Constrained Submodular Maximization via a Non-symmetric Technique

Niv Buchbinder[*]   Moran Feldman[†]

November 11, 2016

Abstract

The study of combinatorial optimization problems with a submodular objective has attracted much attention in recent years. Such problems are important in both theory and practice because their objective functions are very general. Obtaining further improvements for many submodular maximization problems boils down to finding better algorithms for optimizing a relaxation of them known as the multilinear extension.

In this work we present an algorithm for optimizing the multilinear relaxation whose guarantee improves over the guarantee of the best previous algorithm (which was given by Ene and Nguyen (2016)). Moreover, our algorithm is based on a new technique which is, arguably, simpler and more natural for the problem at hand. In a nutshell, previous algorithms for this problem rely on symmetry properties which are natural only in the absence of a constraint. Our technique avoids the need to resort to such properties, and thus, seems to be a better fit for constrained problems.

[*]Department of Statistics and Operations Research, School of Mathematical Sciences, Tel Aviv university, Israel. Email: `niv.buchbinder@gmail.com`.

[†]Depart. of Mathematics and Computer Science, The Open University of Israel. Email: `moranfe@openu.ac.il`.

1 Introduction

The study of combinatorial optimization problems with a submodular objective has attracted much attention in recent years. Such problems are important in both theory and practice because their objective functions are very general—submodular functions generalize, for example, cuts functions of graphs and directed graphs, the mutual information function, matroid weighted rank functions and log-determinants. More specifically, from a theoretical perspective, many well-known problems in combinatorial optimization are in fact submodular maximization problems, including: Max-Cut [30, 33, 38, 40, 56], Max-DiCut [20, 30, 31], Generalized Assignment [10, 14, 22, 27], Max-k-Coverage [19, 41], Max-Bisection [3, 28] and Facility Location [1, 16, 17]. From a practical perspective, submodular maximization problems have found uses in social networks [32, 39], vision [5, 36], machine learning [43, 44, 45, 49, 50] and many other areas (the reader is referred, for example, to a comprehensive survey by Bach [4]).

The techniques used by approximation algorithms for submodular maximization problems usually fall into one of two main approaches. The first approach is combinatorial in nature, and is mostly based on local search techniques and greedy rules. This approach has been used as early as the late 70's for maximizing a monotone submodular function subject to a matroid constraint (some of these works apply only to specific types of matroids) [15, 26, 34, 35, 37, 42, 53, 54]. Later works used this approach to handle also problems with non-monotone submodular objective functions and different constraints [6, 21, 25, 47, 48], yielding in some cases optimal algorithms [6, 55]. However, algorithms based on this approach tend to be highly tailored for the specific structure of the problem at hand, making extensions quite difficult.

The second approach used by approximation algorithms for submodular maximization problems overcomes the above obstacle. This approach resembles a common paradigm for designing approximation algorithms and involves two steps. In the first step a fractional solution is found for a relaxation of the problem, known as the *multilinear relaxation*. In the second step the fractional solution is rounded to obtain an integral one while incurring a bounded loss in the objective. This approach has been used to obtain improved approximations for many problems [8, 11, 12, 24, 46].

Various techniques have been developed for rounding the fractional solution. These techniques tend to be quite flexible, and usually can extend to many related problem. In particular, the Contention Resolution Schemes framework of [12] yields a rounding procedure for every constraint which can be presented as the intersection of a few basic constraints such as knapsack constraints, matroid constraints and matching constraints. Given this wealth of rounding procedures, obtaining further improvements for many important submodular maximization problems (such as maximizing a submodular function subject to a matroid or knapsack constraint) boils down to obtaining improved algorithms for finding a good fractional solution, *i.e.*, optimizing the multilinear relaxation.

1.1 Maximizing the Multilinear Relaxation

At this point we would like to present some terms more formally. A submodular function is a set function $f\colon 2^{\mathcal{N}} \to \mathbb{R}$ obeying $f(A) + f(B) \geq f(A \cup B) + f(A \cap B)$ for any sets $A, B \subseteq \mathcal{N}$. A submodular maximization problem is the problem of finding a set $S \subseteq \mathcal{N}$ maximizing f subject to some constraint. Formally, let $\mathcal{I}$ be the set of subsets of $\mathcal{N}$ obeying the constraint. Then, we are interested in the following problem.

$$\begin{aligned} \max \quad & f(A) \\ \text{s.t.} \quad & A \in \mathcal{I} \subseteq 2^{\mathcal{N}} \end{aligned}$$

A relaxation of the above problem replaces $\mathcal{I}$ with a polytope $P \subseteq [0,1]^{\mathcal{N}}$ containing the

characteristic vectors of all the sets of $\mathcal{I}$. In addition, a relaxation must replace the function f with an extension function $F\colon [0,1]^{\mathcal{N}} \to \mathbb{R}$. Thus, a relaxation is a fractional problem of the following format.

$$\begin{aligned} \max \quad & F(x) \\ \text{s.t.} \quad & x \in P \subseteq [0,1]^{\mathcal{N}} \end{aligned}$$

Defining the "right" extension function, F, for the relaxation is a challenge, as, unlike the linear case, there is no single natural candidate. The objective that turned out to be useful, and is, thus, used by multilinear relaxation is known as the *multilinear extension* (first introduced by [8]). The value $F(x)$ of this extension for any vector $x \in [0,1]^{\mathcal{N}}$ is defined as the expected value of f over a random subset $\mathtt{R}(x) \subseteq \mathcal{N}$ containing every element $u \in \mathcal{N}$ independently with probability x_u. Formally, for every $x \in [0,1]^{\mathcal{N}}$,

$$F(x) = \mathbb{E}[\mathtt{R}(x)] = \sum_{S \subseteq \mathcal{N}} f(S) \prod_{u \in S} x_u \prod_{u \notin S} (1 - x_u) \ .$$

The first algorithm for optimizing the multilinear relaxation was the Continuous Greedy algorithm designed by Calinescu et al. [8]. When the submodular function f is non-negative and monotone[1] and P is solvable[2] this algorithm finds a vector $x \in P$ such that $\mathbb{E}[F(x)] \geq (1 - 1/e - o(1)) \cdot f(OPT)$ (where OPT is the set maximizing f among all sets whose characteristic vectors belongs to P). Interestingly, the guarantee of Continuous Greedy is optimal for monotone functions even when P is a simple cardinality constraint [8, 53].

Optimizing the multilinear relaxation when f is not necessarily monotone proved to be a more challenging task. Initially, several algorithms for specific polytopes were suggested [29, 47, 57]. Later on, improved general algorithms were designed that work whenever f is non-negative and P is down-closed[3] and solvable [13, 24]. Designing algorithms that work in this general setting is highly important as many natural constraints fall into this framework. Moreover, the restriction of the algorithms to down-closed polytopes is unavoidable as Vondrák [57] proved that no algorithm can produce a vector $x \in P$ obeying $\mathbb{E}[F(x)] \geq c \cdot f(OPT)$ for any constant $c > 0$ when P is solvable but not down-closed.

Up until recently, the best algorithm for this general setting was called Measured Continuous Greedy [24]. It guaranteed to produce a vector $x \in P$ obeying $\mathbb{E}[F(x)] \geq (1/e - o(1)) \cdot f(OPT) \approx 0.367 \cdot f(OPT)$ [24]. The natural feel of the guarantee of Measured Continuous Greedy and the fact that it was not improved for a few years made some people suspect that it is optimal. Recently, an evidence against this conjecture was given by [7], which described an algorithm for the special case of a cardinality constraint with an improved approximation guarantee of 0.371. Even more recently, Ene and Nguyen [18] shuttered the conjecture completely. By extending the technique used by [7], they showed that one can get an approximation guarantee 0.372 for every down-closed and solvable polytope P. On the inapproximability side, Oveis Gharan and Vondrák [29] proved that no algorithm can achieve approximation better than 0.478 even when P is the matroid polytope of a partition matroid. Closing the gap between the best algorithm and inapproximability result for this fundamental problem remains an important open problem.

---

[1] A set function $f\colon 2^{\mathcal{N}} \to \mathbb{R}$ is monotone if $f(A) \leq f(B)$ for every $A \subseteq B \subseteq \mathcal{N}$.

[2] A polytope is solvable if one can optimize linear functions over it.

[3] A polytope $P \subseteq [0,1]^{\mathcal{N}}$ is down-closed if $y \in P$ implies that every vector $x \in [0,1]^{\mathcal{N}}$ which is coordinate-wise upper bounded by y must belong to P as well.

1.2 Our Contribution

Our main contribution is an algorithm with an improved guarantee for maximizing the multilinear relaxation.

Theorem 1.1. *There exists a polynomial time algorithm that given a non-negative submodular function $f: 2^{\mathcal{N}} \to \mathbb{R}_{\geq 0}$ and a solvable down-closed polytope $P \subseteq [0,1]^{\mathcal{N}}$ finds a vector $x \in P$ obeying $F(x) \geq 0.385 \cdot f(OPT)$, where $OPT = \arg\max\{f(S) : 1_S \in P\}$ and F is the multilinear extension of f.*

Admittedly, the improvement in the guarantee obtained by our algorithm compared to the 0.372 guarantee of [18] is relatively small. However, the technique underlying our algorithm is very different, and, arguably, much cleaner, than the technique underlying the previous results improving over the natural guarantee of $1/e$ [7, 18]. Moreover, we believe our technique is more natural for the problem at hand, and thus, is likely to yield further improvements in the future. In the rest of this section we explain the intuition on which we base this belief.

The results of [7, 18] are based on the observation that the guarantee of Measured Continuous Greedy improves when the algorithm manages to increase all the coordinates of its solution at a slow rate. Based on this observation, [7, 18] run an instance of Measured Continuous Greedy (or a discretized version of it), and force it to raise the coordinates slowly. If this extra restriction does not affect the behavior of the algorithm significantly, then it produces a solution with an improved guarantee. Otherwise, [7, 18] argue that the point in which the extra restriction affect the behavior of Measured Continuous Greedy reveals a vector $x \in P$ which contains a significant fraction of OPT. Once x is available, one can use the technique of unconstrained submodular maximization, described by [6], that has higher approximation guarantee of $1/2 > 1/e$, to extract from x a vector $0 \leq y \leq x$ of large value. The down-closeness of P guarantees that y belongs to P as well.

Unfortunately, the use of the unconstrained submodular maximization technique in the above approach is very problematic for two reasons. First, this technique is based on ideas that are very different from the ideas used by the analysis of Measured Continuous Greedy. This makes the combination of the two quite involved. Second, on a more abstract level, the unconstrained submodular maximization technique is based on a symmetry which exists in the absence of a constraint since $\bar{f}(S) = f(\mathcal{N} \setminus S)$ is non-negative and submodular whenever f has these properties. However, this symmetry breaks when a constraint is introduced, and thus, the unconstrained submodular maximization technique does not seem to be a good fit for a constrained problem.

Our algorithm replaces the symmetry based unconstrained submodular maximization technique with a local search algorithm. More specifically, it first executes the local search algorithm. If the output of the local search algorithm is good, then our algorithm simply returns it. Otherwise, we observe that the poor value of the output of the local search algorithm guarantees that it is also far from OPT in some sense. Our algorithm then uses this far from OPT solution to guide an instance of Measured Continuous Greedy, and help it avoid bad decisions.

As it turns out, the analysis of Measured Continuous Greedy and the local search algorithm use similar ideas and notions. Thus, the two algorithms combine quite cleanly, as can be observed from Section 3.

2 Preliminaries

Our analysis uses another useful extension of submodular functions. Given a submodular function $f\colon 2^{\mathcal{N}} \to \mathbb{R}$, its Lovász extension is a function $\hat{f}\colon [0,1]^{\mathcal{N}} \to \mathbb{R}$ defined by

$$\hat{f}(x) = \int_0^1 f(T_\lambda(x))d\lambda \ ,$$

where $T_\lambda(x) = \{u \in \mathcal{N} : x_u < \lambda\}$. The Lovász extension has many important applications (see, e.g., [9, 52]), however, in this paper we only use it in the context of the following known result (which is an immediate corollary of the work of [51]).

Lemma 2.1. *Given the multilinear extension F and the Lovász extension $\hat{f}$ of a submodular function $f\colon 2^{\mathcal{N}} \to \mathbb{R}$, it holds that $F(x) \geq \hat{f}(x)$ for every vector $x \in [0,1]^{\mathcal{N}}$.*

We now define some additional notation that we use. Given a set $S \subseteq \mathcal{N}$ and an element $u \in \mathcal{N}$, we denote by $\mathbf{1}_S$ and $\mathbf{1}_u$ the characteristic vectors of the sets S and $\{u\}$, respectively, and by $S + u$ and $S - u$ the sets $S \cup \{u\}$ and $S \setminus \{u\}$, respectively. Given two vectors $x, y \in [0,1]^{\mathcal{N}}$, we denote by $x \vee y$, $x \wedge y$ and $x \circ y$ the coordinate-wise maximum, minimum and multiplication, respectively, of x and y.[4] Finally, given a vector $x \in [0,1]^{\mathcal{N}}$ and an element $u \in \mathcal{N}$, we denote by $\partial_u F(x)$ the derivative of F with respect to u at the point x. The following observation gives a simple formula for $\partial_u F(x)$. This observation holds because F is a multilinear function.

Observation 2.2. *Let $F(x)$ be the multilinear extension of a submodular function $f\colon 2^{\mathcal{N}} \to \mathbb{R}$. Then, for every $u \in \mathcal{N}$ and $x \in [0,1]^{\mathcal{N}}$,*

$$(1 - x_u) \cdot \partial_u F(x) = F(x \vee \mathbf{1}_u) - F(x) \ .$$

In the rest of the paper we assume, without loss of generality, that $\mathbf{1}_u \in P$ for every element $u \in \mathcal{N}$ and that n is larger than any given constant. The first assumption is justified by the observation that every element u violating this assumption can be safely removed from $\mathcal{N}$ since it cannot belong to OPT. The second assumption is justified by the observation that it is possible to find a set S obeying $\mathbf{1}_S \in P$ and $f(S) = f(OPT)$ in constant time when n is a constant.

Another issue that needs to be kept in mind is the representation of submodular functions. We are interested in algorithms whose time complexity is polynomial in $|\mathcal{N}|$. However, the representation of the submodular function f might be exponential in this size; thus, we cannot assume that the representation of f is given as part of the input for the algorithm. The standard way to bypass this difficulty is to assume that the algorithm has access to f through an oracle. We assume the standard *value oracle* that is used in most of the previous works on submodular maximization. This oracle returns, given any subset $S \subseteq \mathcal{N}$, the value $f(S)$.

3 Main Algorithm

In this section we present the algorithm used to prove Theorem 1.1. This algorithm uses two components. The first component is a close variant of a fractional local search algorithm suggested by Chekuri et al. [13] which has the following properties.

---

[4]More formally, for every element $u \in \mathcal{N}$, $(x \vee y)_u = \max\{x_u, y_u\}$, $(x \wedge y)_u = \min\{x_u, y_u\}$ and $(x \circ y)_u = x_u \cdot y_u$.

Lemma 3.1 (Follows from Chekuri et al. [13]). *There exists a polynomial time algorithm which returns vector $x \in P$ such that, with high probability, for every vector $y \in P$,*

$$F(x) \geq \frac{1}{2}F(x \wedge y) + \frac{1}{2}F(x \vee y) - o(1) \cdot f(OPT) \ . \tag{1}$$

Proof. Let $M = \max\{f(u), f(\mathcal{N} - u) : u \in \mathcal{N}\}$, and let a be an arbitrary constant larger than 3. Then, Lemmata 3.7 and 3.8 of Chekuri et al. [13] imply that, with high probability, the fractional local search algorithm they suggest terminates in polynomial time and outputs a vector $x \in P$ obeying, for every vector $y \in P$,

$$2F(x) \geq F(x \wedge y) + F(x \vee y) - \frac{5M}{n^{a-2}} \ .$$

Moreover, the output vector x is in P whenever the fractional local search algorithm terminates.

Our assumption that $\mathbf{1}_u \in P$ for every element $u \in \mathcal{N}$ implies, by submodularity, that $f(S) \leq n \cdot f(OPT)$ for every set $S \subseteq \mathcal{N}$. Since M is the maximum over values of f, we get also $M \leq n \cdot f(OPT)$. Using this observation, and plugging $a = 4$, we get that there exists an algorithm which, with high probability, terminates after $T(n)$ operations (for some polynomial function $T(n)$) and outputs a vector $x \in P$ obeying $2F(x) \geq F(x \wedge y) + F(x \vee y) - \frac{5 \cdot f(OPT)}{n}$ for every vector $y \in P$. Moreover, the output vector x belongs to P whenever the algorithm terminates.

To complete the lemma, we consider a procedure that executes the above algorithm for $T(n)$ operations, and return its output if it terminates within this number of operations. If the algorithm fails to terminate within this number of operations, which happens with a diminishing probability, then the procedure simply returns $1_\varnothing$ (which always belongs to $\mathcal{P}$ since $\mathcal{P}$ is down-closed). One can observe that this procedure has all the properties guaranteed by the lemma. $\square$

The second component of our algorithm is a new auxiliary algorithm which we present and analyze in Section 4. This auxiliary algorithm is the main technical contribution of this paper, and its guarantee is given by the following theorem.

Theorem 3.2. *There exists a polynomial time algorithm that given a vector $z \in [0,1]^{\mathcal{N}}$ and a value $t_s \in [0,1]$ outputs a vector $x \in P$ obeying*

$$\mathbb{E}[F(x)] \geq e^{t_s - 1} \cdot [(2 - t_s - e^{-t_s} - o(1)) \cdot f(OPT) - (1 - e^{-t_s}) \cdot F(z \wedge \mathbf{1}_{OPT}) \tag{2}$$
$$- (2 - t_s - 2e^{-t_s}) \cdot F(z \vee \mathbf{1}_{OPT})] \ .$$

Our main algorithm executes the algorithms suggested by Lemma 3.1 followed by the algorithm suggested by Theorem 3.2. Notice that the second of these algorithms has two parameters in addition to f and P: a parameter z which is set to be the output of the first algorithm, and a parameter t_s which is set to be a constant to be determined later. After the two above algorithms terminate, our algorithm returns the output of the first algorithm with probability p, for a constant p to be determined later, and with the remaining probability it returns the output of the second algorithm.[5] A formal description of our algorithm is given as Algorithm 1. Observe that Lemma 3.1 and Theorem 3.2 imply together that Algorithm 1 is a polynomial time algorithm which always outputs a vector in P.

To prove Theorem 1.1, it remains to analyze the quality of the solution produced by Algorithm 1.

---

[5]Clearly it is always better to return the better of the two solution instead of randomizing between them. However, doing so will require the algorithm to either have an oracle access to F or estimate the values of the solutions using sampling (the later can be done using standard techniques—see, *e.g.*, [8]). For the sake of simplicity, we chose here the easier to analyze approach of randomizing between the two solutions.

Algorithm 1: Main Algorithm(f, P)

1 Execute the algorithm suggested by Lemma 3.1, and let $x_1 \in P$ be its output.
2 Execute the algorithm suggested by Theorem 3.2 with $z = x_1$, and let x_2 be its output.
3 **return** *with probability p the solution x_1, and the solution x_2 otherwise.*

Lemma 3.3. *When its parameters are set to $t_s = 0.372$ and $p = 0.23$, Algorithm 1 produces a solution whose expected value is at least $0.385 \cdot f(OPT)$.*

Proof. Let $\mathcal{E}$ be the event that x_1, the output of the algorithm suggested by Lemma 3.1, satisfies Inequality (1). Since $\mathcal{E}$ is a high probability event, it is enough to prove that, conditioned on $\mathcal{E}$, Algorithm 1 produces a solution whose expected value is at least $c \cdot f(OPT)$ for some constant $c > 0.385$. The rest of the proof of the lemma is devoted to proving the last claim. Throughout it, everything is implicitly conditioned on $\mathcal{E}$.

As we are conditioning on $\mathcal{E}$, we can plug $y = \mathbf{1}_{OPT}$ and, respectively, $y = x_1 \wedge \mathbf{1}_{OPT}$ into Inequality (1) to get

$$F(x_1) \geq \frac{1}{2}F(x_1 \wedge \mathbf{1}_{OPT}) + \frac{1}{2}F(x_1 \vee \mathbf{1}_{OPT}) - o(1) \cdot f(OPT) \tag{3}$$

and

$$F(x_1) \geq F(x_1 \wedge \mathbf{1}_{OPT}) - o(1) \cdot f(OPT) \ , \tag{4}$$

where the last inequality follows by noticing that $x_1 \vee (x_1 \wedge \mathbf{1}_{OPT}) = x_1$. Next, let $\mathbb{E}[F(x_2) \mid x_1]$ denote the expected value of $F(x_2)$ conditioned on the given value of x_1. Inequality (2) guarantees that

$$\mathbb{E}[F(x_2) \mid x_1] \geq e^{t_s - 1} \cdot [(2 - t_s - e^{-t_s} - o(1)) \cdot f(OPT) - (1 - e^{-t_s}) \cdot F(x_1 \wedge \mathbf{1}_{OPT}) \tag{5}$$
$$- (2 - t_s - 2e^{-t_s}) \cdot F(x_1 \vee \mathbf{1}_{OPT})] \ .$$

Recall that Algorithm 1 returns x_1 with probability p, and x_2 otherwise. Hence, the expected value of its output is

$$\mathbb{E}[p \cdot F(x_1) + (1 - p) \cdot \mathbb{E}[F(x_2) \mid x_1]] \ , \tag{6}$$

where the expectation is over x_1.

Optimizing the constants. We would like to derive from Inequalities (3), (4) and (5) the best lower bound we can get on (6). To this end, let p_1 and p_2 be two non-negative numbers such that $p_1 + p_2 = p$, and let $p_3 = 1 - p$. Using the above inequalities and this notation, (6) can now be lower bounded by

$$p_1 \cdot \left[\frac{1}{2}\mathbb{E}[F(x_1 \wedge \mathbf{1}_{OPT})] + \frac{1}{2}\mathbb{E}[F(x_1 \vee \mathbf{1}_{OPT})] - o(1) \cdot f(OPT)\right]$$
$$+ p_2 \cdot [\mathbb{E}[F(x_1 \wedge \mathbf{1}_{OPT})] - o(1) \cdot f(OPT)]$$
$$+ p_3 \cdot e^{t_s - 1} \cdot [(2 - t_s - e^{-t_s} - o(1)) \cdot f(OPT) - (1 - e^{-t_s}) \cdot \mathbb{E}[F(x_1 \wedge \mathbf{1}_{OPT})]$$
$$- (2 - t_s - 2e^{-t_s}) \cdot \mathbb{E}[F(x_1 \vee \mathbf{1}_{OPT})]] \ ,$$

which can be rewritten as

$$\left(\frac{p_1}{2} + p_2 - p_3 \cdot e^{t_s - 1}(1 - e^{-t_s})\right) \cdot \mathbb{E}[F(x_1 \wedge \mathbf{1}_{OPT})]$$
$$+ \left(\frac{p_1}{2} - p_3 \cdot e^{t_s - 1}(2 - t_s - 2e^{-t_s})\right) \cdot \mathbb{E}[F(x_1 \vee \mathbf{1}_{OPT})]$$
$$+ p_3 \cdot e^{t_s - 1}(2 - t_s - e^{-t_s}) \cdot f(OPT) - o(1) \cdot f(OPT) \ .$$

6

To get the most out of this lower bound we need to maximize the coefficient of $f(OPT)$ while keeping the coefficients of $\mathbb{E}[F(x_1 \wedge \mathbf{1}_{OPT})]$ and $\mathbb{E}[F(x_1 \vee \mathbf{1}_{OPT})]$ non-negative (so that they can be ignored due to non-negativity of f). This objective is formalized by the following non-convex program.

$$
\begin{aligned}
\max \quad & p_3 \cdot e^{t_s-1}(2 - t_s - e^{-t_s}) \\
\text{s.t.} \quad & p_1/2 + p_2 - p_3 \cdot e^{t_s-1}(1 - e^{-t_s}) && \geq 0 \\
& p_1/2 - p_3 \cdot e^{t_s-1}(2 - t_s - 2e^{-t_s}) && \geq 0 \\
& p_1 + p_2 + p_3 && = 1 \\
& p_1, p_2, p_3, t_s && \geq 0
\end{aligned}
$$

Solving the program, we get that the best solution is approximately $p_1 = 0.205$, $p_2 = 0.025$, $p_3 = 0.770$ and $t_s = 0.372$, and the objective function value corresponding to this solution is at least 0.3856. Hence, we have managed to lower bound (6) (and thus, also the expected value of the output of Algorithm 1) by $0.3856 \cdot f(OPT)$ for $p = 0.23$ and $t_s = 0.372$, which completes the proof of the lemma. $\qquad\square$

4 Aided Measured Continuous Greedy

In this section we present the algorithm used to prove Theorem 3.2. Proving the above theorem directly is made more involved by the fact that the vector z might be fractional. Instead, we prove the following simplified version of Theorem 3.2 for integral values, and show that the simplified version implies the original one.

Theorem 4.1. *There exists a polynomial time algorithm that given a set $Z \subseteq \mathcal{N}$ and a value $t_s \in [0,1]$ outputs a vector $x \in P$ obeying*

$$
\mathbb{E}[F(x)] \geq e^{t_s-1} \cdot [(2 - t_s - e^{-t_s} - o(1)) \cdot f(OPT) - (1 - e^{-t_s}) \cdot f(Z \cap OPT) \\
- (2 - t_s - 2e^{-t_s}) \cdot f(Z \cup OPT)] \ .
$$

Next is the promised proof that Theorem 4.1 implies Theorem 3.2.

Proof of Theorem 3.2 given Theorem 4.1. Consider an algorithm ALG that given the z and t_s arguments specified by Theorem 3.2 executes the algorithm guaranteed by Theorem 4.1 with the same value t_s and with a random set Z distributed like $\mathtt{R}(z)$. The output of ALG is then the output produced by the algorithm guaranteed by Theorem 4.1. Let us denote this output by x.

Theorem 4.1 guarantees that, for every given Z,

$$
\mathbb{E}[F(x) \mid Z] \geq e^{t_s-1} \cdot [(2 - t_s - e^{-t_s} - o(1)) \cdot f(OPT) - (1 - e^{-t_s}) \cdot f(Z \cap OPT) \\
- (2 - t_s - 2e^{-t_s}) \cdot f(Z \cup OPT)] \ .
$$

To complete the proof we take the expectation over Z over the two sides of the last inequality and observe that

$$
\mathbb{E}[f(Z \cap OPT)] = \mathbb{E}[f(\mathtt{R}(z) \cap OPT)] = \mathbb{E}[f(\mathtt{R}(z \wedge \mathbf{1}_{OPT}))] = F(z \wedge \mathbf{1}_{OPT})
$$

and

$$
\mathbb{E}[f(Z \cup OPT)] = \mathbb{E}[f(\mathtt{R}(z) \cup OPT)] = \mathbb{E}[f(\mathtt{R}(z \vee \mathbf{1}_{OPT}))] = F(z \vee \mathbf{1}_{OPT}) \ . \qquad\square
$$

In the rest of this section we give a non-formal proof of Theorem 4.1. This proof explains the main ideas necessary for proving the theorem, but uses some non-formal simplifications such as allowing a direct oracle access to the multilinear extension F and giving the algorithm in the form of a continuous time algorithm (which cannot be implemented on a discrete computer). There are known techniques for getting rid of these simplifications (see, *e.g.*, [8]), and a formal proof of Theorem 4.1 based on these techniques is given in Appendix A.

The algorithm we use for the non-formal proof of Theorem 4.1 is given as Algorithm 2. This algorithm starts with the empty solution $y(0) = \mathbf{1}_\varnothing$ at time 0, and grows this solution over time until it reaches the final solution $y(1)$ at time 1. The way the solution grows varies over time. During the time range $[t_s, 1)$ the solution grows like in the Measured Continuous Greedy algorithm of [24]. On the other hand, during the earlier time range of $[0, t_s)$ the algorithm pretends that the elements of Z do not exist (by giving them negative marginal profits), and grows the solution in the way Measured Continuous Greedy would have grown it if it was given the ground set $\mathcal{N} \setminus Z$. The value t_s is the time in which the algorithm switches between the two ways it uses to grow its solution, thus, the s in the notation t_s stands for "switch".

Algorithm 2: Aided Measured Continuous Greedy (non-formal)(f, P, Z, t_s)

---

1 Let $y(0) \leftarrow \mathbf{1}_\varnothing$.

2 foreach $t \in [0, 1)$ **do**

3   For each $u \in \mathcal{N}$ let $w_u(t) \leftarrow F(y(t) \vee \mathbf{1}_u) - F(y(t))$.

4   Let $x(t) \leftarrow \begin{cases} \arg\max_{x \in P}\{\sum_{u \in \mathcal{N} \setminus Z} w_u(t) \cdot x_u(t) - \sum_{u \in Z} x_u(t)\} & \text{if } t \in [0, t_s) \ , \\ \arg\max_{x \in P} \{\sum_{u \in \mathcal{N}} w_u(t) \cdot x_u(t)\} & \text{if } t \in [t_s, 1) \ . \end{cases}$

5   Increase $y(t)$ at a rate of $\frac{dy(t)}{dt} = (\mathbf{1}_\mathcal{N} - y(t)) \circ x(t)$.

6 return $y(1)$.

We first note that algorithm outputs a vector in P.

Observation 4.2. $y(1) \in P$.

Proof. Observe that $x(t) \in P$ at each time t, which implies that $(\mathbf{1}_\mathcal{N} - y(t)) \cdot x(t)$ is also in P since P is down-closed. Therefore, $y(1) = \int_0^1 (\mathbf{1}_\mathcal{N} - y(t)) \cdot x(t) dt$ is a convex combination of vectors in P, and thus, belongs to P. $\qquad\square$

The following lemma lower bounds the increase in $F(y(t))$ as a function of t.

Lemma 4.3. *For every $t \in [0, 1)$,*

$$\frac{dF(y(t))}{dt} \geq \begin{cases} F(y(t) \vee \mathbf{1}_{OPT \setminus Z}) - F(y(t)) & \text{if } t \in [0, t_s) \ , \\ F(y(t) \vee \mathbf{1}_{OPT}) - F(y(t)) & \text{if } t \in [t_s, 1) \ . \end{cases}$$

Proof. By the chain rule,

$$\frac{dF(y(t))}{dt} = \sum_{u \in \mathcal{N}} \left(\frac{dy_u(t)}{dt} \cdot \left.\frac{\partial F(y)}{\partial y_u}\right|_{y=y(t)} \right) = \sum_{u \in \mathcal{N}} \left((1 - y_u(t)) \cdot x_u(t) \cdot \left.\frac{\partial F(y)}{\partial y_u}\right|_{y=y(t)} \right) \qquad (7)$$

$$= \sum_{u \in \mathcal{N}} (x_u(t) \cdot [F(y(t) \vee \mathbf{1}_u) - F(y(t))]) = \sum_{u \in \mathcal{N}} x_u(t) \cdot w_u(t) = x(t) \cdot w(t) \ .$$

Consider first the case $t \in [0, t_s)$. During this time period Algorithm 2 chooses $x(t)$ as the vector in P maximizing $\sum_{u \in \mathcal{N} \setminus Z} w_u(t) \cdot x_u(t) - \sum_{u \in Z} x_u(t)$. Since P is down-closed $x(t) = \mathbf{1}_{OPT \setminus Z}$

8

is in P and has value $\mathbf{1}_{OPT \setminus Z} \cdot w(t)$ and thus, we have $x(t) \cdot w(t) \geq \mathbf{1}_{OPT \setminus Z} \cdot w(t)$. Plugging this observation into Equality (7) yields

$$\frac{dF(y(t))}{dt} = x(t) \cdot w(t) \geq \mathbf{1}_{OPT \setminus Z} \cdot w(t) = \sum_{u \in OPT \setminus Z} [F(y(t) \vee \mathbf{1}_u) - F(y(t))]$$

$$\geq F(y(t) \vee \mathbf{1}_{OPT \setminus Z}) - F(y(t)) ,$$

where the last inequality holds by the submodularity of f.

Similarity, when $t \in [t_s, 1)$ Algorithm 2 chooses $x(t)$ as the vector in P maximizing $x(t) \cdot w(t)$. Since $\mathbf{1}_{OPT} \in P$, we get this time $x(t) \cdot w(t) \geq \mathbf{1}_{OPT} \cdot w(t)$. Plugging this observation into Equality (7) yields

$$\frac{dF(y(t))}{dt} = x(t) \cdot w(t) \geq \mathbf{1}_{OPT} \cdot w(t) = \sum_{u \in OPT} [F(y(t) \vee \mathbf{1}_u) - F(y(t))]$$

$$\geq F(y(t) \vee \mathbf{1}_{OPT}) - F(y(t)) ,$$

where the last inequality holds again by the submodularity of f. $\qquad\square$

Lemma 4.4. *For every time $t \in [0, 1)$ and set $A \subseteq \mathcal{N}$ it holds that*

$$F(y(t) \vee \mathbf{1}_A) \geq \left(e^{-\max\{0, t - t_s\}} - e^{-t} \right) \max\{0, f(A) - f(A \cup Z)\} + e^{-t} \cdot f(A) .$$

Proof. First, we note that for every time $t \in [0, 1]$ and element $u \in \mathcal{N}$,

$$y_u(t) \leq \begin{cases} 1 - e^{-t} & \text{if } u \notin Z , \\ 1 - e^{-\max\{0, t - t_s\}} & \text{if } u \in Z . \end{cases} \tag{8}$$

This follows for the following reason. Since $x(t)$ is always in $P \subseteq [0, 1]^{\mathcal{N}}$, $y_u(t)$ obeys the differential inequality

$$\frac{dy(t)}{dt} = (1 - y_u(t)) \cdot x(t) \leq (1 - y_u(t)) .$$

Using the initial condition $y_u(0) = 0$, the solution for this differential inequality is $y_u(t) \leq 1 - e^{-t}$. To get the tighter bound for $u \in Z$, we note that at every time $t \in [0, t_s)$ Algorithm 2 chooses as $x(t)$ a vector maximizing a linear function in P which assigns a negative weight to elements of Z. Since P is down-closed this maximum must have $x_u(t) = 0$ for every element $u \in Z$. This means that $y_u(t) = 0$ whenever $u \in Z$ and $t \in [0, t_s]$. Moreover, plugging the improved initial condition $y_u(t_s) = 0$ into the above differential inequality yields the promised tighter bound also for the range $(t_s, 1]$.

Next, let $\hat{f}$ be the Lovász extension of f. Then, by Lemma 2.1,

$$F(y(t) \vee \mathbf{1}_A) \geq \hat{f}(y(t) \vee \mathbf{1}_A) = \int_0^1 f(T_\lambda(y(t) \vee \mathbf{1}_A)) d\lambda$$

$$\geq \int_{1 - e^{-\max\{0, t - t_s\}}}^{1 - e^{-t}} f(T_\lambda(y(t) \vee \mathbf{1}_A)) d\lambda + \int_{1 - e^{-t}}^{1} f(T_\lambda(y(t) \vee \mathbf{1}_A)) d\lambda \tag{9}$$

$$= \int_{1 - e^{-\max\{0, t - t_s\}}}^{1 - e^{-t}} f(T_\lambda(y(t) \vee \mathbf{1}_A)) d\lambda + e^{-t} \cdot f(A) \tag{10}$$

$$\geq \left(e^{-\max\{0, t - t_s\}} - e^{-t} \right) \max\{0, f(A) - f(A \cup Z)\} + e^{-t} \cdot f(A) . \tag{11}$$

9

Inequality (9) follows by the non-negativity of f. Equality (10) follows since, for $\lambda \in [1 - e^{-t}, 1)$, Inequality (8) guarantees that $y_u(t) \leq \lambda$ for every $u \in \mathcal{N}$, and thus, $T_\lambda(y(t) \vee \mathbf{1}_A) = A$. Finally Inequality (11) follows since, for $\lambda \in [1 - e^{-\max\{0, t-t_s\}}, 1 - e^{-t})$, Inequality (8) guarantees that $y_u(t) \leq \lambda$ for every $u \in Z$, and thus, $T_\lambda(y(t) \vee \mathbf{1}_A) = B(\lambda) \cup A$ for some $B(\lambda) \subseteq \mathcal{N} \setminus Z$. By the non-negativity of f, $f(B(\lambda) \cup A) \geq 0$. Also, by the submodularity and non-negativity of f, for every such set $B(\lambda)$

$$f(B(\lambda) \cup A) \geq f(A) + f(B(\lambda) \cup Z \cup A) - f(Z \cup A) \geq f(A) - f(Z \cup A) \ . \qquad \square$$

Plugging the results of Lemma 4.4 into the lower bound given by Lemma 4.3 on the improvement in $F(y(t))$ as a function of t yields immediately the useful lower bound given by the next corollary.[6]

Corollary 4.5. *For every $t \in [0, 1)$,*

$$\frac{dF(y(t))}{dt} \geq \begin{cases} f(OPT \setminus Z) - (1 - e^{-t}) \cdot f(Z \cup OPT) - F(y(t)) & \text{if } t \in [0, t_s) \ , \\ e^{t_s - t} \cdot f(OPT) - (e^{t_s - t} - e^{-t}) \cdot f(Z \cup OPT) - F(y(t)) & \text{if } t \in [t_s, 1) \ . \end{cases}$$

Using the last corollary we can complete the proof of Theorem 4.1.

Proof of Theorem 4.1. We have already seen that $y(1)$—the output of Algorithm 2—belongs to P. It remains to show that

$$F(y(1)) \geq e^{t_s - 1} \cdot [(2 - t_s - e^{-t_s}) \cdot f(OPT) - (1 - e^{-t_s}) \cdot f(Z \cap OPT)$$
$$- (2 - t_s - 2e^{-t_s}) \cdot f(Z \cup OPT)] \ .$$

Corollary 4.5 describes a differential inequality for $F(y(t))$. Given the boundary condition $F(y(0)) \geq 0$, the solution for this differential inequality within the range $t \in [0, t_s)$ is

$$F(y(t)) \geq (1 - e^{-t}) \cdot f(OPT \setminus Z) - (1 - e^{-t} - te^{-t}) \cdot f(Z \cup OPT) \ .$$

Plugging $t = t_s$ into the last inequality, we get

$$F(y(t_s)) \geq (1 - e^{-t_s}) \cdot f(OPT \setminus Z) - (1 - e^{-t_s} - t_s e^{-t_s}) \cdot f(Z \cup OPT) \ .$$

Let $v = (1 - e^{-t_s}) \cdot f(OPT \setminus Z) - (1 - e^{-t_s} - t_s e^{-t_s}) \cdot f(Z \cup OPT)$ be the right hand side of the last inequality. Next, we solve again the differential inequality given by Corollary 4.5 for the range $t \in [t_s, 1]$ with the boundary condition $F(y(t_s)) \geq v$. The resulting solution is

$$F(y(t)) \geq e^{-t} \left[(t - t_s) \left(e^{t_s} \cdot f(OPT) - (e^{t_s} - 1) \cdot f(Z \cup OPT) \right) + v e^{t_s} \right]$$

Plugging $t = 1$ and the value of v we get

$$F(y(1)) \geq e^{-1} \left[(1 - t_s) \left(e^{t_s} \cdot f(OPT) - (e^{t_s} - 1) \cdot f(Z \cup OPT) \right) + v e^{t_s} \right]$$
$$\geq \frac{1 - t_s}{e} \left(e^{t_s} \cdot f(OPT) - (e^{t_s} - 1) \cdot f(Z \cup OPT) \right) \qquad (12)$$
$$+ e^{t_s - 1} \cdot \{ (1 - e^{-t_s}) \cdot [f(OPT) - f(OPT \cap Z)] - (1 - e^{-t_s} - t_s e^{-t_s}) \cdot f(Z \cup OPT) \}$$
$$= e^{t_s - 1} \cdot [(2 - t_s - e^{-t_s}) \cdot f(OPT) - (1 - e^{-t_s}) \cdot f(Z \cap OPT)$$
$$- (2 - t_s - 2e^{-t_s}) \cdot f(Z \cup OPT)] \ ,$$

where Inequality (12) follows since, by the submodularity and non-negativity of f,

$$f(OPT \setminus Z) \geq f(OPT) - f(OPT \cap Z) + f(\varnothing) \geq f(OPT) - f(OPT \cap Z) \ . \qquad \square$$

---

[6]Note that Corollary 4.5 follows from a weaker version of Lemma 4.4 which only guarantees $F(y(t) \vee \mathbf{1}_A) \geq (e^{-\max\{0, t-t_s\}} - e^{-t}) \cdot [f(A) - f(A \cup Z)] + e^{-t} \cdot f(A)$. We proved the stronger version of the lemma above because it is useful in the formal proof of Theorem 4.1 given in Appendix A.

References

[1] A. A. Ageev and M. I. Sviridenko. An 0.828 approximation algorithm for the uncapacitated facility location problem. *Discrete Appl. Math.*, 93:149–156, July 1999.

[2] Noga Alon and Joel H. Spencer. *The Probabilistic Method*. Wiley Interscience, second edition, 2000.

[3] Per Austrin, Siavosh Benabbas, and Konstantinos Georgiou. Better balance by being biased: A 0.8776-approximation for max bisection. In *SODA*, pages 277–294, 2013.

[4] Francis Bach. Learning with submodular functions: A convex optimization perspective. *Foundations and Trends in Machine Learning*, 6(2-3):145–373, 2013.

[5] Y. Y. Boykov and M. P. Jolly. Interactive graph cuts for optimal boundary & region segmentation of objects in N-D images. In *ICCV*, volume 1, pages 105–112, 2001.

[6] Niv Buchbinder, Moran Feldman, Joseph (Seffi) Naor, and Roy Schwartz. A tight linear time (1/2)-approximation for unconstrained submodular maximization. In *FOCS*, pages 649–658, 2012.

[7] Niv Buchbinder, Moran Feldman, Joseph (Seffi) Naor, and Roy Schwartz. Submodular maximization with cardinality constraints. In *SODA*, pages 1433–1452, 2014.

[8] Gruia Călinescu, Chandra Chekuri, Martin Pál, and Jan Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM J. Comput.*, 40(6):1740–1766, 2011.

[9] Chandra Chekuri and Alina Ene. Approximation algorithms for submodular multiway partition. In *FOCS*, pages 807–816, 2011.

[10] Chandra Chekuri and Sanjeev Khanna. A polynomial time approximation scheme for the multiple knapsack problem. *SIAM J. Comput.*, 35(3):713–728, September 2005.

[11] Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Dependent randomized rounding via exchange properties of combinatorial structures. In *FOCS*, pages 575–584, 2010.

[12] Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. In *STOC*, pages 783–792, 2011.

[13] Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. *SIAM J. Comput.*, 43(6):1831–1879, 2014.

[14] Reuven Cohen, Liran Katzir, and Danny Raz. An efficient approximation for the generalized assignment problem. *Information Processing Letters*, 100(4):162–166, 2006.

[15] M. Conforti and G. Cornuèjols. Submodular set functions, matroids and the greedy algorithm. tight worstcase bounds and some generalizations of the radoedmonds theorem. *Disc. Appl. Math.*, 7(3):251–274, 1984.

[16] G. Cornuejols, M. L. Fisher, and G. L. Nemhauser. Location of bank accounts to optimize float: an analytic study of exact and approximate algorithms. *Management Sciences*, 23:789–810, 1977.

[17] G. Cornuejols, M. L. Fisher, and G. L. Nemhauser. On the uncapacitated location problem. *Annals of Discrete Mathematics*, 1:163–177, 1977.

[18] Alina Ene and Huy L. Nguyen. Constrained submodular maximization: Beyond 1/e. In *FOCS*, 2016.

[19] Uriel Feige. A threshold of ln n for approximating set cover. *J. ACM*, 45(4):634–652, 1998.

[20] Uriel Feige and Michel X. Goemans. Aproximating the value of two prover proof systems, with applications to max 2sat and max dicut. In *ISTCS*, pages 182–189, 1995.

[21] Uriel Feige, Vahab S. Mirrokni, and Jan Vondrák. Maximizing non-monotone submodular functions. *SIAM Journal on Computing*, 40(4):1133–1153, 2011.

[22] Uriel Feige and Jan Vondrák. Approximation algorithms for allocation problems: Improving the factor of $1 - 1/e$. In *FOCS*, pages 667–676, 2006.

[23] Moran Feldman. *Maximization Problems with Submodular Objective Functions*. PhD thesis, Technion – Israel Institute of Technology, June 2013.

[24] Moran Feldman, Joseph Naor, and Roy Schwartz. A unified continuous greedy algorithm for submodular maximization. In *FOCS*, pages 570–579, 2011.

[25] Moran Feldman, Joseph (Seffi) Naor, Roy Schwartz, and Justin Ward. Improved approximations for k-exchange systems. In *ESA*, pages 784–798, 2011.

[26] M. L. Fisher, G. L. Nemhauser, and L. A. Wolsey. An analysis of approximations for maximizing submodular set functions – ii. In *Polyhedral Combinatorics*, volume 8 of *Mathematical Programming Studies*, pages 73–87. Springer Berlin Heidelberg, 1978.

[27] Lisa Fleischer, Michel X. Goemans, Vahab S. Mirrokni, and Maxim Sviridenko. Tight approximation algorithms for maximum general assignment problems. In *SODA*, pages 611–620, 2006.

[28] Alan M. Frieze and Mark Jerrum. Improved approximation algorithms for max k-cut and max bisection. In *IPCO*, pages 1–13, 1995.

[29] Shayan Oveis Gharan and Jan Vondrák. Submodular maximization by simulated annealing. In *SODA*, pages 1098–1117, 2011.

[30] Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6):1115–1145, 1995.

[31] Eran Halperin and Uri Zwick. Combinatorial approximation algorithms for the maximum directed cut problem. In *SODA*, pages 1–7, 2001.

[32] Jason Hartline, Vahab Mirrokni, and Mukund Sundararajan. Optimal marketing strategies over social networks. In *WWW*, pages 189–198, 2008.

[33] Johan Hàstad. Some optimal inapproximability results. *J. ACM*, 48:798–859, July 2001.

[34] D. Hausmann and B. Korte. K-greedy algorithms for independence systems. *Oper. Res. Ser. A-B*, 22(1):219–228, 1978.

[35] D. Hausmann, B. Korte, and T. Jenkyns. Worst case analysis of greedy type algorithms for independence systems. *Math. Prog. Study*, 12:120–131, 1980.

[36] S. Jegelka and J. Bilmes. Submodularity beyond submodular energies: Coupling edges in graph cuts. *2012 IEEE Conference on Computer Vision and Pattern Recognition*, 0:1897–1904, 2011.

[37] T. Jenkyns. The efficacy of the greedy algorithm. *Cong. Num.*, 17:341–350, 1976.

[38] Richard M. Karp. Reducibility among combinatorial problems. In R. E. Miller and J. W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, 1972.

[39] David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *SIGKDD*, pages 137–146, 2003.

[40] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'Donnell. Optimal inapproximability results for max-cut and other 2-variable csps? *SIAM J. Comput.*, 37:319–357, April 2007.

[41] S. Khuller, A. Moss, and J. Naor. The budgeted maximum coverage problem. *Information Processing Letters*, 70(1):39–45, 1999.

[42] B. Korte and D. Hausmann. An analysis of the greedy heuristic for independence systems. *Annals of Discrete Math.*, 2:65–74, 1978.

[43] Andreas Krause, AjitSingh, and Carlos Guestrin. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *J. Mach. Learn. Res.*, 9:235–284, January 2008.

[44] Andreas Krause and Carlos Guestrin. Near-optimal nonmyopic value of information in graphical models. In *UAI*, page 5, 2005.

[45] Andreas Krause, Jure Leskovec, Carlos Guestrin, Jeanne VanBriesen, and Christos Faloutsos. Efficient sensor placement optimization for securing large water distribution networks. *Journal of Water Resources Planning and Management*, 134(6):516–526, November 2008.

[46] Ariel Kulik, Hadas Shachnai, and Tami Tamir. Approximations for monotone and nonmonotone submodular maximization with knapsack constraints. *Math. Oper. Res.*, 38(4):729–739, 2013.

[47] Jon Lee, Vahab S. Mirrokni, Viswanath Nagarajan, and Maxim Sviridenko. Maximizing nonmonotone submodular functions under matroid or knapsack constraints. *SIAM Journal on Discrete Mathematics*, 23(4):2053–2078, 2010.

[48] Jon Lee, Maxim Sviridenko, and Jan Vondrák. Submodular maximization over multiple matroids via generalized exchange properties. In *APPROX*, pages 244–257, 2009.

[49] Hui Lin and Jeff Bilmes. Multi-document summarization via budgeted maximization of submodular functions. In *North American chapter of the Association for Computational Linguistics/Human Language Technology Conference (NAACL/HLT-2010)*, Los Angeles, CA, June 2010.

[50] Hui Lin and Jeff Bilmes. A class of submodular functions for document summarization. In *HLT*, pages 510–520, 2011.

[51] László Lovász. Submodular functions and convexity. In A. Bachem, M. Grötschel, and B. Korte, editors, *Mathematical Programming: the State of the Art*, pages 235–257. Springer, 1983.

[52] L. Lovász M. Grötschel and A. Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatoria*, 1(2):169–197, 1981.

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

[54] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functionsi. *Mathematical Programming*, 14:265–294, 1978.

[55] Maxim Sviridenko. A note on maximizing a submodular set function subject to knapsack constraint. *Operations Research Letters*, 32:41–43, 2004.

[56] Luca Trevisan, Gregory B. Sorkin, Madhu Sudan, and David P. Williamson. Gadgets, approximation, and linear programming. *SIAM J. Comput.*, 29:2074–2097, April 2000.

[57] Jan Vondrák. Symmetry and approximability of submodular maximization problems. *SIAM J. Comput.*, 42(1):265–304, 2013.

A A Formal Proof of Theorem 4.1

In this section we give a formal proof of Theorem 4.1. This proof is based on the same ideas used in the non-formal proof of this theorem in Section 4, but employs also additional known techniques in order to get rid of the issues that make the proof from Section 4 non-formal.

The algorithm we use to prove Theorem 4.1 is given as Algorithm 3. This algorithm is a discrete variant of Algorithm 2. While reading the algorithm, it is important to observe that the choice of the values $\bar{\delta}_1$ and $\bar{\delta}_2$ guarantees that the variable t takes each one of the values t_s and 1 at some point, and thus, the vectors $y(t_s)$ and $y(1)$ are well defined.

Algorithm 3: Aided Measured Continuous Greedy(f, P, Z, t_s)

    // Initialization

1 Let $\bar{\delta}_1 \leftarrow t_s \cdot n^{-4}$ and $\bar{\delta}_2 \leftarrow (1 - t_s) \cdot n^{-4}$.

2 Let $t \leftarrow 0$ and $y(t) \leftarrow \mathbf{1}_\varnothing$.

    // Growing $y(t)$

3 while $t < 1$ **do**

4   **foreach** $u \in \mathcal{N}$ **do**

5   Let $w_u(t)$ be an estimate of $\mathbb{E}[f(u \mid \mathtt{R}(y(t))]$ obtained by averaging the values of $f(u \mid \mathtt{R}(y(t))$ for $r = \lceil 48n^6 \ln(2n) \rceil$ independent samples of $\mathtt{R}(y(t))$.

6   Let $x(t) \leftarrow \begin{cases} \arg\max_{x \in P}\{\sum_{u \in \mathcal{N} \setminus Z} w_u(t) \cdot x_u(t) - \sum_{u \in Z} x_u(t)\} & \text{if } t \in [0, t_s) \ , \\ \arg\max_{x \in P}\{\sum_{u \in \mathcal{N}} w_u(t) \cdot x_u(t)\} & \text{if } t \in [t_s, 1) \ . \end{cases}$

7   Let δ_t be $\bar{\delta}_1$ when $t < t_s$ and $\bar{\delta}_2$ when $t \geq t_s$.

8   Let $y(t + \delta_t) \leftarrow y(t) + \delta_t(\mathbf{1}_{\mathcal{N}} - y(t)) \circ x(t)$.

9   Update $t \leftarrow t + \delta_t$.

10 return $y(1)$.

We begin the analysis of Algorithm 3 by showing that $y(t)$ remains within the cube $[0,1]^{\mathcal{N}}$ throughout the execution of the algorithm. Without this observation, the algorithm is not well-defined.

Observation A.1. *For every value of t, $y(t) \in [0,1]^{\mathcal{N}}$.*

Proof. We prove the observation by induction on t. Clearly the observation holds for $y(0) = \mathbf{1}_\varnothing$. Assume the observation holds for some time t, then, for every $u \in \mathcal{N}$,

$$y_u(t + \delta_t) = y_u(t) + \delta_t(1 - y_u(t)) \cdot x_u(t) \geq 0 \ ,$$

where the inequality holds since the induction hypothesis implies $1 - y_u(t) \in [0,1]$. A similar argument also implies

$$y_u(t + \delta_t) = y_u(t) + \delta_t(1 - y_u(t)) \cdot x_u(t) \leq y_u(t) + (1 - y_u(t)) = 1 \ . \qquad \square$$

Using the last observation it is now possible to prove the following counterpart of Observation 4.2.

Corollary A.2. *Algorithm 3 always outputs a vector in P.*

Proof. Let T be the set of values t takes during the execution of Algorithm 3. We observe that $\sum_{t \in T \setminus \{1\}} \delta_t = 1$, which implies that $\sum_{t \in T \setminus \{1\}} \delta_t \cdot x(t)$ is a convex combination of the vectors $\{x(t) : t \in T \setminus \{1\}\}$. As all these vectors belong to P, and P is convex, any convex combination of them, including $\sum_{t \in T \setminus \{1\}} \delta_t \cdot x(t)$, must be in P.

Next, we rewrite the output of Algorithm 3 as

$$y(1) = \sum_{t \in T \setminus \{1\}} \delta_t(\mathbf{1}_\mathcal{N} - y(t)) \circ x(t) \leq \sum_{t \in T \setminus \{1\}} \delta_t \cdot x(t) \ .$$

By the above discussion the rightmost hand side of this inequality is a vector in P, which implies that $y(1) \in P$ since P is down-closed. $\qquad \square$

The next step towards showing that Algorithm 3 proves Theorem 4.1 is analyzing its approximation ratio. We start this analysis by showing that with high probability all the estimations made by the algorithm are quite accurate. Let $\mathcal{A}$ be the event that $|w_u(t) - \mathbb{E}[f(u \mid \mathtt{R}(y(t)))]| \leq n^{-2} \cdot f(OPT)$ for every $u \in \mathcal{N}$ and time t.

Lemma A.3 (The symmetric version of Theorem A.1.16 in [2])**.** *Let X_i, $1 \leq i \leq k$, be mutually independent with all $\mathbb{E}[X_i] = 0$ and all $|X_i| \leq 1$. Set $S = X_1 + \cdots + X_k$. Then, $\Pr[|S| > a] \leq 2e^{-a^2/2k}$.*

Corollary A.4. $\Pr[\mathcal{A}] \geq 1 - n^{-1}$.

Proof. Consider the calculation of $w_u(t)$ for a given $u \in \mathcal{N}$ and time t. This calculation is done by averaging the value of $f(u \mid \mathtt{R}(y(t)))$ for r independent samples of $\mathtt{R}(y(t))$. Let Y_i denote the value of $f(u \mid \mathtt{R}(y(t)))$ obtained for the i-th sample, and let $X_i = \frac{Y_i - \mathbb{E}[f(u \mid \mathtt{R}(y(t)))]}{2n \cdot f(OPT)}$. Then, by definition,

$$w_u(t) = \frac{\sum_{i=1}^r Y_i}{r} = [2n \cdot f(OPT)] \cdot \frac{\sum_{i=1}^r X_i}{r} + \mathbb{E}[f(u \mid \mathtt{R}(y(t)))] \ .$$

Since Y_i is distributed like $f(u \mid \mathtt{R}(y(t)))$, the definition of X_i guarantees that $\mathbb{E}[X_i] = 0$ for every $1 \leq i \leq r$. Additionally, $|X_i| \leq 1$ for every such i since the absolute values of both Y_i and

$\mathbb{E}[f(u \mid \mathtt{R}(y(t)))]$ are upper bounded by $\max_{S \subseteq \mathcal{N}} f(S) \le n \cdot f(OPT)$ (the last inequality follows from our assumption that $\mathbf{1}_u \in P$ for every element $u \in \mathcal{N}$). Thus, by Lemma A.3,

$$\Pr[|w_u(t) - \mathbb{E}[f(u \mid \mathtt{R}(y(t)))]| > n^{-2} \cdot f(OPT)] = \Pr\left[\left|\sum_{i=1}^{r} X_i\right| > \frac{r}{2n^3}\right] \le 2e^{-[rn^{-3}/2]^2/2r}$$

$$= 2e^{-rn^{-6}/8} \le 2e^{-6\ln(2n)} = 2 \cdot \left(\frac{1}{2n}\right)^6 \le \frac{1}{2n^6} \ .$$

Observe that Algorithm 3 calculates $w_u(t)$ for every combination of element $u \in \mathcal{N}$ and time $t < 1$. Since there are n elements in $\mathcal{N}$ and $2n^4$ times smaller than 1, the union bound implies that the probability that for at least one such value $w_u(t)$ we have $|w_u(t) - \mathbb{E}[f(u \mid \mathtt{R}(y(t)))]| > n^{-2} \cdot f(OPT)$ is upper bounded by

$$\frac{1}{2n^6} \cdot \left(n \cdot 2n^4\right) = \frac{1}{n} \ ,$$

which completes the proof of the corollary. $\qquad\square$

Our next step is to give a lower bound on the increase in $F(y(t))$ as a function of t given $\mathcal{A}$. This lower bound is given by Corollary A.7, which follows from the next two lemmata. The statement and proof of the corollary and the next lemma is easier with the following definition. Let OPT'_t denote the set $OPT \setminus Z$ when $t < t_s$, and OPT otherwise.

Lemma A.5. *Given $\mathcal{A}$, for every time $t < 1$, $\sum_{u \in \mathcal{N}}(1 - y_u(t)) \cdot x_u(t) \cdot \partial_u F(y(t)) \ge F(y(t) \vee \mathbf{1}_{OPT'_t}) - F(y(t)) - O(n^{-1}) \cdot f(OPT)$.*

Proof. Let us calculate the weight of OPT'_t according to the weight function $w(t)$.

$$w(t) \cdot \mathbf{1}_{OPT'_t} = \sum_{u \in OPT'_t} w_u(t) \ge \sum_{u \in OPT'_t} [\mathbb{E}[f(u \mid \mathtt{R}(y(t)))] - n^{-2} \cdot f(OPT)]$$

$$\ge \mathbb{E}\left[\sum_{u \in OPT'_t} f(\mathtt{R}(y(t)) + u) - f(\mathtt{R}(y(t)))\right] - n^{-1} \cdot f(OPT)$$

$$\ge \mathbb{E}\left[f(\mathtt{R}(y(t)) \cup OPT'_t) - f(\mathtt{R}(y(t)))\right] - n^{-1} \cdot f(OPT)$$

$$= F(y(t) \vee \mathbf{1}_{OPT'_t}) - F(y(t)) - n^{-1} \cdot f(OPT) \ ,$$

where the first inequality follows from the definition of $\mathcal{A}$, and the last follows from the submodularity of f. Recall that $x(t)$ is the vector in P maximizing some objective function (which depends on t). For $t < t_s$, the objective function maximized by $x(t)$ assigns the value $w(t) \cdot \mathbf{1}_{OPT \setminus Z} = w(t) \cdot \mathbf{1}_{OPT'_t}$ to the vector $\mathbf{1}_{OPT'_t} \in P$. Similarly, for $t \ge t_s$, the objective function maximized by $x(t)$ assigns the value $w(t) \cdot \mathbf{1}_{OPT} = w(t) \cdot \mathbf{1}_{OPT'_t}$ to the vector $\mathbf{1}_{OPT'_t} \in P$. Thus, the definition of $x(t)$ guarantees that in both cases we have

$$w(t) \cdot x(t) \ge w(t) \cdot \mathbf{1}_{OPT'_t} \ge F(y(t) \vee \mathbf{1}_{OPT'_t}) - F(y(t)) - n^{-1} \cdot f(OPT) \ .$$

Hence,

$$\sum_{u \in \mathcal{N}} (1 - y_u(t)) \cdot x_u(t) \cdot \partial_u F(y(t)) = \sum_{u \in \mathcal{N}} x_u(t) \cdot [F(y(t) \vee \mathbf{1}_u) - F(y(t))]$$

$$= \sum_{u \in \mathcal{N}} x_u(t) \cdot \mathbb{E}[f(u \mid \mathtt{R}(y(t)))]$$

$$\geq \sum_{u \in \mathcal{N}} x_u(t) \cdot [w_u(t) - n^{-2} \cdot f(OPT)] = x(t) \cdot w(t) - n^{-1} \cdot f(OPT)$$

$$\geq F(y(t) \vee \mathbf{1}_{OPT'_t}) - F(y(t)) - 2n^{-1} \cdot f(OPT) \ ,$$

where the first inequality holds by the definition of $\mathcal{A}$ and the second equality holds since

$$F(y(t) \vee \mathbf{1}_u) - F(y(t)) = \mathbb{E}[f(\mathtt{R}(y(t)) + u)] - \mathbb{E}[f(\mathtt{R}(y(t)))] = \mathbb{E}[f(u \mid \mathtt{R}(y(t)))] \ . \qquad \square$$

Lemma A.6 (A rephrased version of Lemma 2.3.7 in [23]). *Consider two vectors $x, x' \in [0,1]^{\mathcal{N}}$ such that $|x_u - x'_u| \leq \delta$ for every $u \in \mathcal{N}$. Then, $F(x') - F(x) \geq \sum_{u \in \mathcal{N}} (x'_u - x_u) \cdot \partial_u F(x) - O(n^3 \delta^2) \cdot \max_{u \in N} f(\{u\})$.*

Corollary A.7. *Given $\mathcal{A}$, for every time $t < 1$, $F(y(t + \delta_t)) - F(y(t)) \geq \delta_t [F(y(t) \vee \mathbf{1}_{OPT'_t}) - F(y(t))] - O(n^{-1} \delta_t) \cdot f(OPT)$.*

Proof. Observe that for every $u \in \mathcal{N}$, $|y_u(t + \delta_t) - y_u(t)| = |\delta_t(1 - y_u(t))x_u(t)| \leq \delta_t$. Hence, by Lemma A.6,

$$F(y(t + \delta_t)) - F(y(t)) \geq \sum_{u \in \mathcal{N}} [y_u(t + \delta_t)) - y_u(t)] \cdot \partial_u F(y(t)) - O(n^3 \delta_t^2) \cdot \max_{u \in N} f(\{u\})$$

$$= \sum_{u \in \mathcal{N}} \delta_t (1 - y_u(t)) \cdot x_u(t) \cdot \partial_u F(y(t)) - O(n^3 \delta_t^2) \cdot \max_{u \in N} f(\{u\}) \ . \qquad (13)$$

Consider the rightmost hand side of the last inequality. By Lemma A.5, the first term on this side can be bounded by

$$\sum_{u \in \mathcal{N}} \delta_t (1 - y_u(t)) \cdot x_u(t) \cdot \partial_u F(y(t)) \geq \delta_t \cdot [F(y(t) \vee \mathbf{1}_{OPT'_t}) - F(y(t)) - O(n^{-1}) \cdot f(OPT)]$$

$$= \delta_t \cdot [F(y(t) \vee \mathbf{1}_{OPT'_t}) - F(y(t))] - O(n^{-1} \delta_t) \cdot f(OPT) \ .$$

On the other hand, the second term of (13) can be bounded by

$$O(n^3 \delta_t^2) \cdot \max_{u \in N} f(\{u\}) = O(n^{-1} \delta_t) \cdot f(OPT)$$

since $\delta_t \leq n^{-4}$ by definition and $\max_{u \in N} f(\{u\}) \leq f(OPT)$ by our assumption that $\mathbf{1}_u \in P$ for every $u \in \mathcal{N}$. $\qquad \square$

The lower bound given by the last corollary is in terms of $F(y(t) \vee \mathbf{1}_{OPT'_t})$. To make this lower bound useful, we need to lower bound the term $F(y(t) \vee \mathbf{1}_{OPT'_t})$. This is done by the following two lemma which corresponds to Lemma 4.4.

Lemma A.8. *[corresponds to Lemma 4.4] For every time $t < 1$ and set $A \subseteq \mathcal{N}$ it holds that*

$$F(y(t) \vee \mathbf{1}_A) \geq \left(e^{-\max\{0, t-t_s\}} - e^{-t} - O(n^{-4}) \right) \cdot \max \{0, f(A) - f(A \cup Z)\}$$

$$+ (e^{-t} - O(n^{-4})) \cdot f(A) \ .$$

The proof of this lemma goes along the same lines as the proof of its corresponding lemma in Section 4, except that the bounds on the coordinates of $y(t)$ used by the proof from Section 4 are replaced with the (slightly weaker) bounds given by the following lemma.

Lemma A.9. *For every time t and element $u \in \mathcal{N}$,*

$$y_u(t) \leq \begin{cases} 1 - e^{-t} + O(n^{-4}) & \text{if } u \notin Z \ , \\ 1 - e^{-\max\{0, t-t_s\}} + O(n^{-4}) & \text{if } u \in Z \ . \end{cases}$$

Proof. Let $\varepsilon = n^{-4}$, and observe that $\delta_t \leq \varepsilon$ for every time t. Our first objective is to prove by induction on t that, if $y_u(\tau) = 0$ for some time $\tau \in [0, 1]$, then $y_u(t) \leq 1 - (1 - \varepsilon)^{(t-\tau)/\varepsilon}$ for every time $t \in [\tau, 1]$. For $t = \tau$ the claim holds because $y_u(\tau) = 0 = 1 - (1 - \varepsilon)^{(\tau-\tau)/\varepsilon}$. Next, assume the claim holds for some t, and let us prove it for $t + \delta_t$.

$$y_u(t + \delta_t) = y_u(t) + \delta_t(1 - y_u(t)) \cdot x_u(t) \leq y_u(t) + \delta_t(1 - y_u(t)) = y_u(t)(1 - \delta_t) + \delta_t$$
$$\leq (1 - (1 - \varepsilon)^{(t-\tau)/\varepsilon})(1 - \delta_t) + \delta_t = 1 - (1 - \delta_t)(1 - \varepsilon)^{(t-\tau)/\varepsilon}$$
$$\leq 1 - (1 - \varepsilon)^{\delta_t/\varepsilon}(1 - \varepsilon)^{(t-\tau)/\varepsilon} = 1 - (1 - \varepsilon)^{(t+\delta_t-\tau)/\varepsilon} \ ,$$

where the last inequality holds since $(1 - x)^{1/x}$ is a decreasing function for $x \in (0, 1]$.

We complete the proof for the case $u \notin Z$ by choosing $\tau = 0$ (clearly $y_u(0) = 0$) and observing that, for every time t,

$$1 - (1 - \varepsilon)^{t/\varepsilon} \leq 1 - [e^{-1}(1 - \varepsilon)]^t = 1 - e^{-t}(1 - \varepsilon)^t \leq 1 - e^{-t}(1 - \varepsilon) = 1 - e^{-t} + O(\varepsilon) \ .$$

It remains to prove the lemma for the case $u \in Z$. Note that at every time $t \in [0, t_s)$ Algorithm 3 chooses as $x(t)$ a vector maximizing a linear function in P which assigns a negative weight to elements of Z. Since P is down-closed this maximum must have $x_u(t) = 0$ for an element $u \in Z$. This means that $y_u(t) = 0$ for $t \in [0, t_s]$. In addition to proving the lemma for this time range, the last inequality also allows us to choose $\tau = t_s$, which gives, for $t \in [t_s, 1]$,

$$y_u(t) \geq 1 - (1 - \varepsilon)^{(t-t_s)/\varepsilon} \geq 1 - e^{t_s-t} + O(\varepsilon) \ . \qquad \square$$

Combining Corollary A.7 with Lemma A.8 gives us the following corollary.

Corollary A.10. *Given $\mathcal{A}$, for every time $t \in [0, t_s)$,*

$$F(y(t + \delta_t)) - F(y(t)) \geq \delta_t[f(OPT \setminus Z) - (1 - e^{-t}) \cdot f(Z \cup OPT) - F(y(t)))]$$
$$- O(n^{-1}\delta_t) \cdot f(OPT)$$

and, for every time $t \in [t_s, 1)$,

$$F(y(t + \delta_t)) - F(y(t)) \geq \delta_t[e^{-t} \cdot f(OPT) + (e^{t_s-t} - e^{-t}) \cdot \max\{f(OPT) - f(Z \cup OPT), 0\}$$
$$- F(y(t))] - O(n^{-1}\delta_t) \cdot f(OPT) \ .$$

Proof. For every time $t \in [0, t_s)$, Corollary A.7 and Lemma A.8 imply together

$$F(y(t + \delta_t)) - F(y(t)) \geq \delta_t[(1 - e^{-t} - O(n^{-4})) \cdot \max\{0, f(OPT \setminus Z) - f(OPT \cup Z)\}$$
$$+ (e^{-t} - O(n^{-4})) \cdot f(OPT \setminus Z)] - O(n^{-1}\delta_t) \cdot f(OPT)$$
$$\geq \delta_t[(1 - O(n^{-4})) \cdot f(OPT \setminus Z) - (1 - e^{-t}) \cdot f(Z \cup OPT)$$
$$- F(y(t)))] - O(n^{-1}\delta_t) \cdot f(OPT) \ .$$

18

We observe that this inequality is identical to the inequality promised for this time range by the corollary, except that it has an extra term of $-\delta_t \cdot O(n^{-4}) \cdot f(OPT \setminus Z)$ on its right hand side. Since $f(OPT \setminus Z)$ is upper bounded by $f(OPT)$, due to the down-closeness of P, the absolute value of this extra term is at most

$$\delta_t \cdot O(n^{-4}) \cdot f(OPT) = O(n^{-1}\delta_t) \cdot f(OPT) \ ,$$

which completes the proof for the time range $t \in [0, t_s)$.

Consider now the time range $t \in [t_s, 1)$. For this time range Corollary A.7 and Lemma A.8 imply together

$$F(y(t+\delta_t)) - F(y(t)) \geq \delta_t [(e^{t_s - t} - e^{-t} - O(n^{-4})) \cdot \max\{0, f(OPT) - f(OPT \cup Z)\}$$
$$+ (e^{-t} - O(n^{-4})) \cdot f(OPT)] - O(n^{-1}\delta_t) \cdot f(OPT) \ .$$

We observe again that this inequality is identical to the inequality promised for this time range by the corollary, except that it has extra terms of $-\delta_t \cdot O(n^{-4}) \cdot f(OPT)$ and $-\delta_t \cdot O(n^{-4}) \cdot \max\{0, f(OPT) - f(OPT \cup Z)\}$ on its right hand side. The corollary now follows since the absolute value of both these terms is upper bounded by $O(n^{-1}\delta_t) \cdot f(OPT)$. □

Corollary A.10 bounds the increase in $F(y(t))$ in terms of $F(y(t))$ itself. Thus, it gives a recursive formula which can be used to lower bound $F(y(t))$. Our remaining task is to solve this formula and get a closed-form lower bound on $F(y(t))$. Let $g(t)$ be defined as follows. $g(0) = 0$ and for every time $t < 1$,

$$g(t+\delta_t)$$
$$= \begin{cases} g(t) + \delta_t [f(OPT \setminus Z) - (1 - e^{-t}) \cdot f(Z \cup OPT) - g(t)] & \text{if } t < t_s \ , \\ g(t) + \delta_t [e^{-t} \cdot f(OPT) + (e^{t_s - t} - e^{-t}) \cdot \max\{f(OPT) - f(Z \cup OPT), 0\} - g(t)] & \text{if } t \geq t_s \ . \end{cases}$$

The next lemma shows that a lower bound on $g(t)$ yields a lower bound on $F(y(t))$.

Lemma A.11. *Given $\mathcal{A}$, for every time t, $g(t) \leq F(y(t)) + O(n^{-1}) \cdot t \cdot f(OPT)$.*

Proof. Let c be the larger constant among the constants hiding behind the big O notations in Corollary A.10. We prove by induction on t that $g(t) \leq F(y(t)) + (ct/n) \cdot f(OPT)$. For $t = 0$, this clearly holds since $g(0) = 0 \leq F(y(0))$. Assume now that the claim holds for some t, and let us prove it for $t + \delta_t$. There are two cases to consider. If $t < t_s$, then the induction hypothesis and Corollary A.10 imply, for a large enough n,

$$\begin{aligned} g(t + \delta_t) &= g(t) + \delta_t [f(OPT \setminus Z) - (1 - e^{-t}) \cdot f(Z \cup OPT) - g(t)] \\ &= (1 - \delta_t)g(t) + \delta_t [f(OPT \setminus Z) - (1 - e^{-t}) \cdot f(Z \cup OPT)] \\ &\leq (1 - \delta_t)[F(y(t)) + (ct/n) \cdot f(OPT)] + \delta_t [f(OPT \setminus Z) - (1 - e^{-t}) \cdot f(Z \cup OPT)] \\ &= F(y(t)) + \delta_t [f(OPT \setminus Z) - (1 - e^{-t}) \cdot f(Z \cup OPT) - F(y(t))] \\ &\quad + (ct/n) \cdot (1 - \delta_t) \cdot f(OPT) \\ &\leq F(y(t + \delta_t)) + (c\delta_t/n) \cdot f(OPT) + (ct/n) \cdot (1 - \delta_t) \cdot f(OPT) \\ &\leq F(y(t + \delta_t)) + [c(t + \delta_t)/n] \cdot f(OPT) \ . \end{aligned}$$

Similarly, if $t \geq t_s$, then we get

$$g(t + \delta_t) = g(t) + \delta_t [e^{-t} \cdot f(OPT) + (e^{t_s - t} - e^{-t}) \cdot \max\{f(OPT) - f(Z \cup OPT), 0\} - g(t)]$$

19

$$\begin{aligned}
&= (1 - \delta_t)g(t) + \delta_t[e^{-t} \cdot f(OPT) + (e^{t_s - t} - e^{-t}) \cdot \max\{f(OPT) - f(Z \cup OPT), 0\}] \\
&\leq (1 - \delta_t)[F(y(t)) + (ct/n) \cdot f(OPT)] \\
&\quad + \delta_t[e^{-t} \cdot f(OPT) + (e^{t_s - t} - e^{-t}) \cdot \max\{f(OPT) - f(Z \cup OPT), 0\}] \\
&= F(y(t)) + \delta_t[e^{-t} \cdot f(OPT) + (e^{t_s - t} - e^{-t}) \cdot \max\{f(OPT) - f(Z \cup OPT), 0\} - F(y(t))] \\
&\quad + (ct/n) \cdot (1 - \delta_t) \cdot f(OPT) \\
&\leq F(y(t + \delta_t)) + (c\delta_t/n) \cdot f(OPT) + (ct/n) \cdot (1 - \delta_t) \cdot f(OPT) \\
&\leq F(y(t + \delta_t)) + [c(t + \delta_t)/n] \cdot f(OPT) \ . \qquad \square
\end{aligned}$$

It remains to find a closed-form expression that lower bounds $g(t)$ (and thus, also $F(y(t))$). Let $h_1(t) \colon [0, t_s] \to \mathbb{R}$ and $h_2(t) \colon [t_s, 1] \to \mathbb{R}$ be defined as follows.

$$h_1(t) = (1 - e^{-t}) \cdot f(OPT \setminus Z) - (1 - e^{-t} - te^{-t}) \cdot f(Z \cup OPT) \ ,$$

and

$$h_2(t) = e^{-t} \cdot \{(t - t_s) \cdot [f(OPT) + (e^{t_s} - 1) \cdot \max\{f(OPT) - f(OPT \cup Z), 0\}] + e^{t_s} \cdot h_1(t_s)\} \ .$$

Lemma A.12. *For every time $t \leq t_s$, $h_1(t) \leq g(t)$.*

Proof. The proof is by induction on t. For $t = 0$, $g(0) = 0 = (1 - e^0) \cdot f(OPT \setminus Z) - (1 - e^0 - 0 \cdot e^0) \cdot f(Z \cup OPT) = h_1(0)$. Assume now that the lemma holds for some $t < t_s$, and let us prove it holds also for $t + \delta_t$. By the induction hypothesis,

$$\begin{aligned}
h_1(t + \delta_t) &= h_1(t) + \int_t^{t + \delta_t} h'(\tau)d\tau \\
&= h_1(t) + \int_t^{t + \delta_t} \{e^{-\tau} \cdot f(OPT \setminus Z) - \tau e^{-\tau} \cdot f(Z \cup OPT)\}d\tau \\
&\leq h_1(t) + \delta_t \cdot \{e^{-t} \cdot f(OPT \setminus Z) - te^{-t} \cdot f(Z \cup OPT)\}d\tau \\
&= (1 - \delta_t)h_1(t) + \delta_t \cdot \{f(OPT \setminus Z) - (1 - e^{-t}) \cdot f(Z \cup OPT)\} \\
&\leq (1 - \delta_t)g(t) + \delta_t \cdot \{f(OPT \setminus Z) - (1 - e^{-t}) \cdot f(Z \cup OPT)\} = g(t + \delta_t) \ ,
\end{aligned}$$

where the first inequality holds since $e^{-\tau}$ is a decreasing function of τ and $\tau e^{-\tau}$ is an increasing function of τ in the range $\tau \in [0, 1]$. $\qquad \square$

Lemma A.13. *For every time $t_s \leq t \leq 1$, $h_2(t) \leq g(t)$.*

Proof. The proof is by induction on t. For $t = t_s$, by Lemma A.12, $h_2(t_s) = h_1(t_s) \leq g(t_s)$. Assume now that the lemma holds for some $t_s \leq t < 1$, and let us prove it holds also for $t + \delta_t$.

To avoid repeating complex expressions, let us denote $A = f(OPT) + (e^{t_s} - 1) \cdot \max\{f(OPT) - f(Z \cup OPT), 0\}$. Notice that A is independent of t. Moreover, using this notation we can rewrite $h_2(t)$ as $h_2(t) = e^{-t} \cdot \{(t - t_s) \cdot A + e^{t_s} \cdot h_1(t_s)\}$. Thus, for every $\tau \in (t_s, 1)$,

$$h'_2(\tau) = -e^{-\tau} \cdot \{(\tau - t_s) \cdot A + e^{t_s} \cdot h_1(t_s)\} + e^{-\tau} \cdot A = e^{-\tau} \cdot \{(1 - \tau + t_s) \cdot A - e^{t_s} \cdot h_1(t_s)\} \ .$$

The definition of A and the non-negativity of f imply immediately that $A \geq 0$. We would like to prove also that $t_s \cdot A - e^{t_s} \cdot h_1(t_s) \geq 0$. There are two cases to consider. First, if $f(OPT) \geq f(Z \cup OPT)$, then

$$t_s \cdot A - e^{t_s} \cdot h_1(t_s) = t_s \cdot f(OPT) + t_s(e^{t_s} - 1) \cdot \max\{f(OPT) - f(Z \cup OPT), 0\}$$

$$- (e^{t_s} - 1) \cdot f(OPT \setminus Z) + (e^{t_s} - 1 - t_s) \cdot f(Z \cup OPT)$$
$$\geq t_s e^{t_s} \cdot f(OPT) - t_s(e^{t_s} - 1) \cdot f(Z \cup OPT)$$
$$- (e^{t_s} - 1) \cdot f(OPT) + (e^{t_s} - 1 - t_s) \cdot f(Z \cup OPT)$$
$$= (t_s e^{t_s} - e^{t_s} + 1) \cdot [f(OPT) - f(Z \cup OPT)] \geq 0 \ .$$

where the inequality uses the fact that $f(OPT) \geq f(OPT \setminus Z)$ because of the down-closure of P. On the other hand, if $f(OPT) < f(Z \cup OPT)$, then

$$t_s \cdot A - e^{t_s} \cdot h_1(t_s) = t_s \cdot f(OPT) - (e^{t_s} - 1) \cdot f(OPT \setminus Z) + (e^{t_s} - 1 - t_s) \cdot f(Z \cup OPT)$$
$$\geq t_s \cdot f(OPT) - (e^{t_s} - 1) \cdot f(OPT) + (e^{t_s} - 1 - t_s) \cdot f(OPT) = 0 \ .$$

Using the above observations and the induction hypothesis, we can now get

$$h_2(t + \delta_t) = h_2(t) + \int_t^{t+\delta_t} h'(\tau)d\tau = h_2(t) + \int_t^{t+\delta_t} e^{-\tau} \cdot \{(1 - \tau + t_s) \cdot A - e^{t_s} \cdot h_1(t_s)\}d\tau$$
$$\leq h_2(t) + \delta_t \cdot e^{-t} \cdot \{(1 - t + t_s) \cdot A - e^{t_s} \cdot h_1(t_s)\} = (1 - \delta_t)h_2(t) + \delta_t \cdot e^{-t} \cdot A$$
$$\leq (1 - \delta_t)g(t) + \delta_t \cdot e^{-t} \cdot A = g(t + \delta_t) \ . \qquad \qquad \square$$

The last two lemmata give us the promised closed-form lower bound on $g(t)$, which can be used to lower bound the approximation ratio of Algorithm 3.

Corollary A.14. $\mathbb{E}[F(y(1))] \geq e^{t_s - 1} \cdot [(2 - t_s - e^{-t_s} - O(n^{-1})) \cdot f(OPT) - (1 - e^{-t_s}) \cdot f(Z \cap OPT) - (2 - t_s - 2e^{-t_s}) \cdot f(Z \cup OPT)]$.

Proof. By Lemma A.11, given $\mathcal{A}$,

$$F(y(1)) \geq g(1) - O(n^{-1}) \cdot f(OPT) \ .$$

By Lemma A.13,

$$g(1) \geq h_2(1)$$
$$= e^{-1} \cdot \{(1 - t_s) \cdot [f(OPT) + (e^{t_s} - 1) \cdot \max\{f(OPT) - f(Z \cup OPT), 0\}]$$
$$+ (e^{t_s} - 1) \cdot f(OPT \setminus Z) - (e^{t_s} - 1 - t_s) \cdot f(Z \cup OPT)\}$$
$$\geq e^{-1} \cdot \{(1 - t_s) \cdot [e^{t_s} \cdot f(OPT) - (e^{t_s} - 1) \cdot f(Z \cup OPT)]$$
$$+ (e^{t_s} - 1) \cdot [f(OPT) - f(Z \cap OPT)] - (e^{t_s} - 1 - t_s) \cdot f(Z \cup OPT)\}$$
$$= e^{t_s - 1} \cdot [(2 - t_s - e^{-t_s}) \cdot f(OPT) - (1 - e^{-t_s}) \cdot f(Z \cap OPT)$$
$$- (2 - t_s - 2e^{-t_s}) \cdot f(Z \cup OPT)] \ ,$$

where the second inequality holds since the submodularity and non-negativity of f imply

$$f(OPT \setminus Z) \geq f(OPT) + f(\varnothing) - f(Z \cap OPT) \geq f(OPT) - f(Z \cap OPT) \ .$$

Combining the above observations we get that, given $\mathcal{A}$,

$$F(y(1)) \geq e^{t_s - 1} \cdot [(2 - t_s - e^{-t_s} - O(n^{-1})) \cdot f(OPT) - (1 - e^{-t_s}) \cdot f(Z \cap OPT)$$
$$- (2 - t_s - 2e^{-t_s}) \cdot f(Z \cup OPT)] \ .$$

Since $F(y(1))$ is always non-negative, this implies, by the law of total expectation,

$$\begin{aligned}
\mathbb{E}[F(y(1))] \geq{} & \Pr[\mathcal{A}] \cdot \{e^{t_s-1} \cdot [(2 - t_s - e^{-t_s} - O(n^{-1})) \cdot f(OPT) - (1 - e^{-t_s}) \cdot f(Z \cap OPT) \\
& - (2 - t_s - 2e^{-t_s}) \cdot f(Z \cup OPT)]\} \\
\geq{} & \{e^{t_s-1} \cdot [(2 - t_s - e^{-t_s} - O(n^{-1})) \cdot f(OPT) - (1 - e^{-t_s}) \cdot f(Z \cap OPT) \\
& - (2 - t_s - 2e^{-t_s}) \cdot f(Z \cup OPT)]\} \\
& - \frac{1}{n} \cdot e^{t_s-1} \cdot (2 - t_s - e^{-t_s} - O(n^{-1})) \cdot f(OPT) \\
={} & e^{t_s-1} \cdot [(2 - t_s - e^{-t_s} - O(n^{-1})) \cdot f(OPT) - (1 - e^{-t_s}) \cdot f(Z \cap OPT) \\
& - (2 - t_s - 2e^{-t_s}) \cdot f(Z \cup OPT)] \enspace,
\end{aligned}$$

where the second inequality holds since $\Pr[\mathcal{A}] \geq 1 - n^{-1}$ by Corollary A.4. $\qquad\square$

Theorem 4.1 now follows immediately by combining Corollaries A.2 and A.14.