. We stress that the union between S and S^* has to be intended in the following sense: run our algorithm, and independently, also the optimal one, both for the same realization ω . The previous inequality is true for any ω . So, by the law of total probability, we also have

$$\mathbb{E}[v(S \cup S^*)] \geq (1 - p) \cdot \mathbb{E}[v(S^*)]. \tag{4}$$

For the next observation, assume our algorithm has picked (and therefore observed) exactly set S . That is, we know only $\omega_{|S}$. We number all items $j \in E$ with positive marginal value with respect to $(S, \omega_{|S})$ by decreasing ratio $v(i \mid (S, \omega_{|S}))/c_i$, i.e., $j_k = \arg \max_{i \in E \setminus [k-1]} v(i \mid (S, \omega_{|S}))/c_i$ for $k = 1, \dots, |E|$, where $[0] = \emptyset$. Given this ordering, we can refer to *the k best-looking items given S in F* for any specific $k \leq |E|$. In particular, for $k = \min\{i \in \mathbb{N} \mid \sum_{\ell=1}^i c_\ell \geq B\}$, we may apply Lemma 2 to get

$$\sum_{i=1}^k v(i \mid (S, \omega_{|S})) \geq \mathbb{E} \left[\sum_{j \in S^*} v(j \mid (S, \omega_{|S})) \mid \omega_{|S} \right] \geq v(S^* \mid (S, \omega_{|S})). \tag{5}$$

For the sake of the future analysis the last element k is considered fractionally, so that $\sum_{\ell=1}^k c_\ell = B$. Note that it could be the case that k is not well defined, as there may not be enough elements with positive marginal to fill the knapsack. If that is the case, just consider k to be the number of elements with positive marginals. In the future, we refer to this (possibly fraction) set of *best-looking items given S fitting into the budget* as D .

The point of inequality (5) is that, given $(S, \omega_{|S})$ and the coin tosses realized so far, the set D is deterministic (as is its cardinality k), while S^* is not, because it corresponds to the set selected by the best adaptive policy. Moreover, in the middle term notice that the conditioning influences the valuation, but *not* the policy, since we are assuming to run it obliviously. This is fundamental for the analysis.

Since this holds for any set S , we can again generalize to the expectation over all possible runs of the algorithm to all the terms in inequality (5), therefore obtaining

$$\mathbb{E} \left[\sum_{i=1}^k v(i | (S, \omega_{|S})) \right] \geq \mathbb{E} [v(S^* | (S, \omega_{|S}))]. \tag{6}$$

We remark that k above is a random variable which depends on S . We use these observations to prove the ratio of our algorithm.

Theorem 5. *For $p_0 = 1/3$ and $p = 1/6$, ADAPTIVEGREEDY yields a 9-approximation of OPT_Ω , while its lazy version achieves a $(9 + \varepsilon)$ -approximation using $O(n\varepsilon^{-1} \log(n/\varepsilon))$ value oracle calls.*

Moreover, when $\max_{i \in E} v(i) \leq \delta \cdot \text{OPT}_\Omega$ for $\delta \in (0, 1/2)$, then for $p_0 = 0$ and $p = (\sqrt{3 - 2\delta} - 1)/2$, ADAPTIVEGREEDY yields a $(4 + 2\sqrt{3} + \varepsilon'_\delta)$ -approximation, where $\varepsilon'_\delta \approx \frac{6\delta(2-\delta)}{(1-\delta)^2}$.

Proof. For any run of the algorithm, i.e. a fixed set S , the corresponding partial realization $\omega_{|S}$ and the coin flips observed, define for convenience the set C as those items in D that have been considered during the algorithm and then not added to S because of the coin flips. Define $U = D \setminus C$. Additionally, define C' to be the set of all items that are considered, but not chosen during the run of our algorithm which have positive expected marginal contribution to S . C captures the best-looking items given S fitting into the budget that we missed due to coin tosses. We have that $C \subseteq C'$, in fact C' contains all the elements with positive marginal value with respect to S that were discarded by the coin tosses, while C contains only the best of those.

We can then split the left-hand side term of inequality (6) into two parts: the sum over C (upper bounded by the sum over C'), and the sum over U . Now we control separately these terms using linear combinations of $v(S)$ and $v(i^*)$.

Lemma 6. $\mathbb{E}[v(S)] \geq p \cdot \mathbb{E} [\sum_{i \in C} v(i | (S, \omega_{|S}))].$

Proof. Since $C \subseteq C'$ and C' contains all considered elements with nonnegative expected contribution to S , it is sufficient to show $\mathbb{E}[v(S)] \geq p \cdot \mathbb{E} [\sum_{i \in C'} v(i | (S, \omega_{|S}))].$

We proceed as in Lemma 4. Let's consider for each $i \in E$ all the events \mathcal{E}_i capturing the story of a run of the algorithm up to the point element i is considered (all the history if it is never considered).

Let G_i be the marginal contribution of element i to the solution set S . If \mathcal{E}_i corresponds to a story in which element i is not considered, then it does not contribute—neither in the left, nor in the right-hand side of the inequality we are trying to prove. Else, let (S_i, ω_i) be the partial solution when it is indeed considered:

$$\mathbb{E}[G_i | \mathcal{E}_i] = p \cdot v(i | (S_i, \omega_i)) \geq p \cdot \mathbb{E}[v(i | (S, \omega_{|S})) | \mathcal{E}_i].$$

The statement follows from the law of total probability with respect to \mathcal{E}_i , and pointwise submodularity of v . □

Lemma 7. $\mathbb{E}[v(S)] + v(i^*) \geq \mathbb{E} [\sum_{i \in U} v(i | (S, \omega_{|S}))].$

Proof. Now let us turn towards the items U that were not considered by the algorithm. The intuition behind the claim is that if they were not considered then they were not good enough, in expectation, to compare with S . The proof, though, has to deal with some probabilistic subtleties.

Let's start fixing a story of the algorithm, i.e., the coin tosses and (S, ω_S) , $S = s_1, s_2, \dots, s_T$, numbered according to their insertion in S , i.e., s_j is the j^{th} element to be added to S .

There are two cases. If during the whole algorithm the elements in U have ratio $\frac{v(i | (S_j, \omega_{|S_j}))}{c_i}$ smaller than that of the item which was instead considered, then one can easily argue, by adaptive submodularity, that:

$$\begin{aligned} \sum_{i \in U} v(i | (S, \omega_{|S})) &\leq \sum_{t=1}^T v(s_t | (S_t, \omega_t)) + v(u_1 | (S, \omega_{|S})) \\ &\leq \sum_{t=1}^T v(s_t | (S_t, \omega_t)) + v(u_1) \leq \sum_{t=1}^T v(s_t | (S_t, \omega_t)) + v(i^*), \end{aligned}$$

where $S_t = (s_1, \dots, s_{t-1})$ and ω_t is the restriction of $\omega_{|S}$ to S_t . Note that the last element u_1 is added to account for the unspent budget by the solution. The first inequality holds because our solution fills all the budget (up to *at most* one item) with densities which are better than all the $v(i | (S, \omega_{|S}))$. We claim that the above inequality holds also in the case in which there is an element in U whose marginal over cost is greater than that of some in S . Such an element can exist because of the budget constraint: during the algorithm it had better marginal over cost, but was discarded because there was not enough room for it. We observe there can exist at most one such element, due to the budget constraint and because its value is upper bounded by u_1 , so the above formula still holds.

Once we know that, we can conclude the proof applying the law of total probability:

$$\mathbb{E} \left[\sum_{i \in U} v(i | (S, \omega_{|S})) \right] \leq \mathbb{E} [v(S)] + v(i^*). \quad \square$$

Combining the two Lemmata we get:

$$\begin{aligned} (1 + \frac{1}{p}) \mathbb{E} [v(S)] + \mathbb{E} [v(i^*)] &\geq \mathbb{E} \left[\sum_{i \in U} v(i | (S, \omega_{|S})) \right] + \mathbb{E} \left[\sum_{i \in C} v(i | (S, \omega_{|S})) \right] = \\ &= \mathbb{E} \left[\sum_{i \in U \cup C} v(i | (S, \omega_{|S})) \right]. \end{aligned}$$

Inequality (6) implies

$$(1 + \frac{1}{p}) \mathbb{E} [v(S)] + \mathbb{E} [v(i^*)] \geq \mathbb{E} \left[\sum_{i \in U \cup C} v(i | (S, \omega_{|S})) \right] \geq \mathbb{E} [v(S^* | (S, \omega_{|S}))].$$

Also, with some rewriting and inequality (4):

$$\mathbb{E} [v(S^* | (S, \omega_{|S}))] = \mathbb{E} [\mathbb{E} [v(S^* \cup S, \omega) - v(S, \omega) | \omega_{|S}]]$$

$$\begin{aligned} &= \mathbb{E}[v(S^* \cup S)] - \mathbb{E}[v(S)] \\ &\geq (1 - p) \cdot \mathbb{E}[v(S^*)] - \mathbb{E}[v(S)] . \end{aligned}$$

Altogether, denoting as OPT the $\mathbb{E}[v(S^*)]$, we get:

$$(2p + 1)\mathbb{E}[v(S)] + p\mathbb{E}[v(i^*)] \geq p(1 - p)\text{OPT} . \tag{7}$$

Let's call ALG the expected value of the solution output by the algorithm. Since the algorithm chooses with a coin flip either the best expected single item or S , we have that

$$\text{ALG} = (1 - p_0)v(S) + p_0v(i^*) .$$

Picking $p_0 = \frac{p}{3p+1}$,

$$\text{ALG} = \frac{2p + 1}{3p + 1}\mathbb{E}[v(S)] + \frac{p}{3p + 1}\mathbb{E}[v(i^*)] \geq \frac{p(1 - p)}{3p + 1}\text{OPT} .$$

The right-hand side is minimized for $p = \frac{1}{3}$, concluding the proof of the first part of the statement.

The lazy version of ADAPTIVEGREEDY is analogous to the non-adaptive setting, both for the orithm and the analysis, so we omit repeating the proof.

In order to prove the last part of the statement, we start from inequality (7) and apply the large instance property:

$$(1 - p)\mathbb{E}[v(S^*)] \leq \left(2 + \frac{1}{p}\right) \mathbb{E}[v(S)] + \mathbb{E}[v(i^*)] \leq \left(2 + \frac{1}{p}\right) \mathbb{E}[v(S)] + \delta \cdot \mathbb{E}[v(S^*)] .$$

Rearranging terms and assuming $p + \delta < 1$ we have that:

$$\mathbb{E}[v(S^*)] \leq \frac{\left(2 + \frac{1}{p}\right)}{1 - p - \delta} \cdot \mathbb{E}[v(S)] .$$

Optimizing for $p \in (0, 1 - \delta)$, we get the claimed result. Specifically, for $p = (\sqrt{3 - 2\delta} - 1)/2$ the approximation factor is $(4 + 2\sqrt{3} + \varepsilon_\delta)$, with

$$\varepsilon_\delta = 2 \left(\frac{\sqrt{3 - 2\delta} + 1}{(1 - \delta)^2} + \frac{1}{1 - \delta} - \sqrt{3} - 2 \right) \approx \frac{6\delta(2 - \delta)}{(1 - \delta)^2} . \quad \square$$

5. Experiments

Out of the numerous applications of submodular maximization subject to a knapsack constraint, we evaluate SAMPLEGREEDY and ADAPTIVEGREEDY on two selected examples, using real and synthetic graph topologies. Variants of these have been studied in a similar context; see Mirzasoleiman et al. (2016).

As our algorithms are randomized, but extremely fast, we use the best output out of 5 iterations. All graphs presenting the experimental results contain error bars, indicating the standard deviation between the different runs of the experiments. This is usually insignificant due to the concentrating effect of the large size of the instances, despite the randomly

initialized weights and inherent randomness of the algorithms used. For all algorithms involved, we use *lazy evaluations* with $\varepsilon = 0.01$. A delicate point is tuning the probabilities of acceptance p (line 9 of ADAPTIVEGREEDY) for improved performance. While the choices of p in Theorems 2 and 5 minimize our analysis of the theoretical worst-case approximation, there are two factors that suggest a value much closer to 1 works best in practice: the small value of any singleton solution, and the much better guarantee of Lemma 5 for most widely used non-monotone submodular objectives. We do not micro-optimize for p but rather choose uniformly at randomly from $[0.9, 1]$.

Video Recommendation: Suppose we have a large collection E of videos from various categories (represented as possibly intersecting subsets $C_1, \dots, C_k \subseteq E$) and we want to design a recommendation system. When a user inputs a subset of categories and a target total length B , the system should return a set of videos from the selected categories of total duration at most B that maximizes an appropriate objective function. (Of course, instead of time here, we could use costs and a budget constraint.) Each video has a rating and there is some measure of similarity between any two videos. We use a weighted graph on E to model the latter: each edge $\{i, j\}$ between two videos i and j has a weight $w_{ij} \in [0, 1]$ capturing the percentage of their similarity. To pave the way for our $v(\cdot)$, we start from the auxiliary objective $f(S) = \sum_{i \in S} \sum_{j \in E} w_{ij} - \lambda \sum_{i \in S} \sum_{j \in S} w_{ij}$, for some $\lambda \geq 1$ (Lin and Bilmes, 2010; Mirzasoleiman et al., 2016). This is a *maximal marginal relevance* inspired objective (Carbinell and Goldstein, 2017) that rewards coverage, while penalizing similarity. For $\lambda = 1$, internal similarities are irrelevant and f becomes a cut function. However, one can penalize similarities even more severely as f is submodular for $\lambda \geq 1$ (e.g., Lin and Bilmes (2010) use $\lambda = 5$).

In order to mimic the effect of a partition matroid constraint, i.e., the avoidance of many videos from the same category, we may use two parameters $\lambda \geq 1, \mu \geq 0$. While λ is as above, μ puts extra weight on similarities between videos that belong to the same category. That leads to a more general auxiliary objective $g(S) = \sum_{i \in S} \sum_{j \in E} w_{ij} - \sum_{i \in S} \sum_{j \in S} (\lambda + \chi_{ij} \mu) w_{ij}$, where χ_{ij} is equal to 1 if there exists ℓ such that $i, j \in C_\ell$ and 0 otherwise. To interpolate between choosing highly rated videos and videos that represent well the whole collection, here we use the submodular function $v(S) = \alpha \sum_{i \in S} \rho_i + \beta g(S)$ for $\alpha, \beta \geq 0$, where ρ_i is the rating of video i . We use $\lambda = 3, \mu = 7$ and set the parameters α, β so that the two terms are of comparable size.

We evaluate SAMPLEGREEDY on an instance based on the latest version of the MovieLens dataset (Harper and Konstan, 2015), which includes 62000 movies, 13816 of which have *both* user-generated tags *and* ratings.

We calculate the weights w_{ij} using these tags while the costs are drawn independently from $U(0, 1)$. These tag vectors are *not* normalized and have no additional structure, other than each coordinate being restricted to $[0, 1]$. We define the similarity w_{ij} between two movies i, j as:

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

In other words, it is the L2 norm of the *coordinate-wise minimum* of t^i and t^j . This metric was chosen so that if *both* movies have a high value in some tag, this counts as a much

stronger similarity than one having a high value and the other a low one. For example, if we consider an inner product metric, any movie with all tags set to 1 would be as similar as possible to all other movies, even though it would include many tags that would be missing from the others. In particular, any movie would appear more similar to the all 1 movie than to itself! Choosing the minimum of both tags avoids this issue. Another possibility would be to normalize each tag vector before taking the inner product, to obtain the *cosine similarity*. Although this alleviates some of the issues, there is some information loss as one movie could meaningfully have higher scores in all tags than another one; tags are not mutually exclusive. Ultimately any sensible metric has advantages and disadvantages and the exact choice has little bearing on our results. The similarity scores are then divided by their maximum as a final normalization step.

We compare against the FANTOM algorithm of Mirzasoleiman et al. (2016) as it is the only other algorithm with a provable approximation guarantee that runs in reasonable time. Continuous greedy approaches (Feldman et al., 2011) or the repeated greedy of Gupta et al. (2010) are prohibitively slow. SAMPLEGREEDY consistently performs better than FANTOM for a wide range of budgets (Fig. 1a). Plotting the number of function evaluations against the budget, SAMPLEGREEDY is much faster (Fig. 1d) despite the fact that it is run 5 times!

The experiment was repeated 5 times. The budget is represented as a fraction of the total cost starting at 1/100 and geometrically increasing to 1/10 in 10 steps. The total computation time was around 3 hours.

Remark 1. The running time of FANTOM for fixed ε is $O(nr \log n)$, where r is the cardinality of the largest feasible solution. For a knapsack constraint this translates to $O(n^2 \log n)$. To be as fair as possible, we implemented FANTOM using lazy evaluations, which improves the number of evaluations of the objective function to $O(n \log^2 n)$ and is indeed much faster in practice, for the knapsack sizes we consider. Even so, our SAMPLEGREEDY is faster by a factor of $\Omega(\log n)$ which, including the improvement in the constants involved, still makes a huge difference. Note that in both Figs. 1d and 1e one can discern the superlinear increase of the function evaluations for FANTOM but not for SAMPLEGREEDY.

Remark 2. The role of our parameter p is different from that of the ε precision parameter in FANTOM. The role of ε in (Mirzasoleiman et al., 2016) is to control the density of the search space grid. The behaviour of the solution for varying ε is roughly monotone, i.e., the smaller the ε , the finer the grid, generally yielding better solutions and requiring more value queries. In our case, the instance-specific optimal p is unknown and reflects, to some extent and roughly speaking, “how non-monotone the function is”. Picking smaller values for ε in our experiments, would slightly improve the quality of the solution output by FANTOM, but would also increase its run time to the point where it would become impractical to run it for any large instances. Choosing $\varepsilon = 1$, we believe that we give a fair and reasonable comparison.

Influence-and-Exploit Marketing: Consider a seller of a single digital good (i.e., producing extra units of the good comes at no extra cost) and a social network on a set E of potential buyers. Suppose that the buyers influence each other and this is quantified by a weight w_{ij} on each edge $\{i, j\}$ between buyers i and j . Each buyer’s value for the good depends on who owns it within her immediate social circle and how they influence

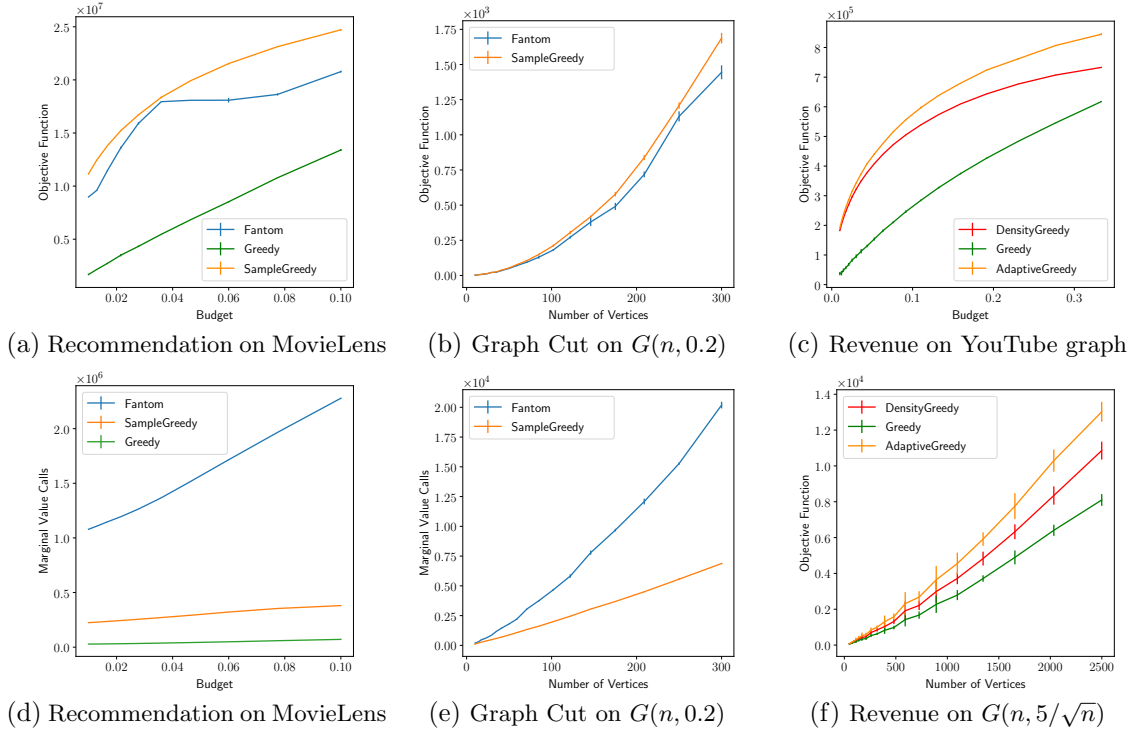


Figure 1: The four plots on the left compare the performance and the number of function evaluations of SAMPLEGREEDY and FANTOM on the video recommendation problem for the MovieLens dataset (a), (d) and on the maximum weighted cut problem on random graphs (b), (e). Since no $\varepsilon \leq 1$ affected the performance of FANTOM noticeably before becoming too computationally expensive, we used $\varepsilon = 1$ to achieve the maximum possible speedup. The plots on the far right illustrate the performance of ADAPTIVEGREEDY (ignoring single item solutions, i.e., $p_0 = 0$) on the influence-and-exploit problem for two distinct topologies: the YouTube graph (c) and random graphs (f). All budgets are shown as fractions of the total cost.

her. A possible revenue-maximizing strategy for the seller is to first give the item for free to a selected set S of influential buyers (influence phase) and then extract revenue by selling to each of the remaining buyers at a price matching their value for the item due to the influential nodes (exploit phase). Here we further assume, similarly to the adaptation of this model by Mirzasoleiman et al. (2016), that each buyer comes with a cost of convincing her to advertise the product to her friends. The seller has a budget B and the set S should be such that $\sum_{i \in S} c_i \leq B$.

We adopt the generalization of the *Concave Graph Model* of Hartline et al. (2008) to non-monotone functions (Babaei et al., 2013). Each buyer $i \in E$ is associated with a non-negative concave function f_i . For any $i \in E$ and any set $S \subset E \setminus \{i\}$ of agents already owning the good, the value of i for it is $v_i(S) = f_i(\sum_{j \in S \cup \{i\}} w_{ij})$. The total potential revenue $v(S) = \sum_{i \in E \setminus S} v_i(S)$ that we aim to maximize is a non-monotone submodular function. Besides the theoretical guarantees for influence-and-exploit marketing in the Bayesian setting (Hartline et al., 2008), there are strong experimental evidence of its performance in practice (Babaei et al., 2013). The problem generalizes naturally to different stochastic versions. We assume that the valuation function of each buyer i is of the form $f_i(x) = a_i \sqrt{x}$

where a_i is drawn independently from a Pareto Type II distribution with $\lambda = 1$, $\alpha = 2$. We only learn the exact value of a buyer when we give the good for free to someone in her neighborhood.

We evaluate ADAPTIVEGREEDY on an instance based on the YouTube graph (Yang and Leskovec, 2015), containing 1,134,890 vertices. The (known) weights are drawn independently from $U(0, 1)$, and the costs are proportional to the sum of the weights of the incident edges. As ADAPTIVEGREEDY is the first adaptive algorithm for the problem, we compare with non-adaptive alternatives like *Greedy*³ and *Density Greedy*⁴ for different values of the budget. ADAPTIVEGREEDY outperforms the alternatives by up to 20% (Fig. 1c).

We observe similar improvements for Erdős-Rényi random graphs of different sizes and edge probability $5/\sqrt{n}$ and a fixed budget of 10% of the total cost (Fig. 1f).

For the YouTube graph, the experiment was repeated 5 times for a budget starting at 1/100 of the total cost and geometrically increasing to 1/3 in 20 steps, leading to a total computation time of 7 hours. For the Erdős-Rényi graph with n vertices and edge probability $5/\sqrt{n}$ it was repeated 10 times, for n starting at 50 and geometrically increasing to 2500 in 20 steps, taking approximately 10 minutes.

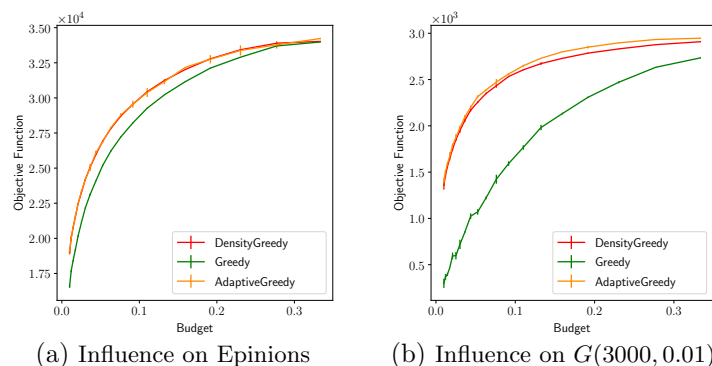


Figure 2: The plots illustrate the performance of ADAPTIVEGREEDY (ignoring single item solutions, i.e., $p_0 = 0$) on the influence maximization problem for two distinct topologies: the Epinions graph (a) and random graph (b). All budgets are shown as fractions of the total cost.

Influence Maximization: This setting is similar to influence-and-exploit marketing, focusing just on maximizing the number of people influenced. Consider a social network where each user could become an ‘influencer’, by posting about the product we are trying to promote. Their friends on the platform would see this and might favourably change their opinion. Similar to Breuer et al. (2020), we assume that each user has an independent probability q of being influenced by any influencer adjacent to them. Formally given the set of influencers S , for the probability that each user i is influenced is:

$$\begin{cases} 0, & \text{for } i \in S \\ 1 - (1 - q)^{N_S(i)}, & \text{for } i \notin S \end{cases} .$$

One difference from Breuer et al. (2020) is that influencers do not count towards maximizing the objective: this makes the instance non-monotone submodular. Moreover, as we did for

3. The simple greedy algorithm that in each step picks the element with the largest marginal value.
 4. The greedy part of Wolsey’s algorithm (Wolsey, 1982).

influence-and-exploit marketing, we consider the *adaptive* version: whenever an influencer is added to the set S , any of their uninfluenced neighbours might indendently become influenced with probability q , updating the expected marginals accordingly.

In Figure 2 we evaluate ADAPTIVEGREEDY on an instance based on the Epinions network graph (Richardson et al., 2003), containing 75,879 nodes and 508,837 edges. We have set $q = 0.2$ and the cost of each node is drawn from $U[0, 1]$. As before, we compare ADAPTIVEGREEDY to Greedy and Density Greedy, repeating the experiment 5 times for budgets ranging from $1/100$ to $1/30$ of the total cost. We observe that ADAPTIVEGREEDY is tied with Density Greedy and performs significantly better than Greedy. The most likely reason for this decreased separation between the different algorithms is that this objective is more benign: especially for small budgets, the non-monotone aspect is muted and the knowledge obtained from adding an element to S is not as useful as it was for influence-and-exploit marketing. We also repeat the same experiment for the $G(3000, 0.01)$ random graph.

Maximum Weighted Cut: Beyond the above applications, we would like to compare SAMPLEGREEDY to FANTOM with respect to both their performance and the number of value oracle calls as n grows. We turn to *weighted cut functions*—one of the most prominent subclasses of non-monotone submodular functions—on dense Erdős–Rényi random graphs with edge probability 0.2. The weights and the costs are drawn independently and uniformly from $[0, 1]$ and the budget is fixed to 15% of the total cost. Again SAMPLEGREEDY consistently performs better than FANTOM, albeit by 5–15% (Fig. 1b). In terms of running time, there is a large difference in favor of SAMPLEGREEDY (even for multiple runs), while the superlinear increase for FANTOM is evident (Fig. 1e). The experiment was repeated 10 times for n starting at 10 and increasing geometrically to 300 in 20 steps, requiring approximately 5 minutes.

Since here SAMPLEGREEDY and FANTOM are quite close to each other in terms of performance and Greedy would lie only slightly below the plot of SAMPLEGREEDY, so to improve the readability of Fig. 1b we have removed Greedy from this comparison.

Remark 3. Based on the theoretical query complexities, one would expect the comparison between FANTOM and SAMPLEGREEDY in Figs. 1d and 1e to be qualitatively similar to $n \log^2 n$ vs. $n \log n$. However, while FANTOM clearly exhibits a superlinear dependence of the query complexity on the input size, SAMPLEGREEDY does not. The reason for this is that, in practice, lazy evaluations often result in much less than $\log n$ evaluations per element. So, what we see in Figs. 1d and 1e is closer to $n \log n$ vs. n .

A roadmap for practitioners. Although we consider our work to be a theoretical paper, SAMPLEGREEDY is the first constant factor approximation algorithm for non-monotone submodular maximization with knapsack constraints that needs only $O(n \log n)$ value queries and *can be applied in practice* wherever Density Greedy can. Moreover, in the considered experiments it outperforms the state of the art, FANTOM. In some instances, the heuristic Density Greedy outputs marginally better solutions than SAMPLEGREEDY, albeit at the expense of offering no theoretical approximation guarantee (as we have argued in the Introduction). From a practical point of view, a reasonable strategy is to start with Density Greedy (SampleGreedy for $p = 1$), and then gradually reduce p , tuning it to the effective non-monotonicity of the instance at hand. Note that as soon as p is a constant factor away from 1, SampleGreedy guarantees a constant factor approximation to the optimal solution.

6. Conclusions

The proposed random greedy method yields a considerable improvement over state-of-the-art algorithms, especially, but not exclusively, regarding the handling of huge instances. With all the subtleties of our work affecting solely our analysis, the algorithm remains strikingly simple and we are confident this will also contribute to its use in practice. Simultaneously, this very simplicity translates into a generality that can be employed to achieve comparably good results for a variety of settings; we demonstrated this in the case of the adaptive submodularity setting.

Specifically, we expect that our approach can be directly utilised to improve the performance and running time of algorithms that now use some variant of the algorithm of Gupta et al. (2010). Such examples include the distributed algorithm of da Ponte Barbosa et al. (2015) and the streaming algorithm of Mirzasoleiman et al. (2018) in the case of a knapsack constraint. We further suspect that the same algorithmic principle can be applied in the presence of incentives. This would largely improve the current state of the art in budget-feasible mechanism design for non-monotone objectives (Chen et al., 2011; Amanatidis et al., 2019).

A different direction would be to try other greedy algorithms for monotone objectives as a starting point. For instance, the 2-approximation algorithm of Yaroslavtsev et al. (2020) could potentially yield a better approximation ratio for the standard algorithmic setting. Unfortunately, it does not seem possible to translate more involved algorithms like that one to the adaptive setting, where one has to commit to all of their past choices.

Finally, a major question here is whether the same high level approach is valid even in the presence of additional combinatorial constraints. In particular, is it possible to achieve similar guarantees as FANTOM for a p -system and multiple knapsack constraints using only $O(n \log n)$ value queries?

Acknowledgments

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

References

- Alaluf, N., Ene, A., Feldman, M., Nguyen, H. L., and Suh, A. (2020). Optimal streaming algorithms for submodular maximization with cardinality constraints. In *Proceedings of the 47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8-11, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 168 of *LIPICs*, pages 6:1–6:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Amanatidis, G., Fusco, F., Lazos, P., Leonardi, S., Marchetti-Spaccamela, A., and Reiffenhäuser, R. (2021). Submodular maximization subject to a knapsack constraint: Combinatorial algorithms with near-optimal adaptive complexity. In Meila, M. and Zhang, T., editors, *Proceedings of the 38th International Conference on Machine Learning, ICML*

- 2021, 18-24 July 2021, Virtual Event, volume 139 of *Proceedings of Machine Learning Research*, pages 231–242. PMLR.
- Amanatidis, G., Fusco, F., Lazos, P., Leonardi, S., and Reiffenhäuser, R. (2020). 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*.
- Amanatidis, G., Kleer, P., and Schäfer, G. (2019). Budget-feasible mechanism design for non-monotone submodular objectives: Offline and online. In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019*, pages 901–919. ACM.
- Babaei, M., Mirzasoleiman, B., Jalili, M., and Safari, M. A. (2013). Revenue maximization in social networks through discounting. *Social Netw. Analys. Mining*, 3(4):1249–1262.
- Badanidiyuru, A., Mirzasoleiman, B., Karbasi, A., and Krause, A. (2014). Streaming submodular maximization: massive data summarization on the fly. In *The 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '14, New York, NY, USA - August 24 - 27, 2014*, pages 671–680. ACM.
- Balkanski, E., Rubinstein, A., and Singer, Y. (2019). An exponential speedup in parallel running time for submodular maximization without loss in approximation. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 283–302. SIAM.
- Balkanski, E. and Singer, Y. (2018). 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.
- Boodaghians, S., Fusco, F., Lazos, P., and Leonardi, S. (2020). Pandora’s box problem with order constraints. In *EC '20: The 21st ACM Conference on Economics and Computation, Virtual Event, Hungary, July 13-17, 2020*, pages 439–458. ACM.
- Bradac, D., Singla, S., and Zuzic, G. (2019). (near) optimal adaptivity gaps for stochastic multi-value probing. In *APPROX-RANDOM*, volume 145 of *LIPICs*, pages 49:1–49:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Breuer, A., Balkanski, E., and Singer, Y. (2020). The FAST algorithm for submodular maximization. In *ICML*, volume 119 of *Proceedings of Machine Learning Research*, pages 1134–1143. PMLR.
- Buchbinder, N., Feldman, M., Naor, J., and Schwartz, R. (2014). 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.
- Carbinell, J. and Goldstein, J. (2017). The use of MMR, diversity-based reranking for reordering documents and producing summaries. *SIGIR Forum*, 51(2):209–210.

- Chekuri, C., Gupta, S., and Quanrud, K. (2015a). Streaming algorithms for submodular function maximization. In *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I*, volume 9134 of *Lecture Notes in Computer Science*, pages 318–330. Springer.
- Chekuri, C., Jayram, T. S., and Vondrák, J. (2015b). On multiplicative weight updates for concave and submodular function maximization. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science, ITCS 2015, Rehovot, Israel, January 11-13, 2015*, pages 201–210. ACM.
- Chekuri, C., Vondrák, J., and Zenklusen, R. (2014). Submodular function maximization via the multilinear relaxation and contention resolution schemes. *SIAM J. Comput.*, 43(6):1831–1879.
- Chen, N., Gravin, N., and Lu, P. (2011). On the approximability of budget feasible mechanisms. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 685–699.
- da Ponte Barbosa, R., Ene, A., Nguyen, H. L., and Ward, J. (2015). The power of randomization: Distributed submodular maximization on massive datasets. In *ICML*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 1236–1244. JMLR.org.
- Das, A. and Kempe, D. (2008). 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.
- Das, A. and Kempe, D. (2018). Approximate submodularity and its applications: Subset selection, sparse approximation and dictionary selection. *Journal of Machine Learning Research*, 19(3):1–34.
- Dütting, P., Fusco, F., Lattanzi, S., Norouzi-Fard, A., and Zadimoghaddam, M. (2022). Deletion robust submodular maximization over matroids. *CoRR*, abs/2201.13128.
- Edmonds, J. (1971). Matroids and the greedy algorithm. *Math. Program.*, 1(1):127–136.
- Ene, A. and Nguyen, H. L. (2019). A nearly-linear time algorithm for submodular maximization with a knapsack constraint. In *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019, Patras, Greece*, volume 132 of *LIPICs*, pages 53:1–53:12. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Ene, A., Nguyen, H. L., and Vladu, A. (2019). 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.
- Fadaei, S., Fazli, M., and Safari, M. (2011). Maximizing non-monotone submodular set functions subject to different constraints: Combined algorithms. *Oper. Res. Lett.*, 39(6):447–451.

- Fahrbach, M., Mirrokni, V. S., and Zadimoghaddam, M. (2019a). 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.
- Fahrbach, M., Mirrokni, V. S., and Zadimoghaddam, M. (2019b). 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.
- Feige, U. (1998). A threshold of $\ln(n)$ for approximating set cover. *J. ACM*, 45(4):634–652.
- Feige, U., Mirrokni, V. S., and Vondrák, J. (2011). Maximizing non-monotone submodular functions. *SIAM J. Comput.*, 40(4):1133–1153.
- Feldman, M., Harshaw, C., and Karbasi, A. (2017). Greed is good: Near-optimal submodular maximization via greedy optimization. In *Proceedings of the 30th Conference on Learning Theory, COLT 2017, Amsterdam, The Netherlands, 7-10 July 2017*, volume 65 of *Proceedings of Machine Learning Research*, pages 758–784. PMLR.
- Feldman, M., Karbasi, A., and Kazemi, E. (2018). Do less, get more: Streaming submodular maximization with subsampling. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems, NeurIPS 2018, 3-8 December 2018, Montréal, Canada*, pages 730–740.
- Feldman, M., Naor, J., and Schwartz, R. (2011). 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.
- Feldman, M. and Zenklusen, R. (2018). The submodular secretary problem goes linear. *SIAM J. Comput.*, 47(2):330–366.
- Golovin, D. and Krause, A. (2011). Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *J. Artif. Intell. Res.*, 42:427–486.
- Gotovos, A., Karbasi, A., and Krause, A. (2015). Non-monotone adaptive submodular maximization. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 1996–2003. AAAI Press.
- Gupta, A., Nagarajan, V., and Singla, S. (2017). Adaptivity gaps for stochastic probing: Submodular and XOS functions. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 1688–1702. SIAM.
- Gupta, A., Roth, A., Schoenebeck, G., and Talwar, K. (2010). 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 *LNCS*, pages 246–257. Springer.
- Harper, F. M. and Konstan, J. A. (2015). The MovieLens datasets: History and context. *ACM Trans. Interact. Intell. Syst.*, 5(4).
- Hartline, J. D., Mirrokni, V. S., and Sundararajan, M. (2008). 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.
- Kempe, D., Kleinberg, J. M., and Tardos, É. (2015). Maximizing the spread of influence through a social network. *Theory of Computing*, 11:105–147.
- Khanna, R., Elenberg, E. R., Dimakis, A. G., Negahban, S. N., and Ghosh, J. (2017). 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.
- Khuller, S., Moss, A., and Naor, J. S. (1999). The budgeted maximum coverage problem. *Information processing letters*, 70(1):39–45.
- Krause, A., Singh, A. P., and Guestrin, C. (2008). Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9:235–284.
- Kuhnle, A. (2021). Nearly linear-time, parallelizable algorithms for non-monotone submodular maximization. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Virtual Event, February 2-9, 2021*, pages 8200–8208. AAAI Press.
- Kulik, A., Shachnai, H., and Tamir, T. (2013). Approximations for monotone and nonmonotone submodular maximization with knapsack constraints. *Math. Oper. Res.*, 38(4):729–739.
- Lee, J., Mirrokni, V. S., Nagarajan, V., and Sviridenko, M. (2010). Maximizing nonmonotone submodular functions under matroid or knapsack constraints. *SIAM J. Discrete Math.*, 23(4):2053–2078.
- Lin, H. and Bilmes, J. A. (2010). Multi-document summarization via budgeted maximization of submodular functions. In *Human Language Technologies: Conference of the North American Chapter of the Association of Computational Linguistics, Proceedings, June 2-4, 2010, Los Angeles, California, USA*, pages 912–920. The Association for Computational Linguistics.
- Minoux, M. (1978). Accelerated greedy algorithms for maximizing submodular set functions. In *Optimization Techniques*, pages 234–243, Berlin, Heidelberg. Springer Berlin Heidelberg.

- Mirzasoleiman, B., Badanidiyuru, A., and Karbasi, A. (2016). 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*, pages 1358–1367. JMLR.org.
- Mirzasoleiman, B., Jegelka, S., and Krause, A. (2018). Streaming non-monotone submodular maximization: Personalized video summarization on the fly. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 1379–1386. AAAI Press.
- Mirzasoleiman, B., Karbasi, A., Sarkar, R., and Krause, A. (2013). 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, NIPS 2013, December 5-8, 2013, Lake Tahoe, Nevada, United States*, pages 2049–2057.
- Mitrovic, M., Kazemi, E., Feldman, M., Krause, A., and Karbasi, A. (2019). Adaptive sequence submodularity. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pages 5353–5364.
- Nemhauser, G. L., Wolsey, L. A., and Fisher, M. L. (1978). An analysis of approximations for maximizing submodular set functions - I. *Math. Program.*, 14(1):265–294.
- Richardson, M., Agrawal, R., and Domingos, P. (2003). Trust management for the semantic web. In *International semantic Web conference*, pages 351–368. Springer.
- Shperberg, S. S. and Shimony, S. E. (2017). Some properties of batch value of information in the selection problem. *J. Artif. Intell. Res.*, 58:777–796.
- Sipos, R., Swaminathan, A., Shivaswamy, P., and Joachims, T. (2012). Temporal corpus summarization using submodular word coverage. In *21st ACM International Conference on Information and Knowledge Management, CIKM'12, Maui, HI, USA, October 29 - November 02, 2012*, pages 754–763. ACM.
- Sviridenko, M. (2004). A note on maximizing a submodular set function subject to a knapsack constraint. *Oper. Res. Lett.*, 32(1):41–43.
- Tang, S. (2021a). Beyond pointwise submodularity: Non-monotone adaptive submodular maximization in linear time. *Theor. Comput. Sci.*, 850:249–261.
- Tang, S. (2021b). Beyond pointwise submodularity: Non-monotone adaptive submodular maximization subject to a knapsack constraint. *CoRR*, abs/2104.04853.
- Tschiatschek, S., Iyer, R. K., Wei, H., and Bilmes, J. A. (2014). 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, NIPS 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 1413–1421.

- Wolsey, L. A. (1982). Maximising real-valued submodular functions: Primal and dual heuristics for location problems. *Math. Oper. Res.*, 7(3):410–425.
- Yang, J. and Leskovec, J. (2015). Defining and evaluating network communities based on ground-truth. *Knowledge and Information Systems*, 42(1):181–213.
- Yaroslavtsev, G., Zhou, S., and Avdiukhin, D. (2020). “Bring your own greedy” + Max: Near-optimal $1/2$ -approximations for submodular knapsack. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics, AISTATS 2020, 26-28 August 2020, Online*, volume 108 of *Proceedings of Machine Learning Research*, pages 3263–3274. PMLR.