

Non-Monotone DR-Submodular Function Maximization

Tasuku Soma

The University of Tokyo
tasuku_soma@mist.i.u-tokyo.ac.jp

Yuichi Yoshida

National Institute of Informatics, *and*
Preferred Infrastructure, Inc.
yyoshida@nii.ac.jp

Abstract

We consider non-monotone DR-submodular function maximization, where DR-submodularity (diminishing return submodularity) is an extension of submodularity for functions over the integer lattice based on the concept of the diminishing return property. Maximizing non-monotone DR-submodular functions has many applications in machine learning that cannot be captured by submodular set functions. In this paper, we present a $\frac{1}{2+\epsilon}$ -approximation algorithm with a running time of roughly $O(\frac{n}{\epsilon} \log^2 B)$, where n is the size of the ground set, B is the maximum value of a coordinate, and $\epsilon > 0$ is a parameter. The approximation ratio is almost tight and the dependency of running time on B is exponentially smaller than the naive greedy algorithm. Experiments on synthetic and real-world datasets demonstrate that our algorithm outputs almost the best solution compared to other baseline algorithms, whereas its running time is several orders of magnitude faster.

Introduction

Submodular functions have played a key role in various tasks in machine learning, statistics, social science, and economics. A set function $f : 2^E \rightarrow \mathbb{R}$ with a ground set E is *submodular* if

$$f(X \cup \{e\}) - f(X) \geq f(Y \cup \{e\}) - f(Y)$$

for arbitrary sets $X, Y \subseteq E$ with $X \subseteq Y$, and an element $e \in E \setminus Y$. The importance and usefulness of submodularity in these areas are due to the fact that submodular functions naturally capture the *diminishing return property*. Various important functions in these areas such as the entropy function, coverage function, and utility functions satisfy this property. See, e.g., (Krause and Golovin 2014; Fujishige 2005).

Recently, *maximizing* (non-monotone) submodular functions has attracted particular interest in the machine learning community. In contrast to *minimizing* submodular functions, which can be done in polynomial time, maximizing submodular functions is NP-hard in general. However, we can achieve a constant factor approximation for various settings. Notably, (Buchbinder et al. 2012) presented a very elegant *double greedy* algorithm for (unconstrained) submodular function maximization, which was the first algorithm

achieving $\frac{1}{2}$ -approximation, and this approximation ratio is tight (Feige, Mirrokni, and Vondrak 2011). Applications of non-monotone submodular function maximization include efficient sensor placement (Krause, Singh, and Guestrin 2008), privacy in online services (Krause and Horvitz 2008), and maximum entropy sampling (Ko, Lee, and Queyranne 1995).

The models and applications mentioned so far are built upon submodular *set functions*. Although set functions are fairly powerful for describing problems such as variable selection, we sometimes face difficult situations that cannot be cast with set functions. For example, in the budget allocation problem (Alon, Gamzu, and Tennenholtz 2012), we would like to decide how much budget should be set aside for each ad source, rather than whether we use the ad source or not. A similar issue arises when we consider models allowing multiple choices of an element in the ground set.

To deal with such situations, several generalizations of submodularity have been proposed. (Soma et al. 2014) devised a general framework for maximizing *monotone submodular functions on the integer lattice*, and showed that the budget allocation problem and its variant fall into this framework. In their framework, functions are defined over the integer lattice \mathbb{Z}_+^E and therefore effectively represent discrete allocations of budget. Regarding the original motivation for the diminishing return property, one can naturally consider its generalization to the integer lattice: a function $f : \mathbb{Z}_+^E \rightarrow \mathbb{R}$ satisfying

$$f(\mathbf{x} + \boldsymbol{\chi}_e) - f(\mathbf{x}) \geq f(\mathbf{y} + \boldsymbol{\chi}_e) - f(\mathbf{y})$$

for $\mathbf{x} \leq \mathbf{y}$ and $e \in E$, where $\boldsymbol{\chi}_e \in \mathbb{R}^E$ is the vector with $\boldsymbol{\chi}_e(e) = 1$ and $\boldsymbol{\chi}_e(a) = 0$ for every $a \neq e$. Such functions are called *diminishing return submodular (DR-submodular) functions* (Soma and Yoshida 2015) or *coordinate-wise concave submodular functions* (Milgrom and Strulovici 2009). DR-submodular functions have found various applications in generalized sensor placement (Soma and Yoshida 2015) and (a natural special case of) the budget allocation problem (Soma et al. 2014).

As a related notion, a function is said to be *lattice submodular* if

$$f(\mathbf{x}) + f(\mathbf{y}) \geq f(\mathbf{x} \vee \mathbf{y}) + f(\mathbf{x} \wedge \mathbf{y})$$

for arbitrary \mathbf{x} and \mathbf{y} , where \vee and \wedge are coordinate-wise max and min, respectively. Note that DR-submodularity is

stronger than *lattice submodularity* in general (see, e.g., (Soma et al. 2014)). Nevertheless, we consider the DR-submodularity to be a “natural definition” of submodularity, at least for the applications mentioned so far, because the diminishing return property is crucial in these real-world scenarios.

Our contributions

We design a novel polynomial-time approximation algorithm for maximizing (non-monotone) DR-submodular functions. More precisely, we consider the optimization problem

$$\begin{aligned} & \text{maximize} && f(\mathbf{x}) \\ & \text{subject to} && \mathbf{0} \leq \mathbf{x} \leq \mathbf{B}, \end{aligned} \quad (1)$$

where $f : \mathbb{Z}_+^E \rightarrow \mathbb{R}_+$ is a non-negative DR-submodular function, $\mathbf{0}$ is the zero vector, and $\mathbf{B} \in \mathbb{Z}_+^E$ is a vector representing the maximum values for each coordinate. When \mathbf{B} is the all-ones vector, this is equivalent to the original (unconstrained) submodular function maximization. We assume that f is given as an evaluation oracle; when we specify $\mathbf{x} \in \mathbb{Z}_+^E$, the oracle returns the value of $f(\mathbf{x})$.

Our algorithm achieves $\frac{1}{2+\epsilon}$ -approximation for any constant $\epsilon > 0$ in $O(\frac{|E|}{\epsilon} \cdot \log(\frac{\Delta}{\delta}) \log B \cdot (\theta + \log B))$ time, where δ and Δ are the minimum positive marginal gain and maximum positive values, respectively, of f , $B = \|\mathbf{B}\|_\infty := \max_{e \in E} \mathbf{B}(e)$, and θ is the running time of evaluating (the oracle for) f . To our knowledge, this is the first polynomial-time algorithm achieving (roughly) $\frac{1}{2}$ -approximation.

We also conduct numerical experiments on the revenue maximization problem using real-world networks. The experimental results show that the solution quality of our algorithm is comparable to other algorithms. Furthermore, our algorithm runs several orders of magnitude faster than other algorithms when B is large.

DR-submodularity is necessary for obtaining polynomial-time algorithms with a meaningful approximation guarantee; if f is only lattice submodular, then we cannot obtain constant-approximation in polynomial time. To see this, it suffices to observe that an *arbitrary* univariate function is lattice submodular, and therefore finding an (approximate) maximum value must invoke $O(B)$ queries. We note that representing an integer B requires $\lceil \log_2 B \rceil$ bits. Hence, the running time of $O(B)$ is pseudopolynomial rather than polynomial.

Fast simulation of the double greedy algorithm

Naturally, one can reduce the problem (1) to maximization of a submodular set function by simply duplicating each element e in the ground set into $\mathbf{B}(e)$ distinct copies and defining a set function over the set of all the copies. One can then run the double greedy algorithm (Buchbinder et al. 2012) to obtain $\frac{1}{2}$ -approximation. This reduction is simple but has one large drawback; the size of the new ground set is $\sum_{e \in E} \mathbf{B}(e)$, a pseudopolynomial in B . Therefore, this naive double greedy algorithm does not scale to a situation where B is large.

For scalability, we need an additional trick that reduces the pseudo-polynomial running time to a polynomial one. In *monotone* submodular function maximization on the integer lattice, (Soma and Yoshida 2015; 2016) provide such a speedup trick, which effectively combines the *decreasing threshold technique* (Badanidiyuru and Vondrák 2014) with binary search. However, a similar technique does not apply to our setting, because the double greedy algorithm works differently from (single) greedy algorithms for monotone submodular function maximization. The double greedy algorithm examines each element in a *fixed* order and marginal gains are used to decide whether to include the element or not. In contrast, the greedy algorithm chooses each element in *decreasing* order of marginal gains, and this property is crucial for the decreasing threshold technique.

We resolve this issue by splitting the set of all marginal gains into polynomially many small intervals. For each interval, we approximately execute multiple steps of the double greedy algorithm at once, as long as the marginal gains remain in the interval. Because the marginal gains do not change (much) within the interval, this simulation can be done with polynomially many queries and polynomial-time overhead. To our knowledge, this speedup technique is not known in the literature and is therefore of more general interest.

Very recently, (Ene and Nguyen 2016) pointed out that a DR-submodular function $f : \{0, 1, \dots, B\}^E \rightarrow \mathbb{R}_+$ can be expressed as a submodular set function g over a polynomial-sized ground set, which turns out to be $E \times \{0, 1, \dots, k-1\}$, where $k = \lceil \log_2(B+1) \rceil$. Their idea is representing $\mathbf{x}(e)$ in binary form for each $e \in E$, and bits in the binary representations form the new ground set. We may want to apply the double greedy algorithm to g in order to get a polynomial-time approximation algorithm. However, this strategy has the following two drawbacks: (i) The value of $g(E \times \{0, 1, \dots, k-1\})$ is defined as $f(\mathbf{x})$, where $\mathbf{x}(e) = 2^k - 1$ for every $e \in E$. This means that we have to extend the domain of f . (ii) More crucially, the double greedy algorithm on g may return a large set such as $E \times \{0, 1, \dots, k-1\}$ whose corresponding vector $\mathbf{x} \in \mathbb{Z}_+^E$ may violate the constraint $\mathbf{x} \leq \mathbf{B}$. Although we can resolve these issues by introducing a knapsack constraint, it is not a practical solution because existing algorithms for knapsack constraints (Lee et al. 2009; Chekuri, Vondrák, and Zenklusen 2014) are slow and have worse approximation ratios than $1/2$.

Notations For an integer $n \in \mathbb{N}$, $[n]$ denotes the set $\{1, \dots, n\}$. For vectors $\mathbf{x}, \mathbf{y} \in \mathbb{Z}^E$, we define $f(\mathbf{x} \mid \mathbf{y}) := f(\mathbf{x} + \mathbf{y}) - f(\mathbf{y})$. The ℓ_1 -norm and ℓ_∞ -norm of a vector $\mathbf{x} \in \mathbb{Z}^E$ are defined as $\|\mathbf{x}\|_1 := \sum_{e \in E} |\mathbf{x}(e)|$ and $\|\mathbf{x}\|_\infty := \max_{e \in E} |\mathbf{x}(e)|$, respectively.

Related work

As mentioned above, there have been many efforts to maximize submodular functions on the integer lattice. Perhaps the work most related to our interest is (Gottschalk and Peis 2015), in which the authors considered maximizing

lattice submodular functions over the bounded integer lattice and designed $\frac{1}{3}$ -approximation pseudopolynomial-time algorithm. Their algorithm was also based on the double greedy algorithm, but does not include a speeding up technique, as proposed in this paper.

In addition there are several studies on the *constrained* maximization of submodular functions (Feige, Mirrokni, and Vondrak 2011; Buchbinder et al. 2014; Buchbinder and Feldman 2016), although we focus on the unconstrained case. Many algorithms for maximizing submodular functions are randomized, but a very recent work (Buchbinder and Feldman 2016) devised a derandomized version of the double greedy algorithm. (Gotovos, Karbasi, and Krause 2015) considered maximizing non-monotone submodular functions in the *adaptive* setting, a concept introduced in (Golovin and Krause 2011).

A continuous analogue of DR-submodular functions is considered in (Bian et al. 2016).

Algorithms

In this section, we present a polynomial-time approximation algorithm for maximizing (non-monotone) DR-submodular functions. We first explain a simple adaption of the double greedy algorithm for (set) submodular functions to our setting, which runs in pseudopolynomial time. Then, we show how to achieve a polynomial number of oracle calls. Finally, we provide an algorithm with a polynomial running time (details are placed in the full version).

Pseudopolynomial-time algorithm

Algorithm 1 is an immediate extension of the double greedy algorithm for maximizing submodular (set) functions (Buchbinder et al. 2012) to our setting. We start with $\mathbf{x} = \mathbf{0}$ and $\mathbf{y} = \mathbf{B}$, and then for each $e \in E$, we tighten the gap between $\mathbf{x}(e)$ and $\mathbf{y}(e)$ until they become exactly the same. Let $\alpha = f(\chi_e | \mathbf{x})$ and $\beta = f(-\chi_e | \mathbf{y})$. We note that

$$\alpha + \beta = f(\mathbf{x} + \chi_e) - f(\mathbf{x}) - (f(\mathbf{y}) - f(\mathbf{y} - \chi_e)) \geq 0$$

holds from the DR-submodularity of f . Hence, if $\beta < 0$, then $\alpha > 0$ must hold, and we increase $\mathbf{x}(e)$ by one. Similarly, if $\alpha < 0$, then $\beta > 0$ must hold, and we decrease $\mathbf{y}(e)$ by one. When both of them are non-negative, we increase $\mathbf{x}(e)$ by one with probability $\frac{\alpha}{\alpha+\beta}$, or decrease $\mathbf{y}(e)$ by one with the complement probability $\frac{\beta}{\alpha+\beta}$.

Theorem 1. *Algorithm 1 is a $\frac{1}{2}$ -approximation algorithm for (1) with time complexity $O(\|\mathbf{B}\|_1 \cdot \theta + \|\mathbf{B}\|_1)$, where θ is the running time of evaluating f .*

We omit the proof as it is a simple modification of the analysis of the original algorithm.

Algorithm with polynomially many oracle calls

In this section, we present an algorithm with a polynomial number of oracle calls.

Our strategy is to simulate Algorithm 1 without evaluating the input function f many times. A key observation is that, at Line 4 of Algorithm 1, we do not need to know the exact

Algorithm 1 Pseudopolynomial-time algorithm

Input: $f : \mathbb{Z}_+^E \rightarrow \mathbb{R}_+$, $\mathbf{B} \in \mathbb{Z}_+^E$.

Assumption: f is DR-submodular.

```

1:  $\mathbf{x} \leftarrow \mathbf{0}$ ,  $\mathbf{y} \leftarrow \mathbf{B}$ .
2: for  $e \in E$  do
3:   while  $\mathbf{x}(e) < \mathbf{y}(e)$  do
4:      $\alpha \leftarrow f(\chi_e | \mathbf{x})$  and  $\beta \leftarrow f(-\chi_e | \mathbf{y})$ .
5:     if  $\beta < 0$  then
6:        $\mathbf{x}(e) \leftarrow \mathbf{x}(e) + 1$ .
7:     else if  $\alpha < 0$  then
8:        $\mathbf{y}(e) \leftarrow \mathbf{y}(e) - 1$ .
9:     else
10:       $\mathbf{x}(e) \leftarrow \mathbf{x}(e) + 1$  with probability  $\frac{\alpha}{\alpha+\beta}$  and
       $\mathbf{y}(e) \leftarrow \mathbf{y}(e) - 1$  with the complement probability  $\frac{\beta}{\alpha+\beta}$ . If  $\alpha = \beta = 0$ , we assume  $\frac{\alpha}{\alpha+\beta} = 1$ .
11:    end if
12:  end while
13: end for
14: return  $\mathbf{x}$ .
```

value of $f(\chi_e | \mathbf{x})$ and $f(-\chi_e | \mathbf{y})$; good approximations to them are sufficient to achieve an approximation guarantee close to $\frac{1}{2}$. To exploit this observation, we first design an algorithm that outputs (sketches of) approximations to the functions $g(b) := f(\chi_e | \mathbf{x} + b\chi_e)$ and $h(b) := f(-\chi_e | \mathbf{y} - b\chi_e)$. Note that g and h are non-increasing functions in b because of the DR-submodularity of f .

To illustrate this idea, let us consider a non-increasing function $\phi : \{0, 1, \dots, B-1\} \rightarrow \mathbb{R}$ and suppose that ϕ is non-negative (ϕ is either g or h later on). Let δ and Δ be the minimum and the maximum positive values of ϕ , respectively. Then, for each $\delta \leq \tau \leq \Delta$ of the form $\delta(1+\epsilon)^k$, we find the minimum b_τ such that $\phi(b_\tau) < \tau$ (we regard $\phi(B) = -\infty$). From the non-increasing property of ϕ , we then have $\phi(b) \geq \tau$ for any $b < b_\tau$. Using the set of pairs $\{(\tau, b_\tau)\}_\tau$, we can obtain a good approximation to ϕ . The details are provided in Algorithm 2.

Lemma 2. *For any $\phi : \{0, 1, \dots, B-1\} \rightarrow \mathbb{R}$ and $\epsilon > 0$, Algorithm 2 outputs a set of pairs $\{(b_\tau, \tau)\}_\tau$ from which, for any $b \in \{0, 1, \dots, B-1\}$, we can reconstruct a value v in $O(\log B)$ time such that $v \leq \phi(b) < (1+\epsilon)v$ if $\phi(b) > 0$ and $v = 0$ otherwise. The time complexity of Algorithm 2 is $O(\frac{1}{\epsilon} \log(\frac{\Delta}{\delta}) \log B \cdot \theta)$ if ϕ has a positive value, where δ and Δ are the minimum and maximum positive values of ϕ , respectively, and θ is the running time of evaluating ϕ , and is $O(\log B \cdot \theta)$ otherwise.*

Proof. Let $S = \{(b_\tau, \tau)\}_\tau$ be the set of pairs output by Algorithm 2. Our reconstruction algorithm is as follows: Given $b \in \{0, 1, \dots, B-1\}$, let (b_{τ^*}, τ^*) be the pair with the minimum b_{τ^*} , where $b < b_{\tau^*}$. Note that such a b_{τ^*} always exists because a pair of the form (B, \cdot) is always added to S . We then output τ^* . The time complexity of this reconstruction algorithm is clearly $O(\log B)$.

We now show the correctness of the reconstruction algorithm. If $\phi(b) > 0$, then, in particular, we have $\phi(b) \geq \delta$. Then, τ^* is the maximum value of the form $\delta(1+\epsilon)^k$ at most

Algorithm 2 Sketching subroutine for Algorithm 3

Input: $\phi : \{0, 1, \dots, B-1\} \rightarrow \mathbb{R}, \epsilon > 0$.
Assumption: ϕ is non-increasing.
1: $S \leftarrow \emptyset$. We regard $\phi(B) = -\infty$.
2: Find the minimum $b_0 \in \{0, 1, \dots, B\}$ with $\phi(b_0) \leq 0$ by binary search.
3: **if** $b_0 \geq 1$ **then** # ϕ has a positive value.
4: $\Delta \leftarrow \phi(0)$ and $\delta \leftarrow \phi(b_0 - 1)$.
5: **for** ($\tau \leftarrow \delta; \tau \leq \Delta; \tau \leftarrow (1 + \epsilon)\tau$) **do**
6: Find the minimum $b_\tau \in \{0, 1, \dots, B\}$ with $\phi(b_\tau) < \tau$ by binary search.
7: $S \leftarrow S \cup \{(b_\tau, \tau)\}$
8: **end for**
9: **if** $b_\delta \neq B$ **then**
10: $S \leftarrow S \cup \{(B, 0)\}$.
11: **end if**
12: **else** # ϕ is non-positive.
13: $S \leftarrow S \cup \{(B, 0)\}$.
14: **end if**
15: **return** S .

$\phi(b)$. Hence, we have $\tau^* \leq \phi(b) < (1 + \epsilon)\tau^*$. If $\phi(b) \leq 0$, $(b_{\tau^*}, \tau^*) = (B, 0)$ and we output zero.

Finally, we analyze the time complexity of Algorithm 2. Each binary search requires $O(\log B)$ time. The number of binary searches performed is $O(\log_{1+\epsilon} \frac{\Delta}{\delta}) = O(\frac{1}{\epsilon} \log \frac{\Delta}{\delta})$ when ϕ has a positive value and 1 when ϕ is non-positive. Hence, we have the desired time complexity. \square

We can regard Algorithm 2 as providing a value oracle for a function $\tilde{\phi} : \{0, 1, \dots, B-1\} \rightarrow \mathbb{R}_+$ that is an approximation to the input function $\phi : \{0, 1, \dots, B-1\} \rightarrow \mathbb{R}$.

We now describe our algorithm for maximizing DR-submodular functions. The basic idea is similar to Algorithm 1, but when we need $f(\chi_e | \mathbf{x})$ and $f(-\chi_e | \mathbf{y})$, we use approximations to them instead. Let α and β be approximations to $f(\chi_e | \mathbf{x})$ and $f(-\chi_e | \mathbf{y})$, respectively, obtained by Algorithm 2. Then, we increase $\mathbf{x}(e)$ by one with probability $\frac{\alpha}{\alpha+\beta}$ and decrease $\mathbf{y}(e)$ by one with the complement probability $\frac{\beta}{\alpha+\beta}$. The details are given in Algorithm 3.

We now analyze Algorithm 3. An *iteration* refers to an iteration in the while loop from Line 5. We have $\|\mathbf{B}\|_1$ iterations in total. For $k \in \{1, \dots, \|\mathbf{B}\|_1\}$, let \mathbf{x}_k and \mathbf{y}_k be \mathbf{x} and \mathbf{y} , respectively, right after the k th iteration. Note that $\mathbf{x}_{\|\mathbf{B}\|_1} = \mathbf{y}_{\|\mathbf{B}\|_1}$ is the output of the algorithm. We define $\mathbf{x}_0 = \mathbf{0}$ and $\mathbf{y}_0 = \mathbf{B}$ for convenience.

Let \mathbf{o} be an optimal solution. For $k \in \{0, 1, \dots, \|\mathbf{B}\|_1\}$, we then define $\mathbf{o}_k = (\mathbf{o} \vee \mathbf{x}_k) \wedge \mathbf{y}_k$. Note that $\mathbf{o}_0 = \mathbf{o}$ holds and $\mathbf{o}_{\|\mathbf{B}\|_1}$ equals the output of the algorithm. We have the following key lemma.

Lemma 3. For every $k \in [\|\mathbf{B}\|_1]$, we have

$$\begin{aligned} & \mathbf{E}[f(\mathbf{o}_{k-1}) - f(\mathbf{o}_k)] \\ & \leq \frac{1+\epsilon}{2} \mathbf{E}[f(\mathbf{x}_k) - f(\mathbf{x}_{k-1}) + f(\mathbf{y}_k) - f(\mathbf{y}_{k-1})] \quad (2) \end{aligned}$$

Proof. Fix $k \in [\|\mathbf{B}\|_1]$ and let e be the element of interest in the k th iteration. Let α and β be the values in Line 6 in

Algorithm 3 Algorithm with polynomially many queries

Input: $f : \mathbb{Z}_+^E \rightarrow \mathbb{R}_+, \mathbf{B} \in \mathbb{Z}_+^E, \epsilon > 0$.
Assumption: f is DR-submodular.
1: $\mathbf{x} \leftarrow \mathbf{0}, \mathbf{y} \leftarrow \mathbf{B}$.
2: **for** $e \in E$ **do**
3: Define $g, h : \{0, 1, \dots, B-1\} \rightarrow \mathbb{R}$ as $g(b) = f(\chi_e | \mathbf{x} + b\chi_e)$ and $h(b) = f(-\chi_e | \mathbf{y} - b\chi_e)$.
4: Let \tilde{g} and \tilde{h} be approximations to g and h , respectively, given by Algorithm 2.
5: **while** $\mathbf{x}(e) < \mathbf{y}(e)$ **do**
6: $\alpha \leftarrow \tilde{g}(\mathbf{x}(e))$ and $\beta \leftarrow \tilde{h}(\mathbf{B}(e) - \mathbf{y}(e))$.
7: $\mathbf{x}(e) \leftarrow \mathbf{x}(e) + 1$ with probability $\frac{\alpha}{\alpha+\beta}$ and $\mathbf{y}(e) \leftarrow \mathbf{y}(e) - 1$ with the complement probability $\frac{\beta}{\alpha+\beta}$. If $\alpha = \beta = 0$, we assume $\frac{\alpha}{\alpha+\beta} = 1$.
8: **end while**
9: **end for**
10: **return** \mathbf{x} .

the k th iteration. We then have

$$\begin{aligned} & \mathbf{E}[f(\mathbf{x}_k) - f(\mathbf{x}_{k-1}) + f(\mathbf{y}_k) - f(\mathbf{y}_{k-1})] \\ & = \frac{\alpha}{\alpha+\beta} f(\chi_e | \mathbf{x}_{k-1}) + \frac{\beta}{\alpha+\beta} f(-\chi_e | \mathbf{y}_{k-1}) \\ & \geq \frac{\alpha}{\alpha+\beta} \alpha + \frac{\beta}{\alpha+\beta} \beta = \frac{\alpha^2 + \beta^2}{\alpha + \beta}, \quad (3) \end{aligned}$$

where we use the guarantee in Lemma 2 in the inequality.

We next establish an upper bound of $\mathbf{E}[f(\mathbf{o}_{k-1}) - f(\mathbf{o}_k)]$. As $\mathbf{o}_k = (\mathbf{o} \vee \mathbf{x}_k) \wedge \mathbf{y}_k$, conditioned on a fixed \mathbf{o}_{k-1} , we obtain

$$\begin{aligned} & \mathbf{E}[f(\mathbf{o}_{k-1}) - f(\mathbf{o}_k)] \\ & = \frac{\alpha}{\alpha+\beta} \left(f(\mathbf{o}_{k-1}) - f(\mathbf{o}_{k-1} \vee \mathbf{x}_k(e)\chi_e) \right) \\ & \quad + \frac{\beta}{\alpha+\beta} \left(f(\mathbf{o}_{k-1}) - f(\mathbf{o}_{k-1} \wedge \mathbf{y}_k(e)\chi_e) \right). \quad (4) \end{aligned}$$

Claim 4. (4) $\leq \frac{(1+\epsilon)\alpha\beta}{\alpha+\beta}$.

Proof. We show this claim by considering the following three cases.

If $\mathbf{x}_k(e) \leq \mathbf{o}_{k-1}(e) \leq \mathbf{y}_k(e)$, then (4) is zero.

If $\mathbf{o}_{k-1}(e) < \mathbf{x}_k(e)$, then $\mathbf{o}_k(e) = \mathbf{o}_{k-1}(e) + 1$, and the first term of (4) is

$$\begin{aligned} & f(\mathbf{o}_{k-1}) - f(\mathbf{o}_{k-1} \vee \mathbf{x}_k(e)\chi_e) \\ & = f(\mathbf{o}_{k-1}) - f(\mathbf{o}_{k-1} + \chi_e) \\ & \leq f(\mathbf{y}_{k-1} - \chi_e) - f(\mathbf{y}_{k-1}) \\ & = f(-\chi_e | \mathbf{y}_{k-1}) \\ & \leq (1+\epsilon)\beta. \end{aligned}$$

Here, the first inequality uses the DR-submodularity of f and the fact that $\mathbf{o}_{k-1} \leq \mathbf{y}_{k-1} - \chi_e$, and the second inequality uses the guarantee in Lemma 2. The second term of (4) is zero, and hence we have (4) $\leq \frac{(1+\epsilon)\alpha\beta}{\alpha+\beta}$.

If $\mathbf{y}_k(e) < \mathbf{o}_{k-1}(e)$, then by a similar argument, we have (4) $\leq \frac{(1+\epsilon)\alpha\beta}{\alpha+\beta}$. \square

We now return to proving Lemma 3. By Claim 4,

$$(4) \leq \frac{(1+\epsilon)\alpha\beta}{\alpha+\beta} \leq \frac{1+\epsilon}{2} \left(\frac{\alpha^2 + \beta^2}{\alpha + \beta} \right) \leq \frac{1+\epsilon}{2}. \quad (3),$$

which indicates the desired result. \square

Theorem 5. *Algorithm 3 is a $\frac{1}{2+\epsilon}$ -approximation algorithm for (1) with time complexity $O(\frac{|E|}{\epsilon} \cdot \log(\frac{\Delta}{\delta}) \log \|\mathbf{B}\|_\infty \cdot \theta + \|\mathbf{B}\|_1 \log \|\mathbf{B}\|_\infty)$, where δ and Δ are the minimum positive marginal gain and the maximum positive value, respectively, of f and θ is the running time of evaluating f .*

Proof. Summing up (2) for $k \in [|\mathbf{B}|_1]$, we get

$$\begin{aligned} & \sum_{k=1}^{|\mathbf{B}|_1} \mathbf{E}[f(\mathbf{o}_{k-1}) - f(\mathbf{o}_k)] \\ & \leq \frac{1+\epsilon}{2} \sum_{k=1}^{|\mathbf{B}|_1} \mathbf{E}[f(\mathbf{x}_k) - f(\mathbf{x}_{k-1}) + f(\mathbf{y}_k) - f(\mathbf{y}_{k-1})]. \end{aligned}$$

The above sum is telescopic, and hence we obtain

$$\begin{aligned} & \mathbf{E}[f(\mathbf{o}_0) - f(\mathbf{o}_{|\mathbf{B}|_1})] \\ & \leq \frac{1+\epsilon}{2} \mathbf{E}[f(\mathbf{x}_{|\mathbf{B}|_1}) - f(\mathbf{x}_0) + f(\mathbf{y}_{|\mathbf{B}|_1}) - f(\mathbf{y}_0)] \\ & \leq \frac{1+\epsilon}{2} \mathbf{E}[f(\mathbf{x}_{|\mathbf{B}|_1}) + f(\mathbf{y}_{|\mathbf{B}|_1})] \\ & = (1+\epsilon) \mathbf{E}[f(\mathbf{x}_{|\mathbf{B}|_1})]. \end{aligned}$$

The second inequality uses the fact that f is non-negative, and the last equality uses $\mathbf{y}_{|\mathbf{B}|_1} = \mathbf{x}_{|\mathbf{B}|_1}$. Because $\mathbf{E}[f(\mathbf{o}_0) - f(\mathbf{o}_{|\mathbf{B}|_1})] = f(\mathbf{o}) - \mathbf{E}[f(\mathbf{x}_{|\mathbf{B}|_1})]$, we obtain that $\mathbf{E}[f(\mathbf{x}_{|\mathbf{B}|_1})] \geq \frac{1}{2+\epsilon} f(\mathbf{o})$.

We now analyze the time complexity. We only query the input function f inside of Algorithm 2, and the number of oracle calls is $O(\frac{|E|}{\epsilon} \log(\frac{\Delta}{\delta}) \log B)$ by Lemma 2. Note that we invoke Algorithm 2 with g and h , and the minimum positive values of g and h are at least the minimum positive marginal gain δ of f . The number of iterations is $\|\mathbf{B}\|_1$, and we need $O(\log B)$ time to access \tilde{g} and \tilde{h} . Hence, the total time complexity is as stated. \square

Remark 6. *We note that even if f is not a non-negative function, the proof of Theorem 5 works as long as $f(\mathbf{x}_0) \geq 0$ and $f(\mathbf{y}_0) \geq 0$, that is, $f(\mathbf{0}) \geq 0$ and $f(\mathbf{B}) \geq 0$. Hence, given a DR-submodular function $f : \mathbb{Z}_+^E \rightarrow \mathbb{R}$ and $\mathbf{B} \in \mathbb{Z}_+^E$, we can obtain a $\frac{1}{2+\epsilon}$ -approximation algorithm for the following problem:*

$$\begin{aligned} & \text{maximize} && f(\mathbf{x}) - \min\{f(\mathbf{0}), f(\mathbf{B})\} \\ & \text{subject to} && \mathbf{0} \leq \mathbf{x} \leq \mathbf{B}, \end{aligned} \quad (5)$$

This observation is useful, as the objective function often takes negative values in real-world applications.

Polynomial-time algorithm

In many applications, the running time needed to evaluate the input function is a bottleneck, and hence Algorithm 3 is already satisfactory. However, it is theoretically interesting to reduce the total running time to a polynomial, and we show the following. The proof is deferred to the full version.

Theorem 7. *There exists a $\frac{1}{2+2\epsilon}$ -approximation algorithm with time complexity $\tilde{O}(\frac{|E|}{\epsilon} \log(\frac{\Delta}{\delta}) \log \|\mathbf{B}\|_\infty \cdot (\theta + \log \|\mathbf{B}\|_\infty))$, where δ and Δ are the minimum positive marginal gain and the maximum positive value, respectively of f and θ is the running time of evaluating f . Here $\tilde{O}(T)$ means $O(T \log^c T)$ for some $c \in \mathbb{N}$.*

Experiments

In this section, we show our experimental results and the superiority of our algorithm with respect to other baseline algorithms.

Experimental setting

We conducted experiments on a Linux server with an Intel Xeon E5-2690 (2.90 GHz) processor and 256 GB of main memory. All the algorithms were implemented in C# and were run using Mono 4.2.3.

We compared the following four algorithms:

- **Single Greedy (SG):** We start with $\mathbf{x} = \mathbf{0}$. For each element $e \in E$, as long as the marginal gain of adding χ_e to the current solution \mathbf{x} is positive, we add it to \mathbf{x} . The reason that we do not choose the element with the maximum marginal gain is to reduce the number of oracle calls, and our preliminary experiments showed that such a tweak does not improve the solution quality.
- **Double Greedy (DG, Algorithm 1).**
- **Lattice Double Greedy (Lattice-DG):** The $1/3$ -approximation algorithm for maximizing non-monotone lattice submodular functions (Gottschalk and Peis 2015).
- **Double Greedy with a polynomial number of oracle calls with error parameter $\epsilon > 0$ (Fast-DG $_\epsilon$, Algorithm 3).**

We measure the efficiency of an algorithm by the number of oracle calls instead of the total running time. Indeed, the running time for evaluating the input function is the dominant factor of the total running, because objective functions in typical machine learning tasks contain sums over all data points, which is time consuming. Therefore, we do not consider the polynomial-time algorithm (Theorem 7) here.

Revenue maximization

In this application, we consider revenue maximization on an (undirected) social network $G = (V, W)$, where $W = (w_{ij})_{i,j \in V}$ represents the weights of edges. The goal is to offer for free or advertise a product to users so that the revenue increases through their word-of-mouth effect on others. If we invest x units of cost on a user $i \in V$, the user becomes an advocate of the product (independently from other users) with probability $1 - (1-p)^x$, where $p \in (0, 1)$ is a parameter. This means that, for investing a unit cost to i , we have an extra chance that the user i becomes an advocate

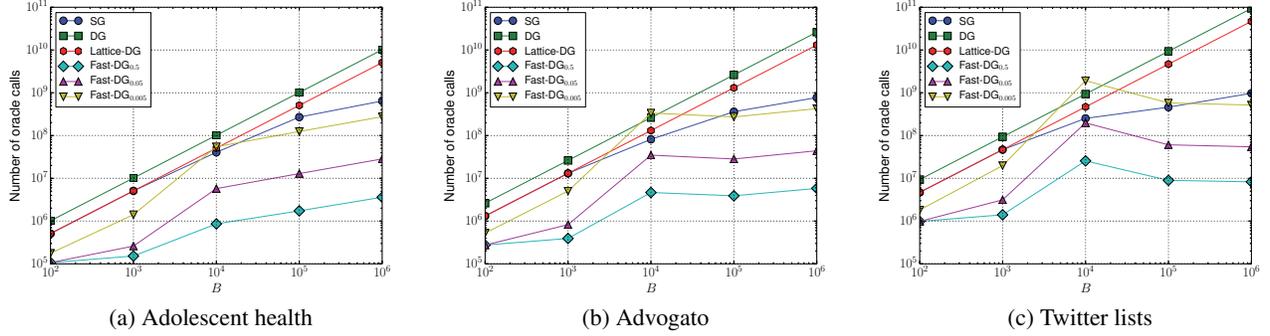


Figure 1: Number of oracle calls

Table 1: Objective values (Our methods are highlighted in gray.)

	Adolescent health				
	$B = 10^2$	10^3	10^4	10^5	10^6
SG	280.55	2452.16	7093.73	7331.42	7331.50
DG	280.55	2452.16	7124.90	7332.96	7331.50
Lattice-DG	215.39	1699.66	6808.97	6709.11	5734.30
Fast-DG _{0.5}	280.55	2452.16	7101.14	7331.36	7331.48
Fast-DG _{0.05}	280.55	2452.16	7100.86	7331.36	7331.48
Fast-DG _{0.005}	280.55	2452.16	7100.83	7331.36	7331.48
	Advogato				
	$B = 10^2$	10^3	10^4	10^5	10^6
SG	993.15	8680.87	25516.05	27325.78	27326.01
DG	993.15	8680.87	25330.91	27329.39	27326.01
Lattice-DG	753.93	6123.39	24289.09	24878.94	21674.35
Fast-DG _{0.5}	993.15	8680.87	25520.83	27325.75	27325.98
Fast-DG _{0.05}	993.15	8680.87	25520.52	27325.75	27325.98
Fast-DG _{0.005}	993.15	8680.87	25520.47	27325.75	27325.98
	Twitter lists				
	$B = 10^2$	10^3	10^4	10^5	10^6
SG	882.43	7713.07	22452.61	25743.26	25744.02
DG	882.43	7713.07	22455.97	25751.42	25744.02
Lattice-DG	675.67	5263.87	20918.89	20847.48	15001.19
Fast-DG _{0.5}	882.43	7713.07	22664.65	25743.06	25743.88
Fast-DG _{0.05}	882.43	7713.07	22658.58	25743.06	25743.88
Fast-DG _{0.005}	882.43	7713.07	22658.07	25743.06	25743.88

with probability p . Let $S \subseteq V$ be a set of users who advocate the product. Note that S is a random set. Following a simplified version of the model introduced by (Hartline, Mirrokni, and Sundararajan 2008), the revenue is defined as $\sum_{i \in S} \sum_{j \in V \setminus S} w_{ij}$. Let $f : \mathbb{Z}_+^E \rightarrow \mathbb{R}$ be the expected revenue obtained in this model, that is,

$$\begin{aligned}
 f(\mathbf{x}) &= \mathbf{E} \left[\sum_{i \in S} \sum_{j \in V \setminus S} w_{ij} \right] \\
 &= \sum_{i \in S} \sum_{j \in V \setminus S} w_{ij} (1 - (1 - p)^{x(i)}) (1 - p)^{x(j)}.
 \end{aligned}$$

It is not hard to show that f is non-monotone DR-submodular function (see the full version for the proof).

In our experiment, we used three networks, Adolescent health (2,539 vertices and 12,969 edges), Advogato (6,541 vertices and 61,127 edges), and Twitter lists (23,370 vertices and 33,101 edges), all taken from (Kunegis 2013). We regard all the networks as undirected. We set $p = 0.0001$, and set $w_{ij} = 1$ when an edge exists between i and j and $w_{ij} = 0$

otherwise. We imposed the constraint $0 \leq x(e) \leq B$ for every $e \in E$, where B is chosen from $\{10^2, \dots, 10^6\}$.

Table 1 shows the objective values obtained by each method. As can be seen, except for Lattice-DG, which is clearly the worst, the choice of a method does not much affect the obtained objective value for all the networks. Notably, even when ϵ is as large as 0.5, the objective values obtained by Fast-DG are almost the same as SG and DG.

Figure 1 illustrates the number of oracle calls of each method. The number of oracle calls of DG and Lattice-DG is linear in B , whereas that of Fast-DG slowly grows. Although the number of oracle calls of SG also slowly grows, it is always orders of magnitude larger than that of Fast-DG with $\epsilon = 0.5$ or $\epsilon = 0.05$.

In summary, Fast-DG_{0.5} achieves almost the best objective value, whereas the number of oracle calls is two or three orders of magnitude smaller than those of the other methods when B is large.

Conclusions

In this paper, we proposed a polynomial-time $\frac{1}{2+\epsilon}$ -approximation algorithm for non-monotone DR-submodular function maximization. Our experimental results on the revenue maximization problem showed the superiority of our method against other baseline algorithms.

Maximizing a submodular set function under constraints is well studied (Lee et al. 2009; Gupta et al. 2010; Chekuri, Vondrák, and Zenklusen 2014; Mirzasoleiman et al. 2016). An intriguing open question is whether we can obtain polynomial-time algorithms for maximizing DR-submodular functions under constraints such as cardinality constraints, polymatroid constraints, and knapsack constraints.

Acknowledgments T. S. is supported by JSPS Grant-in-Aid for Research Activity Start-up. Y. Y. is supported by JSPS Grant-in-Aid for Young Scientists (B) (No. 26730009), MEXT Grant-in-Aid for Scientific Research on Innovative Areas (No. 24106003), and JST, ERATO, Kawarabayashi Large Graph Project.

References

- Alon, N.; Gamzu, I.; and Tennenholtz, M. 2012. Optimizing budget allocation among channels and influencers. In *WWW*, 381–388.
- Badanidiyuru, A., and Vondrák, J. 2014. Fast algorithms for maximizing submodular functions. In *SODA*, 1497–1514.
- Bian, Y.; Mirzasoleiman, B.; Buhmann, J. M.; and Krause, A. 2016. Guaranteed non-convex optimization: Submodular maximization over continuous domains. *CoRR* abs/1606.05615.
- Buchbinder, N., and Feldman, M. 2016. Deterministic algorithms for submodular maximization problems. In *SODA*, 392–403.
- Buchbinder, N.; Feldman, M.; Naor, J. S.; and Schwartz, R. 2012. A tight linear time $(1/2)$ -approximation for unconstrained submodular maximization. In *FOCS*, 649–658.
- Buchbinder, N.; Feldman, M.; Naor, J. S.; and Schwartz, R. 2014. Submodular maximization with cardinality constraints. In *SODA*, 1433–1452.
- Chekuri, C.; Vondrák, J.; and Zenklusen, R. 2014. Submodular function maximization via the multilinear relaxation and contention resolution schemes. *SIAM Journal on Computing* 43(6):1831–1879.
- Ene, A., and Nguyen, H. L. 2016. A reduction for optimizing lattice submodular functions with diminishing returns. *CoRR* abs/1606.08362.
- Feige, U.; Mirrokni, V. S.; and Vondrak, J. 2011. Maximizing non-monotone submodular functions. *SIAM J. on Comput.* 40(4):1133–1153.
- Fujishige, S. 2005. *Submodular Functions and Optimization*. Elsevier, 2nd edition.
- Golovin, D., and Krause, A. 2011. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *Journal of Artificial Intelligence Research* 427–486.
- Gotovos, A.; Karbasi, A.; and Krause, A. 2015. Non-monotone adaptive submodular maximization. In *AAAI*, 1996–2003.
- Gottschalk, C., and Peis, B. 2015. Submodular function maximization on the bounded integer lattice. In *WAOA*, 133–144.
- Gupta, A.; Roth, A.; Schoenebeck, G.; and Talwar, K. 2010. Constrained non-monotone submodular maximization: offline and secretary algorithms. In *WINE*, 246–257.
- Hartline, J.; Mirrokni, V.; and Sundararajan, M. 2008. Optimal marketing strategies over social networks. In *WWW*, 189–198.
- Ko, C. W.; Lee, J.; and Queyranne, M. 1995. An exact algorithm for maximum entropy sampling. *Operations Research* 684–691.
- Krause, A., and Golovin, D. 2014. Submodular function maximization. In *Tractability: Practical Approaches to Hard Problems*. Cambridge University Press. 71–104.
- Krause, A., and Horvitz, E. 2008. A utility-theoretic approach to privacy and personalization. In *AAAI*, 1181–1188.
- Krause, A.; Singh, A.; and Guestrin, C. 2008. Near-optimal sensor placements in gaussian processes: theory, efficient algorithms and empirical studies. *The Journal of Machine Learning Research* 9:235–284.
- Kunegis, J. 2013. Konect - the koblenz network collection. In *WWW Companion*, 1343–1350.
- Lee, J.; Mirrokni, V. S.; Nagarajan, V.; and Sviridenko, M. 2009. Non-monotone submodular maximization under matroid and knapsack constraints. In *STOC*, 323–332.
- Milgrom, P., and Strulovici, B. 2009. Substitute goods, auctions, and equilibrium. *Journal of Economic Theory* 144(1):212–247.
- Mirzasoleiman, B.; CH, E.; Karbasi, A.; and EDU, Y. 2016. Fast constrained submodular maximization: Personalized data summarization. In *ICLM '16: Proceedings of the 33rd International Conference on Machine Learning (ICML)*.
- Soma, T., and Yoshida, Y. 2015. A generalization of submodular cover via the diminishing return property on the integer lattice. In *NIPS*, 847–855.
- Soma, T., and Yoshida, Y. 2016. Maximizing submodular functions with the diminishing return property over the integer lattice. In *IPCO*, 325–336.
- Soma, T.; Kakimura, N.; Inaba, K.; and Kawarabayashi, K. 2014. Optimal budget allocation: Theoretical guarantee and efficient algorithm. In *ICML*, 351–359.