

Guess Free Maximization of Submodular and Linear Sums

Moran Feldman*

October 10, 2018

Abstract

We consider the problem of maximizing the sum of a monotone submodular function and a linear function subject to a general solvable polytope constraint. Recently, Sviridenko et al. [16] described an algorithm for this problem whose approximation guarantee is optimal in some intuitive and formal senses. Unfortunately, this algorithm involves a guessing step which makes it less clean and significantly affects its time complexity. In this work we describe a clean alternative algorithm that uses a novel weighting technique in order to avoid the problematic guessing step while keeping the same approximation guarantee as the algorithm of [16].

Keywords: submodular maximization, continuous greedy, curvature

*Department of Mathematics and Computer Science, The Open University of Israel. E-mail: moranfe@openu.ac.il

1 Introduction

The last decade has seen a surge of work on submodular maximization problems. Arguably, the main factor that allowed this surge was the invention of the multilinear relaxation for submodular maximization problems as well as algorithms for (approximately) solving this relaxation [4, 6, 8, 9, 11]. The invention of the multilinear relaxation was so influential because it allowed algorithms for submodular maximization to use the technique of first solving a relaxed version of the problem, and then rounding the fractional solution obtained. This technique is well-known, and it is often used in the design of algorithms for other kinds of problems; but, prior to the invention of the multilinear relaxation, it was not known how to apply it to submodular maximization problems.

An algorithm based on the above mentioned technique usually has two main components: a solver that (approximately) solves the relaxation and a rounding procedure. Historically, the first solver described for multilinear relaxations was the Continuous Greedy algorithm that solves such relaxations up to an approximation ratio of $1 - 1/e$ when the objective function is non-negative and monotone (in addition to being submodular) [6].¹ While the invention of continuous greedy was very significant, one can note that unlike standard solvers for more familiar relaxations such as LPs and SDPs, the approximation ratio of continuous greedy is quite far from 1. Unfortunately, a hardness result due to [14] implies that its approximation ratio cannot be improved in general.

This situation motivates the question of how well can one approximate multilinear relaxations whose objective includes both monotone submodular and linear components. Specifically, it is interesting to know whether the approximation ratio that can be achieved in such cases improves gradually as the linear component of the objective becomes more prominent. Recently, Sviridenko et al. [16] answered this question in the affirmative. More formally, they considered the following problem. Given a non-negative monotone submodular function $g: 2^{\mathcal{N}} \rightarrow \mathbb{R}_{\geq 0}$, a linear function ℓ and a solvable polytope $P \subseteq [0, 1]^{\mathcal{N}}$,² find a point $x \in P$ that approximately maximizes $G(x) + \ell(x)$, where G is the multilinear extension of G (see Section 2 for a definition). Sviridenko et al. [16] described a variant of continuous greedy that, given an instance of this problem, outputs a vector x obeying the inequality $G(x) + \ell(x) \geq (1 - e^{-1}) \cdot g(OPT) + \ell(OPT)$ up to a small error term, where OPT is an optimal integral solution for the problem.³

Intuitively, the result of Sviridenko et al. [16] is tight since it approximates the submodular component of the objective up an approximation ratio of $1 - 1/e$ and the linear component up to a ratio of 1. More formally, Sviridenko et al. [16] showed that their result is tight since it yields optimal approximation ratios for two problems of interest: maximizing a non-negative monotone submodular function with a bounded curvature subject to a matroid constraint, and minimizing a non-negative non-increasing supermodular function with a bounded curvature subject to the same kind of constraint.

1.1 Our Result

Despite being optimal in terms of its approximation guarantee, in the senses described above, the algorithm of Sviridenko et al. [16] suffers from a significant drawback. Namely, it is based on guessing the contribution of the linear component of the objective to the optimal solution, and this guessing step is quite problematic for the following reasons.

¹A set function $f: 2^{\mathcal{N}} \rightarrow \mathbb{R}$ is *monotone* if $f(S) \leq f(T)$ for every two sets $S \subseteq T \subseteq \mathcal{N}$ and *submodular* if $f(S \cup \{u\}) - f(S) \geq f(T \cup \{u\}) - f(T)$ for every two such sets and element $u \in \mathcal{N} \setminus T$.

²A polytope P is solvable if one can optimize linear functions subject to it.

³Technically, Sviridenko et al. [16] considered only the special case of the problem in which P is a matroid polytope, and designed two algorithms for this case. However, one of these algorithms (the continuous greedy based one) trivially extends to arbitrary solvable polytopes.

- The guessing is done by enumerating $\Theta(n\varepsilon^{-1}\log n)$ different possible values, and thus, increases the time complexity of the algorithm by this factor. Moreover, to guarantee that one of the enumerated values is a good enough guess, the set of values tried is constructed in a non-trivial way, which is then reflected in the complexity of the algorithm’s analysis.
- The original continuous greedy algorithm repeatedly maximizes linear functions subject to the constraint polytope P . While this computational step is quite slow in general, for many cases of interest (for example, when P is a matroid polytope) it can be implemented very efficiently. In contrast, the algorithm of [16] maximizes linear functions subject to the intersection of P with a polytope defined by the guessed value, which might be a very slow operation even when optimizing linear functions subject to P itself is fast. Moreover, various techniques have been described to speed up continuous greedy when it is applied to a matroid polytope [3, 5], and these techniques fail to apply to the algorithm of [16] because it considers the intersection of P with another polytope rather than P itself.

In this work we present a clean alternative algorithm that has the same approximation guarantee as the algorithm of Sviridenko et al. [16], but avoids the guessing step and all the problems resulting from it. In a nut shell, our algorithm is a modification of continuous greedy in which the weight assigned to each component of the objective function varies over time. At the beginning of an execution of the algorithm, the linear component has much more weight than the monotone submodular component, and over time their weights become equal. Intuitively, this kind of weighting makes sense because the standard analysis of continuous greedy for submodular functions uses a lower bound on the gain of the algorithm in its later steps which decreases if the algorithm has already made a significant gain in earlier steps. Thus, any gain from the submodular component of the objective that is obtained early in the algorithm’s execution is partially cancelled by the resulting decrease in the gain guaranteed in later steps of the execution. In contrast, gain obtained from the linear component of the objective in the same early steps of the execution does not suffer from such partial cancellation, and thus, should get more weight.

1.2 Additional Related Work

When the linear function is non-negative, its sum with the monotone submodular function is still monotone and submodular. Thus, in this case the work of Sviridenko et al. [16] can be viewed as improving the guarantee of continuous greedy in a special case. More recently, Soma and Yoshida [15] used an algorithm based on a similar technique to improve over the guarantee of continuous greedy in the more general case in which the monotone submodular objective can be decomposed into a monotone submodular component and a significant M^\natural -concave component. In an earlier work, Feldman et al. [11] took the complementing approach of using properties of the constraint polytope, rather than the objective function, to improve over the guarantee of continuous greedy. Specifically, they described a variant of continuous greedy, named Measured Continuous Greedy, which achieves an improved approximation ratio when the constraint is dense (in some sense).

As mentioned above, an algorithm that works by solving a relaxation and then rounding the solution has two main components: a relaxation solver and a rounding procedure. All the discussion up to this point was devoted to relaxation solvers because the current work is about such a solver and also because, unlike solvers, rounding procedures tend to be very problem specific. Nevertheless, there are a few more noticeable such procedures. A large portion of the work done so far on submodular maximization has been in the context of matroid constraints, for which there are two known rounding procedures that do not lose anything in the objective: Pipage Rounding [6] and Swap Rounding [7]. In another line of work, Chekuri et al. [8] designed a framework called

“contention resolution schemes” which yields a rounding procedure for every constraint that can be presented as the intersection of few simple constraints. Later works extended the contention resolution schemes framework into online and stochastic settings [1, 12, 13].

2 Preliminaries

In this section we describe the notation that we use and give a few relevant definitions. Using these definitions we then formally describe the guarantee of the algorithm we analyze.

Given a set S and an element u , we use $S + u$ and $S - u$ as shorthands for the union $S \cup \{u\}$ and the expression $S \setminus \{u\}$, respectively. If we are also given a set function f , then the marginal contribution of u to S with respect to f is denoted by $f(u \mid S) \triangleq f(S + u) - f(S)$. Notice that using this notation we get that a function $f: 2^{\mathcal{N}} \rightarrow \mathbb{R}$ is submodular if and only if for every two sets $S \subseteq T \subseteq \mathcal{N}$ and element $u \in \mathcal{N} \setminus T$ it holds that $f(u \mid S) \geq f(u \mid T)$. Occasionally, we are also interested in the marginal contribution of a set T to a set S with respect to f , which we denote by $f(T \mid S) \triangleq f(S \cup T) - f(S)$.

The *multilinear* extension of a set function $f: 2^{\mathcal{N}} \rightarrow \mathbb{R}$ is a function $F: [0, 1]^{\mathcal{N}} \rightarrow \mathbb{R}$ whose value for a vector $x \in [0, 1]^{\mathcal{N}}$ is defined as

$$F(x) = \mathbb{E}[f(\mathbf{R}(x))] = \sum_{S \subseteq \mathcal{N}} \left(\prod_{u \in S} x_u \right) \cdot \left(\prod_{u \in \mathcal{N} \setminus S} (1 - x_u) \right) \cdot f(S) ,$$

where $\mathbf{R}(x)$ is a random set containing every element $u \in \mathcal{N}$ with probability x_u , independently. One can observe that F is an extension of f in the sense that for every set $S \subseteq \mathcal{N}$, if we denote by $\mathbf{1}_S$ the characteristic vector of S , then it holds that $f(S) = F(\mathbf{1}_S)$. Additionally, observe that the rightmost hand side of F ’s definition implies that F is indeed a multilinear function, as suggested by its name. The multilinearity of F implies that for every vector $y \in [0, 1]^{\mathcal{N}}$ and element $u \in \mathcal{N}$ the partial derivative of F with respect to y is given by

$$\left. \frac{\partial F(x)}{\partial x_u} \right|_{x=y} = F(y \vee \mathbf{1}_{\{u\}}) - F(y \wedge \mathbf{1}_{\mathcal{N}-u}) = \mathbb{E}[f(u \mid \mathbf{R}(y) - u)] ,$$

where the vector operations \vee and \wedge represent coordinate-wise maximum and minimum, respectively.

A linear function is defined by a vector $\ell \in \mathbb{R}^{\mathcal{N}}$. We abuse notation and identify the vector ℓ with the linear function it defines. Accordingly, we denote the value of the function for a vector $x \in [0, 1]^{\mathcal{N}}$ as $\ell(x) \triangleq \ell \cdot x$. In a further abuse of notation, given a set $S \subseteq \mathcal{N}$, we use $\ell(S)$ as a shorthand for $\ell(\mathbf{1}_S) = \sum_{u \in S} \ell_u$.

An instance of the problem we consider in this work consists of a non-negative monotone submodular function $g: 2^{\mathcal{N}} \rightarrow \mathbb{R}_{\geq 0}$, a linear function $\ell: 2^{\mathcal{N}} \rightarrow \mathbb{R}$ and a solvable polytope $P \subseteq [0, 1]^{\mathcal{N}}$. We make the standard assumption that the submodular function g , whose description might be exponential in terms of the size n of \mathcal{N} , is accessible to the algorithm through a value oracle that given a set $S \subseteq \mathcal{N}$ returns $f(S)$. The objective of the problem is to find a vector $x \in P$ maximizing $G(x) + \ell(x)$, where G is the multilinear extension of g . The result that we prove for this problem is given by the next theorem. Let OPT be the set corresponding to an optimal integral solution for the problem, *i.e.*, $OPT = \arg \max_{S \subseteq \mathcal{N}, \mathbf{1}_S \in P} \{g(S) + \ell(S)\}$, and let $m = \max_{u \in \mathcal{N}} \{g(u \mid \emptyset)\}$.

Theorem 1. *There exists a polynomial time algorithm for the above problem that given a value $\varepsilon \in (0, 1)$ outputs a vector $x \in P$ such that with high probability $G(x) + \ell(x) \geq (1 - e^{-1}) \cdot g(OPT) + \ell(OPT) - O(\varepsilon) \cdot m$.*

Theorem 1 is very similar to the corresponding result of Sviridenko et al. [16]. However, there are two differences between the two. First, Sviridenko et al. [16] considered only matroid polytopes, for which there are known lossless rounding methods [6, 7], and thus, their result is stated in terms of sets rather than vectors. Second, the error term of [16] depends also on $\max_{u \in \mathcal{N}} |\ell_u|$, which is unnecessary for the analysis of our cleaner algorithm.

3 Algorithm

In this section we give a non-formal proof of Theorem 1. This proof demonstrates our new ideas, but uses some non-formal simplifications such as allowing a direct oracle access to the multilinear extension G of g 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.*, [6]), and for completeness, Appendix A includes a formal proof of Theorem 1 based on these techniques.

The algorithm we use for the non-formal proof of Theorem 1 is given as Algorithm 1. Like the original continuous greedy algorithm of [6], this algorithm grows a solution $y(t)$ over time. The solution starts as $\mathbf{1}_\emptyset$ at time $t = 0$, and the output of the algorithm is the solution at time $t = 1$. Our algorithm differs, however, from the original continuous greedy algorithm in the method used to determine the direction in which the solution is grown at every given time point. Specifically, our algorithm defines a weight vector $w(t)$ for every time $t \in [0, 1)$ based on the derivatives of the multilinear extension G . It then looks for a vector $z(t)$ in P maximizing a *weighted* combination of $w(t)$ with the linear function ℓ , and this vector $z(t)$ determines the direction in which the solution $y(t)$ is grown.

Algorithm 1: Distorted Continuous Greedy(g, ℓ, P)

```

1 Let  $y(0) \leftarrow \mathbf{1}_\emptyset$ .
2 foreach time  $t \in [0, 1)$  do
3   For each  $u \in \mathcal{N}$ , let  $w_u(t) \leftarrow \left. \frac{\partial G(x)}{\partial x_u} \right|_{x=y(t)}$ .
4   Let  $z(t)$  be the vector in  $P$  maximizing  $z(t) \cdot (e^{t-1} \cdot w(t) + \ell)$ .
5   Increase  $y(t)$  at a rate of  $\frac{dy(t)}{dt} = z(t)$ .
6 return  $y(1)$ .
```

We begin the analysis of Algorithm 1 by observing that its output is a vector in P .

Observation 3.1. $y(1) \in P$.

Proof. By definition, $z(t)$ is a vector in P for every time $t \in [0, 1)$. Hence, $y(1) = \int_0^1 z(t)dt$ is a convex combination of vectors in P , and thus, belongs to P by the convexity of P . \square

Let us consider now the function $\Phi(t) \triangleq e^{t-1} \cdot G(y(t)) + \ell(y(t))$. This function is a central component in our analysis of the approximation ratio of Algorithm 1. The following technical lemma gives an expression for the derivative of this important function.

Lemma 3.2.

$$\frac{d\Phi(t)}{dt} = e^{t-1} \cdot G(y(t)) + z(t) \cdot (e^{t-1} \cdot w(t) + \ell) .$$

Proof. By the chain rule,

$$\begin{aligned}
\frac{d\Phi(t)}{dt} &= e^{t-1} \cdot G(y(t)) + e^{t-1} \cdot \frac{dG(y(t))}{dt} + \frac{d\ell(y(t))}{dt} \\
&= e^{t-1} \cdot G(y(t)) + e^{t-1} \cdot \sum_{u \in \mathcal{N}} \frac{dy_u(t)}{dt} \cdot \frac{\partial G(x)}{\partial x_u} \Big|_{x=y(t)} + \sum_{u \in \mathcal{N}} \frac{dy_u(t)}{dt} \cdot \frac{\partial \ell(x)}{\partial x_u} \Big|_{x=y(t)} \\
&= e^{t-1} \cdot G(y(t)) + e^{t-1} \cdot \sum_{u \in \mathcal{N}} z_u(t) \cdot w_u(t) + \sum_{u \in \mathcal{N}} z_u(t) \cdot \ell_u \\
&= e^{t-1} \cdot G(y(t)) + z(t) \cdot (e^{t-1} \cdot w(t) + \ell) . \quad \square
\end{aligned}$$

The next lemma lower bounds the expression given by the last lemma for the derivative of $\Phi(t)$.

Lemma 3.3. *For every $t \in [0, 1)$,*

$$e^{t-1} \cdot G(y(t)) + z(t) \cdot (e^{t-1} \cdot w(t) + \ell) \geq e^{t-1} \cdot g(OPT) + \ell(OPT) .$$

Proof. Recall that $z(t)$ is chosen by Algorithm 1 as the vector in P maximizing $z(t) \cdot (e^{t-1} \cdot w(t) + \ell)$. Since $\mathbf{1}_{OPT} \in P$, we get

$$z(t) \cdot (e^{t-1} \cdot w(t) + \ell) \geq \mathbf{1}_{OPT} \cdot (e^{t-1} \cdot w(t) + \ell) = e^{t-1} \cdot \sum_{u \in OPT} w_u(t) + \ell(OPT) .$$

Note now that, by the multilinearity of G ,

$$\begin{aligned}
\sum_{u \in OPT} w_u(t) &= \sum_{u \in OPT} \frac{\partial G(x)}{\partial x_u} \Big|_{x=y(t)} = \sum_{u \in OPT} [G(y(t) \vee \mathbf{1}_{\{u\}}) - G(y(t) \wedge \mathbf{1}_{\mathcal{N}-u})] \\
&\geq \sum_{u \in OPT} [G(y(t) \vee \mathbf{1}_{\{u\}}) - G(y(t))] \geq G(y(t) \vee \mathbf{1}_{OPT}) - G(y(t)) \geq g(OPT) - G(y(t)) ,
\end{aligned}$$

where the first and last inequalities follow from the monotonicity of g , and the remaining inequality follows from its submodularity.

Combining the two above inequalities yields

$$z(t) \cdot (e^{t-1} \cdot w(t) + \ell) \geq e^{t-1} \cdot [g(OPT) - G(y(t))] + \ell(OPT) ,$$

and the lemma now follows by adding $e^{t-1} \cdot G(y(t))$ to both sides of this inequality. \square

We are now ready to prove Theorem 1.

Proof of Theorem 1. Observation 3.1 shows that $y(1) \in P$. Thus, to prove the theorem it only remains to prove $G(y(1)) + \ell(y(1)) \geq (1 - e^{-1}) \cdot g(OPT) + \ell(OPT)$.

Lemmata 3.2 and 3.3 prove together that

$$\frac{d\Phi(t)}{dt} = e^{t-1} \cdot G(y(t)) + z(t) \cdot (e^{t-1} \cdot w(t) + \ell) \geq e^{t-1} \cdot g(OPT) + \ell(OPT) .$$

Integrating both sides of this inequality from $t = 0$ to $t = 1$, we get

$$\Phi(1) - \Phi(0) \geq (1 - e^{-1}) \cdot g(OPT) + \ell(OPT) ,$$

and the theorem now follows by noticing that

$$\Phi(1) = G(y(1)) + \ell(y(1)) \quad \text{and} \quad \Phi(0) = e^{-1} \cdot G(y(0)) + \ell(y(0)) \geq 0$$

(the last inequality holds since g is non-negative and $\ell(y(0)) = \ell(\mathbf{1}_\emptyset) = 0$). \square

References

- [1] Marek Adamczyk and Michal Włodarczyk. Random order contention resolution schemes. *CoRR*, abs/1804.02584, 2018.
- [2] Noga Alon and Joel H. Spencer. *The Probabilistic Method*. Wiley Interscience, second edition, 2000.
- [3] Ashwinkumar Badanidiyuru and Jan Vondrák. Fast algorithms for maximizing submodular functions. In *SODA*, pages 1497–1514, 2014.
- [4] Niv Buchbinder and Moran Feldman. Constrained submodular maximization via a non-symmetric technique. *CoRR*, abs/1611.03253, 2016.
- [5] Niv Buchbinder, Moran Feldman, and Roy Schwartz. Comparing apples and oranges: Query trade-off in submodular maximization. *Math. Oper. Res.*, 42(2):308–329, 2017.
- [6] 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.
- [7] Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Dependent randomized rounding via exchange properties of combinatorial structures. In *FOCS*, pages 575–584, 2010.
- [8] 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.
- [9] Alina Ene and Huy L. Nguyen. Constrained submodular maximization: Beyond $1/e$. In *FOCS*, pages 248–257, 2016.
- [10] Moran Feldman. Maximizing symmetric submodular functions. *ACM Trans. Algorithms*, 13(3):39:1–39:36, 2017.
- [11] Moran Feldman, Joseph Naor, and Roy Schwartz. A unified continuous greedy algorithm for submodular maximization. In *FOCS*, pages 570–579, 2011.
- [12] Moran Feldman, Ola Svensson, and Rico Zenklusen. Online contention resolution schemes. In *SODA*, pages 1014–1033, 2016.
- [13] Anupam Gupta and Viswanath Nagarajan. A stochastic probing problem with applications. In *IPCO*, pages 205–216, 2013.
- [14] G. L. Nemhauser and L. A. Wolsey. Best algorithms for approximating the maximum of a submodular set function. *Math. Oper. Res.*, 3(3):177–188, 1978.
- [15] Tasuku Soma and Yuichi Yoshida. A new approximation guarantee for monotone submodular function maximization via discrete convexity. *CoRR*, abs/1709.02910, 2017.
- [16] Maxim Sviridenko, Jan Vondrák, and Justin Ward. Optimal approximation for submodular and supermodular optimization with bounded curvature. *Math. Oper. Res.*, 42(4), 2017.

A Formal Proof of Theorem 1

In this section we give a formal proof of Theorem 1. The algorithm that we use for this proof is given as Algorithm 2. Notice that this algorithm considers only discrete times that are integer multiples of a value δ chosen in a way guaranteeing two things. First, that $\delta \leq \min\{1/2, \varepsilon n^{-2}\}$, and second, that 1 is an integer multiple of δ .

Algorithm 2: Distorted Continuous Greedy – Formal(g, ℓ, P, ε)

```

1 Let  $y(0) \leftarrow \mathbf{1}_\emptyset$ ,  $t \leftarrow 0$  and  $\delta \leftarrow \lceil 2 + n^2/\varepsilon \rceil^{-1}$ .
2 while  $t < 1$  do
3   For each  $u \in \mathcal{N}$ , let  $w_u(t)$  be an estimate for  $\mathbb{E}[g(u \mid \mathbf{R}(y(t)) - u)]$  obtained by averaging
   the value of the expression within this expectation for  $r = \lceil -2n^2\varepsilon^{-2} \ln(\delta/n^2) \rceil$ 
   independent samples of  $\mathbf{R}(y(t))$ .
4   Let  $z(t)$  be the vector in  $P$  maximizing  $z(t) \cdot [(1 + \delta)^{(t-1)/\delta} \cdot w(t) + \ell]$ .
5   Let  $y(t + \delta) \leftarrow y(t) + \delta \cdot z(t)$ .
6   Update  $t \leftarrow t + \delta$ .
7 return  $y(1)$ .
```

Let T be the set of times considered by Algorithm 2, *i.e.*, $T = \{i\delta \mid i \in \mathbb{Z}, 0 \leq i < \delta^{-1}\}$. The following observation, which corresponds to Observation 3.1, shows that the output of Algorithm 2 is feasible.

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

Proof. By definition, $z(t)$ is a vector in P for every time $t \in T$. Observe also that $|T| = \delta^{-1}$. Hence, $y(1) = \sum_{t \in T} \delta \cdot z(t)$ is a convex combination of vectors in P , and thus, belongs to P by the convexity of P . \square

Next, we need to lower bound the probability that any of the estimates made by Algorithm 2 has a significant error. This is done by Lemma A.3, whose proof is based on the following known lemma.

Lemma A.2 (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 $S = X_1 + \dots + X_k$. Then, $\Pr[|S| > a] \leq 2e^{-a^2/2k}$.*

Let \mathcal{E} be the event that

$$|w_u(t) - \mathbb{E}[g(u \mid \mathbf{R}(y(t)) - u)]| \leq \frac{\varepsilon m}{n}$$

for every every element $u \in \mathcal{N}$ and time $t \in T$.

Lemma A.3. $\Pr[\mathcal{E}] \geq 1 - 2n^{-1}$, and hence, \mathcal{E} is a high probability event.

Proof. Consider an arbitrary element $u \in \mathcal{N}$ and time $t \in T$, and let us denote by R_i the i -th independent sample of $\mathbf{R}(y(t))$ used for calculating $w_u(t)$. We now define for every $1 \leq i \leq r$

$$X_i = \frac{[g(u \mid R_i - u)] - \mathbb{E}[g(u \mid \mathbf{R}(y(t)) - u)]}{m}.$$

Clearly $\mathbb{E}[X_i] = 0$ due to the linearity of the expectation. Additionally, note that $X_i \in [-1, 1]$ because the monotonicity of g guarantees that $g(u \mid R_i - u)$ and $\mathbb{E}[g(u \mid \mathbf{R}(y(t)) - u)]$ are both

non-negative, and the submodularity of g guarantees that these expressions are upper bounded by $f(u \mid \emptyset) \leq m$. Thus, by Lemma A.2,

$$\begin{aligned} \Pr\left[|w_u(t) - \mathbb{E}[g(u \mid \mathbf{R}(y(t)) - u)]| > \frac{\varepsilon m}{n}\right] &= \Pr\left[\frac{m}{r} \cdot \left|\sum_{i=1}^r X_i\right| > \frac{\varepsilon m}{n}\right] = \Pr\left[\left|\sum_{i=1}^r X_i\right| > \frac{r\varepsilon}{n}\right] \\ &\leq 2e^{-(r\varepsilon n^{-1})^2/2r} = 2e^{-r\varepsilon^2/(2n^2)} \leq 2e^{\ln(\delta/n^2)} = \frac{2\delta}{n^2} . \end{aligned}$$

Using the union bound, we now get that the probability that there is any pair of element $u \in \mathcal{N}$ and time $t \in T$ for which

$$|w_u(t) - \mathbb{E}[g(u \mid \mathbf{R}(y(t)) - u)]| > \frac{\varepsilon m}{n}$$

is at most

$$|\mathcal{N}| \cdot |T| \cdot \frac{2\delta}{n^2} = \frac{2}{n} . \quad \square$$

Let us define now $\Phi(t) \triangleq (1 + \delta)^{(t-1)/\delta} \cdot G(y(t)) + \ell(y(t))$. Lemma A.5 bounds the rate in which this expression increases as a function of t (and thus, can be viewed as a counterpart of Lemma 3.2). The following technical lemma is used in the proof of Lemma A.5. Since similar lemmata have been proved in other places (see, for example, [10, 11]), we defer the proof of this lemma to Appendix B.

Lemma A.4. *Given two vectors $y, y' \in [0, 1]^{\mathcal{N}}$ such that $0 \leq y'_u - y_u \leq \delta \leq 1$ and a non-negative monotone submodular function $f: 2^{\mathcal{N}} \rightarrow \mathbb{R}_{\geq 0}$ whose multilinear extension is F ,*

$$F(y') - F(y) \geq \sum_{u \in \mathcal{N}} (y'_u - y_u) \cdot \left. \frac{\partial F(x)}{\partial x_u} \right|_{x=y} - n^2 \delta^2 \cdot \max_{u \in \mathcal{N}} f(u \mid \emptyset) .$$

Lemma A.5. *If the event \mathcal{E} happens, then, for every time $t \in T$,*

$$\frac{\Phi(t + \delta) - \Phi(t)}{\delta} \geq (1 + \delta)^{(t-1)/\delta} \cdot G(y(t)) + z(t) \cdot [(1 + \delta)^{(t-1)/\delta} \cdot w(t) + \ell] - 2\varepsilon m .$$

Proof. Since $y(t + \delta) - y(t) = \delta z(t)$, and every coordinate of $z(t)$ is between 0 and 1, we get by Lemma A.4 that

$$\begin{aligned} G(y(t + \delta)) - G(y(t)) &\geq \sum_{u \in \mathcal{N}} (y(t + \delta) - y(t)) \cdot \left. \frac{\partial G(x)}{\partial x_u} \right|_{x=y(t)} - n^2 \delta^2 \cdot \max_{u \in \mathcal{N}} g(u \mid \emptyset) \quad (1) \\ &\geq \sum_{u \in \mathcal{N}} \delta z_u(t) \cdot \mathbb{E}[g(u \mid \mathbf{R}(y(t)) - u)] - \varepsilon \delta m \\ &\geq \sum_{u \in \mathcal{N}} \delta z_u(t) \cdot [w_u(t) - \varepsilon m/n] - \varepsilon \delta m \geq \delta z(t) \cdot w(t) - 2\varepsilon \delta m , \end{aligned}$$

where the second inequality hold since $\delta \leq \varepsilon n^{-2}$ by definition, and the third inequality holds since we assume that the event \mathcal{E} happened.

Using the linearity of ℓ and the definition of Φ , we now get

$$\begin{aligned} &\frac{\Phi(t + \delta) - \Phi(t)}{\delta} \\ &= \frac{[(1 + \delta)^{(t+\delta-1)/\delta} \cdot G(y(t + \delta)) - (1 + \delta)^{(t-1)/\delta} \cdot G(y(t))] + [\ell(y(t + \delta)) - \ell(y(t))]}{\delta} \\ &= \frac{(1 + \delta)^{(t-1)/\delta} \cdot [G(y(t + \delta)) - G(y(t))]}{\delta} + (1 + \delta)^{(t-1)/\delta} \cdot G(y(t + \delta)) + \ell \cdot z(t) . \end{aligned}$$

Plugging Inequality (1) and the inequality $G(y(t+\delta)) \geq G(y(t))$ (which holds due to monotonicity) into the last equality, we get

$$\frac{\Phi(t+\delta) - \Phi(t)}{\delta} \geq (1+\delta)^{(t-1)/\delta} \cdot G(y(t)) + z(t) \cdot [(1+\delta)^{(t-1)/\delta} \cdot w(t) + \ell] - 2\epsilon m \cdot (1+\delta)^{(t-1)/\delta} .$$

The lemma now follows from the last inequality by observing that $(1+\delta)^{(t-1)/\delta} \leq 1$ since $(t-1)/\delta \leq 0$. \square

The last lemma gives a lower bound on the increase in $\Phi(t)$ as a function of t . Unfortunately, this lower bound depends on a lot of entities (such as $z(t)$ and $w(t)$), and thus, it is difficult to use it. The following lemma allows us to simplify the lower bound.

Lemma A.6. *If the event \mathcal{E} happens, then $z(t) \cdot [(1+\delta)^{(t-1)/\delta} \cdot w(t) + \ell] \geq (1+\delta)^{(t-1)/\delta} \cdot [g(OPT) - G(y(t))] + \ell(OPT) - \epsilon m$.*

Proof. Recall that $z(t)$ is the vector in P maximizing $z(t) \cdot [(1+\delta)^{(t-1)/\delta} \cdot w(t) + \ell]$. Since $\mathbf{1}_{OPT} \in P$, we get

$$\begin{aligned} z(t) \cdot [(1+\delta)^{(t-1)/\delta} \cdot w(t) + \ell] &\geq \mathbf{1}_{OPT} \cdot [(1+\delta)^{(t-1)/\delta} \cdot w(t) + \ell] \\ &= (1+\delta)^{(t-1)/\delta} \cdot \sum_{u \in OPT} w(u) + \ell(OPT) \\ &\geq (1+\delta)^{(t-1)/\delta} \cdot \sum_{u \in OPT} \{\mathbb{E}[g(u \mid \mathbf{R}(y(t)) - u)] - \epsilon m/n\} + \ell(OPT) \\ &\geq (1+\delta)^{(t-1)/\delta} \cdot \sum_{u \in OPT} \mathbb{E}[g(u \mid \mathbf{R}(y(t)) - u)] + \ell(OPT) - \epsilon m , \end{aligned}$$

where the second inequality holds since we assume that the event \mathcal{E} happened. Observe now that the submodularity and monotonicity of f yield

$$\begin{aligned} \sum_{u \in OPT} \mathbb{E}[g(u \mid \mathbf{R}(y(t)) - u)] &\geq \sum_{u \in OPT} \mathbb{E}[g(u \mid \mathbf{R}(y(t)))] \geq \mathbb{E}[g(OPT \mid \mathbf{R}(y(t)))] \\ &= \mathbb{E}[g(OPT \cup \mathbf{R}(y(t))) - g(\mathbf{R}(y(t)))] \geq g(OPT) - G(y(t)) . \end{aligned}$$

The lemma now follows by combining the two above inequalities. \square

Combining the last two lemmata, we immediately get the following corollary, which is the promised simplified lower bound on the increase in $\Phi(t)$ as a function of t .

Corollary A.7.

$$\frac{\Phi(t+\delta) - \Phi(t)}{\delta} \geq (1+\delta)^{(t-1)/\delta} \cdot g(OPT) + \ell(OPT) - 3\epsilon m .$$

Using the last corollary, we can now get a lower bound on the value of $G(y(1)) + \ell(y(1))$ conditioned on the event \mathcal{E} .

Lemma A.8. *If the event \mathcal{E} happens, then $G(y(1)) + \ell(y(1)) \geq (1 - e^{-1}) \cdot g(OPT) + \ell(OPT) - 4\epsilon m$.*

Proof. Observe that

$$\begin{aligned}
G(y(1)) + \ell(y(1)) &= \Phi(1) = \Phi(0) + \sum_{t \in T} [\Phi(t + \delta) - \Phi(t)] \\
&\geq \Phi(0) + \sum_{t \in T} \left[\delta(1 + \delta)^{(t-1)/\delta} \cdot g(OPT) + \delta \cdot \ell(OPT) - 3\varepsilon\delta m \right] \\
&\geq \sum_{t \in T} \delta(1 + \delta)^{(t-1)/\delta} \cdot g(OPT) + \ell(OPT) - 3\varepsilon m,
\end{aligned} \tag{2}$$

where the first inequality holds due to Corollary A.7 and the second inequality holds since $|T| = \delta^{-1}$ and $\Phi(0) = (1 + \delta)^{-1/\delta} \cdot g(\emptyset) + \ell(\emptyset) = (1 + \delta)^{-1/\delta} \cdot g(\emptyset) \geq 0$ because ℓ is linear and g is non-negative.

We now need to lower bound the sum on the rightmost hand side of the last inequality. Notice that this sum can be presented as the sum of a geometrical series as follows.

$$\begin{aligned}
\sum_{t \in T} \delta(1 + \delta)^{(t-1)/\delta} &= \sum_{i=0}^{\delta^{-1}-1} \delta(1 + \delta)^{(i\delta-1)/\delta} = \delta(1 + \delta)^{-\delta^{-1}} \cdot \sum_{i=0}^{\delta^{-1}-1} (1 + \delta)^i \\
&= \delta(1 + \delta)^{-\delta^{-1}} \cdot \frac{1 - (1 + \delta)^{\delta^{-1}}}{1 - (1 + \delta)} = 1 - (1 + \delta)^{-\delta^{-1}} \\
&\geq 1 - e^{-1}(1 - \delta)^{-1} \geq 1 - e^{-1}(1 + 2\delta) \geq 1 - e^{-1} - \varepsilon/n,
\end{aligned}$$

where the first inequality holds since it is known that $(1 + 1/a)^a \geq e(1 - 1/a)$ for every $a \geq 1$ (and in particular for $a = \delta^{-1}$), the second inequality holds since $(1 - a)^{-1} \leq 1 + 2a$ for every $a \leq 1/2$, and the last inequality holds since $\delta \leq \varepsilon/n^2$. The lemma now follows by plugging the last inequality into Inequality (2) and observing that, by the submodularity of g ,

$$g(OPT) \leq g(\emptyset) + \sum_{u \in OPT} g(u \mid \emptyset) \leq g(\emptyset) + mn. \quad \square$$

We are now ready to prove Theorem 1.

Proof of Theorem 1. Observation A.1 shows that $y(1) \in P$. Additionally, Lemmata A.3 and A.8 show together that with high probability

$$G(y(1)) + \ell(y(1)) \geq (1 - e^{-1}) \cdot g(OPT) + \ell(OPT) - 4\varepsilon m = (1 - e^{-1}) \cdot g(OPT) + \ell(OPT) - O(\varepsilon) \cdot m. \quad \square$$

B Proof of Lemma A.4

In this section we prove Lemma A.4. Let us begin by recalling the lemma itself.

Lemma A.4. *Given two vectors $y, y' \in [0, 1]^{\mathcal{N}}$ such that $0 \leq y'_u - y_u \leq \delta \leq 1$ and a non-negative monotone submodular function $f: 2^{\mathcal{N}} \rightarrow \mathbb{R}_{\geq 0}$ whose multilinear extension is F ,*

$$F(y') - F(y) \geq \sum_{u \in \mathcal{N}} (y'_u - y_u) \cdot \left. \frac{\partial F(x)}{\partial x_u} \right|_{x=y} - n^2 \delta^2 \cdot \max_{u \in \mathcal{N}} f(u \mid \emptyset).$$

Let us denote the elements of \mathcal{N} by u_1, u_2, \dots, u_n in an arbitrary order. We define $y^{(i)}$ for every integer $0 \leq i \leq n$ as the vector in $[0, 1]^{\mathcal{N}}$ that agrees with y' on the coordinates 1 to i and with y on the remaining coordinates. Note that this definition implies, in particular, $y^{(0)} = y$ and $y^{(n)} = y'$. The next lemma bounds the amount by which the partial derivative $\frac{\partial F(x)}{\partial x_u}$ can differ between the points $x = y$ and $x = y'$.

Lemma B.1. For every integer $0 \leq i \leq n$ and element $u \in \mathcal{N}$,

$$\left. \frac{\partial F(x)}{\partial x_u} \right|_{x=y^{(i)}} \geq \left. \frac{\partial F(x)}{\partial x_u} \right|_{x=y} - n\delta \cdot f(u \mid \emptyset) .$$

Proof. For the sake of the proof, we assume that $\mathbf{R}(y^{(i)})$ is formed from $\mathbf{R}(y)$ using the following process. Every element of $\mathcal{N} \setminus \mathbf{R}(y)$ is added to a set D with probability of $1 - (1 - y_u^{(i)})/(1 - y_u)$. Then, $\mathbf{R}(y^{(i)})$ is chosen as $\mathbf{R}(y) \cup D$. Observe that every element $u \in \mathcal{N}$ gets into D with probability $y_u^{(i)} - y_u \leq \delta$, independently, and thus, $\mathbf{R}(y) \cup D$ indeed has the distribution that $\mathbf{R}(y^{(i)})$ should have.

Using the above definitions, we get

$$\begin{aligned} \left. \frac{\partial F(x)}{\partial x_u} \right|_{x=y^{(i)}} &= \mathbb{E}[f(u \mid \mathbf{R}(y^{(i)}) - u)] = \mathbb{E}[f(u \mid \mathbf{R}(y) \cup D - u)] \\ &\geq \Pr[D = \emptyset] \cdot \mathbb{E}[f(u \mid \mathbf{R}(y) - u) \mid D = \emptyset] , \end{aligned}$$

where the inequality follows from the law of total expectation and the monotonicity of f . Additionally, by the submodularity of f we also get

$$f(u \mid \mathbf{R}(y) - u) \leq f(u \mid \emptyset) .$$

Combining this inequality with the previous one yields

$$\begin{aligned} \left. \frac{\partial F(x)}{\partial x_u} \right|_{x=y^{(i)}} &+ \Pr[D \neq \emptyset] \cdot f(u \mid \emptyset) \\ &\geq \Pr[D = \emptyset] \cdot \mathbb{E}[f(u \mid \mathbf{R}(y) - u) \mid D = \emptyset] + \Pr[D \neq \emptyset] \cdot \mathbb{E}[f(u \mid \mathbf{R}(y) - u) \mid D \neq \emptyset] \\ &= \mathbb{E}[f(u \mid \mathbf{R}(y) - u)] = \left. \frac{\partial F(x)}{\partial x_u} \right|_{x=y} . \end{aligned}$$

One can verify that the last inequality will imply the lemma if we have an upper bound of $n\delta$ on $\Pr[D \neq \emptyset]$. Thus, all we are left to do is to prove this upper bound. Since elements belong to D with probability at most δ and independently,

$$\Pr[D \neq \emptyset] = 1 - \prod_{u \in \mathcal{N}} \Pr[u \notin D] \leq 1 - \prod_{u \in \mathcal{N}} (1 - \delta) = 1 - (1 - \delta)^n \leq n\delta . \quad \square$$

We are now ready to prove Lemma A.4.

Proof of Lemma A.4. Observe that for every integer $1 \leq i \leq n$ the vectors $y^{(i-1)}$ and $y^{(i)}$ differ only in coordinate i (in which they differ by $y'_u - y_u$). Recalling that $y^{(0)} = y$, $y^{(n)} = y'$ and F is multilinear, this observation yields

$$\begin{aligned} F(y') - F(y) &= \sum_{i=1}^n (y'_{u_i} - y_{u_i}) \cdot \left. \frac{\partial F(x)}{\partial x_{u_i}} \right|_{x=y^{(i-1)}} \\ &\geq \sum_{u \in \mathcal{N}} (y'_u - y_u) \cdot \left[\left. \frac{\partial F(x)}{\partial x_u} \right|_{x=y} - n\delta \cdot f(u \mid \emptyset) \right] \\ &\geq \sum_{u \in \mathcal{N}} (y'_u - y_u) \cdot \left. \frac{\partial F(x)}{\partial x_u} \right|_{x=y} - n^2 \delta^2 \cdot \max_{u \in \mathcal{N}} f(u \mid \emptyset) , \end{aligned}$$

where the first inequality follows from Lemma B.1 and the second inequality holds by the monotonicity of f and the fact that $y'_u - y_u \leq \delta$ for every $u \in \mathcal{N}$. \square