
On the Reducibility of Submodular Functions

Jincheng Mei

Department of Computing Science
University of Alberta
Edmonton, AB, Canada, T6G 2E8

Hao Zhang

The Robotics Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213, USA

Bao-Liang Lu*

Department of Computer Sci. and Eng.
Key Lab of SEC for IICE
Shanghai Jiao Tong University
Shanghai 200240, China

Abstract

The scalability of submodular optimization methods is critical for their usability in practice. In this paper, we study the *reducibility* of submodular functions, a property that enables us to reduce the solution space of submodular optimization problems without performance loss. We introduce the concept of reducibility using *marginal gains*. Then we show that by adding perturbation, we can endow irreducible functions with reducibility, based on which we propose the *perturbation-reduction* optimization framework. Our theoretical analysis proves that given the perturbation scales, the reducibility gain could be computed, and the performance loss has additive upper bounds. We further conduct empirical studies and the results demonstrate that our proposed framework significantly accelerates existing optimization methods for irreducible submodular functions with a cost of only small performance losses.

1 INTRODUCTION

Submodularity naturally arises in a number of machine learning problems, such as active learning [10], clustering [21], and dictionary selection [5]. The scalability of submodular optimization methods is critical in practice, thus has drawn much attention from the research community. For example, Iyer et al. [12] propose $\mathcal{O}(n^2)$ general optimization methods based on the semidifferential. Wei et al. [24] combine approximation with pruning to accelerate the greedy algorithm for uniform matroid constrained submodular

maximization. Mirzasoleiman et al. [20] and Pan et al. [23] use distributed implementation to accelerate existing optimization methods. Other techniques, including stochastic sampling [19] and decomposable assumption [14], are also applied to scale up submodular optimization methods. While in this paper, we focus on the reducibility of submodular functions, a favourable property that can substantially improve the scalability of submodular optimization methods. The reducibility can directly reduce the solution space of the submodular optimization problems, while preserve all the optima in the reduced space, thereby enables us to accelerate the optimization process without incurring performance loss.

Recent research shows that for some submodular functions, by evaluating marginal gains, reduction can be applied for unconstrained maximization [8], unconstrained minimization [6, 12], and uniform matroid constrained maximization [24]. By leveraging the reducibility, a variety of methods have been developed to scale up the optimization of *reducible* submodular functions [6, 8, 12, 18]. While existing works mainly focus on reducible functions, there exist a number of *irreducible* submodular functions widely applied in practice, for which existing methods can only provide vacuous reduction.

In this paper, we investigate the problem that whether irreducible functions can also exploit this favorable property. We firstly introduce the concept of reducibility using marginal gains over the endpoint sets of a given lattice. Then for irreducible functions, we transform them to reducible functions by adding random noise to perturb the marginal gains, after which we perform lattice reduction for the perturbed functions and solve the original functions on the reduced lattice. Theoretical results show that given the perturbation scales, the reducibility gain is lower bounded, and the performance loss has additive upper bounds. The empirical results demonstrate that there exist useful perturbation scale intervals in practice, which enables us

Appearing in Proceedings of the 19th International Conference on Artificial Intelligence and Statistics (AISTATS) 2016, Cadiz, Spain. JMLR: W&CP volume 51. Copyright 2016 by the authors.

to significantly accelerate existing optimization methods with small performance losses.

In summary, this paper has the following contributions. Firstly, we introduce the concept of reducibility, and propose the perturbation-reduction framework. Secondly, we theoretically analyze our proposed method. In particular, for the reducibility gain, we propose a lower bound in terms of the perturbation scale. For the performance loss, we propose both deterministic and probabilistic upper bounds. The deterministic bound provides the understanding of relationship between the reducibility gain and performance loss, while the probabilistic bound can explain the experimental results. Finally, we empirically show that the proposed method is applicable for a variety of commonly used irreducible submodular functions.

In the sequel, we organize the paper as follows. In Section 2, we introduce the definitions and the existing reduction algorithms. In Section 3, we propose our perturbation based method. Theoretical analysis and empirical results are presented in Section 4 and Section 5, respectively. In Section 6, we review some related works. Section 7 comes to our conclusion.

2 REDUCIBILITY

2.1 Notations and Definitions

Given a finite set $N = \{1, 2, \dots, n\}$, and a set function $f : 2^N \mapsto \mathbb{R}$. f is said to be *submodular* [6] if $\forall X, Y \subseteq N$, $f(X) + f(Y) \geq f(X \cap Y) + f(X \cup Y)$. An equivalent definition of submodularity is the *diminishing return* property: $\forall A \subseteq B \subseteq N$, $\forall i \in N \setminus B$, $f(i|A) \geq f(i|B)$, where $f(i|A) \triangleq f(A + i) - f(A)$ is called the *marginal gain* of element i with respect to set A . To simplify the notation, we denote $A \cup \{i\}$ by $A + i$, and $A \setminus \{i\}$ by $A - i$. Given $A \subseteq B \subseteq N$, the (*set interval*) *lattice* is defined as $[A, B] \triangleq \{S \mid A \subseteq S \subseteq B\}$.

Suppose $N = \{1, 2, \dots, n\}$, and $f : 2^N \mapsto \mathbb{R}$ is a submodular function. In this paper, we focus on unconstrained submodular optimization problems,

$$\text{Problem 1: } \min_{X \subseteq N} f(X), \quad \text{Problem 2: } \max_{X \subseteq N} f(X).$$

Problem 1 can be exactly solved in polynomial time [22], while Problem 2 is NP-hard since some of its special cases (e.g., Max Cut) are NP-hard.

For convenience of the presentation, we will use P1 and P2 to refer to Problem 1 and Problem 2, respectively. Similarly, for the following algorithms, we will use A1 and A2 to refer to Algorithm 1 and Algorithm 2. The reference holds for P3, P4, A3 and A4.

Define $\mathcal{X}_{min} \triangleq \{X_* \subseteq N \mid f(X_*) \leq f(X), \forall X \subseteq N\}$ as the optima set of P1. Similarly, define $\mathcal{X}_{max} \triangleq \{X^* \subseteq$

$N \mid f(X^*) \geq f(X), \forall X \subseteq N\}$. Obviously, we have $\mathcal{X}_{min} \subseteq [\emptyset, N]$ and $\mathcal{X}_{max} \subseteq [\emptyset, N]$.

We now give the definition of reducibility. For P1 (P2), we say the objective function f is *reducible for minimization (maximization)* if $\exists [S, T] \subset [\emptyset, N]$, where $[S, T]$ can be obtained in $\mathcal{O}(n^p)$ function evaluations, such that $\mathcal{X}_{min} \subseteq [S, T]$ ($\mathcal{X}_{max} \subseteq [S, T]$). Note that if we can only find $[S, T]$ in $\mathcal{O}(2^n)$ time, the reduction is meaningless since $\mathcal{O}(2^n)$ time is enough for us to find all the optima. The ratio $1 - \frac{|T \setminus S|}{|N|} \in (0, 1]$ is called the *reduction rate*.

2.2 Algorithms

Existing works on reduction for unconstrained submodular optimization can be summarized by the following two algorithms, both of which terminate in $\mathcal{O}(n^2)$ time. The brief review of existing works can be found in Section 6.

Algorithm 1 Reduction for Minimization

Input: $N, f, X_0 \leftarrow \emptyset, Y_0 \leftarrow N, t \leftarrow 0$.

Output: $[X_t, Y_t]$.

- 1: Find $U_t = \{i \in Y_t \setminus X_t \mid f(i|X_t) < 0\}$.
 - 2: $X_{t+1} \leftarrow X_t \cup U_t$.
 - 3: Find $D_t = \{j \in Y_t \setminus X_t \mid f(j|Y_t - j) > 0\}$.
 - 4: $Y_{t+1} \leftarrow Y_t \setminus D_t$.
 - 5: If $X_{t+1} = X_t$ and $Y_{t+1} = Y_t$, terminate.
 - 6: $t \leftarrow t + 1$. Go to Step 1.
-

Algorithm 2 Reduction for Maximization

Input: $N, f, X_0 \leftarrow \emptyset, Y_0 \leftarrow N, t \leftarrow 0$.

Output: $[X_t, Y_t]$.

- 1: Find $U_t = \{i \in Y_t \setminus X_t \mid f(i|X_t) < 0\}$.
 - 2: $Y_{t+1} \leftarrow Y_t \setminus U_t$.
 - 3: Find $D_t = \{j \in Y_t \setminus X_t \mid f(j|Y_t - j) > 0\}$.
 - 4: $X_{t+1} \leftarrow X_t \cup D_t$.
 - 5: If $X_{t+1} = X_t$ and $Y_{t+1} = Y_t$, terminate.
 - 6: $t \leftarrow t + 1$. Go to Step 1.
-

Proposition 1. *Suppose $f : 2^N \mapsto \mathbb{R}$ is submodular. After each iteration of A1 (A2), we have $\mathcal{X}_{min} \subseteq [X_t, Y_t]$ ($\mathcal{X}_{max} \subseteq [X_t, Y_t]$).*

We prove Proposition 1 in the supplementary material. According to Proposition 1, if the output of A1 (A2) satisfies $[X_t, Y_t] \subset [\emptyset, N]$, then f is reducible.

According to A1 (A2), if $U_0 = D_0 = \emptyset$, then we have $X_1 = X_0$ and $Y_1 = Y_0$. The algorithm will terminate after the first iteration and the output is $[X_0, Y_0] = [\emptyset, N]$, which provides a vacuous reduction. In this case, we say that f is *irreducible with respect to A1 (A2)*. For convenience, we directly say f is *irreducible*. Thereby, we conclude two points from the above al-

gorithms. First, by the definition of U_t and D_t , the reducibility of f can be determined by the signs of marginal gains with respect to the endpoint sets of the current working lattice. Second, the reducibility of f for minimization and maximization are actually the same property. Specially, suppose in a certain iteration, A1 and A2 have the same working lattice $[S, T]$. According to the algorithms, they also have the same U_t and D_t , which determine whether f is reducible after the current iteration.

Proposition 2. *Given a submodular function $f : 2^N \mapsto \mathbb{R}$, and a lattice $[S, T]$. $\forall i \in T \setminus S$, Define $K_i = \text{sgn}\{f(i|S)\} \cdot \text{sgn}\{f(i|T-i)\}$. Then f is reducible on $[S, T]$ with respect to A1 (A2) if and only if*

$$K = \max_{i \in T \setminus S} K_i > 0. \quad (1)$$

Proof. Suppose $[X_0, Y_0] = [S, T]$ in A1 (A2). Then f is reducible if and only if the algorithm does not terminate after its first iteration, *i.e.*, $U_0 \neq \emptyset$ or $D_0 \neq \emptyset$. Suppose $U_0 \neq \emptyset$ happens, *i.e.*, $\exists i \in T \setminus S$, $f(i|S) < 0$. According to submodularity, $f(i|T-i) \leq f(i|S) < 0$. We have $K \geq K_i = \text{sgn}\{f(i|S)\} \cdot \text{sgn}\{f(i|T-i)\} = 1 > 0$. Suppose $D_0 \neq \emptyset$ happens, *i.e.*, $\exists j \in T \setminus S$, $f(j|T-j) > 0$. According to submodularity, $f(j|S) \geq f(j|T-j) > 0$. We have $K \geq K_j = \text{sgn}\{f(j|S)\} \cdot \text{sgn}\{f(j|T-j)\} = 1 > 0$. \square

According to Proposition 2, the reducibility of f for minimization (maximization) can be obtained by (1). Thus we say f is *reducible with respect to A1 (A2)* if (1) holds. Similarly, without ambiguity in this paper, we directly say f is *reducible* if (1) holds.

3 PERTURBATION REDUCTION

Given a reducible submodular function, we can use A1 and A2 to provide useful reduction. Unfortunately, there still exist many irreducible submodular functions, some of which are listed in the experimental section. Given a submodular function f , which is irreducible on $[S, T]$. According to Proposition 2, $\forall i \in T \setminus S$, we have $f(i|S) \geq 0$ and $f(i|T-i) \leq 0$.

If we expect A1 and A2 to provide nontrivial reduction, we need to guarantee that (1) holds for some elements without changing the submodularity of the objective function. A natural way is to add random noise r^1 to perturb the original function as follows,

$$\text{Problem 3 : } \min_{X \subseteq N} g(X) \triangleq \min_{X \subseteq N} f(X) + r(X),$$

$$\text{Problem 4 : } \max_{X \subseteq N} g(X) \triangleq \max_{X \subseteq N} f(X) + r(X),$$

¹ $r : 2^N \mapsto \mathbb{R}$ is a modular function, and $r(X) \triangleq \sum_{i \in X} r(i)$.

where $\forall i \in N$, $r(i) \in \mathbb{R}$ is generated uniformly at random in $[-t, t]$ for some $t \geq 0$. By appropriately choosing the value of t , we can ensure $g(i|S) < 0$ or $g(i|T-i) > 0$ hold for some $i \in T \setminus S$. Thus we have (1) holds, indicating that g is reducible. At the same time, as r is a modular function, the submodularity of g still holds.

Algorithm 3 Perturbation-Reduction Minimization

Input: $N, f, [S, T]$ where $\mathcal{X}_{min} \subseteq [S, T]$.

Output: An approximate solution X_*^p .

- 1: If f is reducible on $[S, T]$, $[X_0, Y_0] \leftarrow [S, T]$, run A1 for f , $[S, T] \leftarrow [X_t, Y_t]$.
 - 2: Generate r . Let $g = f + r$. $[X_0, Y_0] \leftarrow [S, T]$, run A1 for g , $[S, T] \leftarrow [X_t, Y_t]$.
 - 3: Solve $X_*^p \in \arg \min_{X \in [S, T]} f(X)$.
-

Algorithm 4 Perturbation-Reduction Maximization

Input: $N, f, [S, T]$ where $\mathcal{X}_{max} \subseteq [S, T]$.

Output: An approximate solution X_p^* .

- 1: If f is reducible on $[S, T]$, $[X_0, Y_0] \leftarrow [S, T]$, run A2 for f , $[S, T] \leftarrow [X_t, Y_t]$.
 - 2: Generate r . Let $g = f + r$. $[X_0, Y_0] \leftarrow [S, T]$, run A2 for g , $[S, T] \leftarrow [X_t, Y_t]$.
 - 3: Solve $X_p^* \in \arg \max_{X \in [S, T]} f(X)$.
-

We propose our perturbation based method for minimization and maximization in A3 and A4, respectively. For an irreducible submodular function f on a given lattice $[S, T]$, we first perturb the objective function to make it reducible, *i.e.*, $g \triangleq f + r$. A1 or A2 are then employed to obtain the reduced lattice of g . Finally we solve the original problems of f on the reduced lattice exactly or approximately using existing methods.

It is worth mentioning that, though we mainly focus on irreducible functions, our methods also work for reducible ones, as they are special cases of irreducible functions. Particularly, given a reducible function f on $[S, T]$, of which the reduction rate is less than 1, after A1 (A2) terminates, we can get a sublattice $[P, Q] \subseteq [S, T]$ so that f is irreducible on $[P, Q]$.

4 THEORETICAL ANALYSIS

By perturbing the irreducible submodular function, we transform P1 (P2) into P3 (P4). This makes the objective function reducible while leads the solution to be inexact. Correspondingly, our theoretical analysis of the method focuses on two main aspects: the reducibility gain and the performance loss incurred by perturbation.

4.1 Reducibility Gain

Suppose $f : 2^N \mapsto \mathbb{R}$ is an irreducible submodular function on $[S, T]$, and $g \triangleq f + r$ as defined in P3 and P4. Since f is irreducible, $\forall i \in T \setminus S$, we have $f(i|S) \geq 0$ and $f(i|T-i) \leq 0$.

Proposition 3. *Given a submodular function $f : 2^N \mapsto \mathbb{R}$, which is irreducible on $[S, T]$. Define $m\{f, [S, T]\} \triangleq \min_{i \in T \setminus S} \min\{f(i|S), -f(i|T-i)\}$. If $t \leq m\{f, [S, T]\}$, then g is irreducible on $[S, T]$.*

Proof. Since $m\{f, [S, T]\} \geq 0$, we suppose $0 \leq t \leq m\{f, [S, T]\}$. $\forall i \in T \setminus S$, we have $g(i|S) = f(i|S) + r(i) \geq f(i|S) - t \geq f(i|S) - m\{f, [S, T]\} \geq 0$, and $g(i|T-i) = f(i|T-i) + r(i) \leq f(i|T-i) + t \leq f(i|T-i) + m\{f, [S, T]\} \leq 0$, which implies that g is also irreducible on $[S, T]$. \square

Proposition 3 indicates that if the perturbation scale t is small enough, there is no reducibility gain. This is intuitively reasonable since we have $g \rightarrow f$ when $t \rightarrow 0$.

To lower bound the reducibility gain of adding perturbation, we generalize the concept of *curvature* [4, 13] for non-monotone irreducible submodular functions.

Definition 1. *Given a submodular function $f : 2^N \mapsto \mathbb{R}$, the curvature of f on $[S, T]$ is defined as,*

$$c\{f, [S, T]\} = \max_{i \in T \setminus S, f(i|S) > 0} \frac{f(i|S) - f(i|T-i)}{f(i|S)}.$$

Note that for any irreducible submodular function f on $[S, T]$, we have $c\{f, [S, T]\} \geq 1$.

Theorem 1. *Suppose $t > m\{f, [S, T]\}$, denote $s = |T \setminus S| > 0$, $k = \sum_{i \in T \setminus S} f(i|S)$, $c = c\{f, [S, T]\}$. The reduction rate in expectation of g is at least $1 - \frac{ck}{2ts}$.*

Proof. Suppose $T \setminus S = \{1, 2, \dots, s\}$, $\forall i \in T \setminus S$, we define a random variable H_i as,

$$H_i = \begin{cases} 1 & \text{if } K_i > 0, \\ 0 & \text{otherwise.} \end{cases}$$

H_i indicates whether i can be reduced from $T \setminus S$ or not. Define $H = \sum_{i \in T \setminus S} H_i$ as the total number of the reduced elements. We firstly lower bound $\mathbb{E}(H)$ by the total number of the reduced elements after the first iteration round of A1(A2),

$$\begin{aligned} \mathbb{E}(H) &= \sum_{i \in T \setminus S} \mathbb{E}(H_i) = \sum_{i \in T \setminus S} \Pr\{H_i = 1\} \\ &\geq \frac{1}{2t} \cdot \sum_{i \in T \setminus S} \max\{0, t - f(i|S)\} + \max\{0, t + f(i|T-i)\} \\ &\geq s - \frac{c}{2t} \cdot \sum_{i \in T \setminus S} f(i|S) = s - \frac{ck}{2t}. \end{aligned}$$

Consequently, the reduction rate in expectation is $\frac{\mathbb{E}(H)}{s} \geq 1 - \frac{ck}{2ts}$. \square

Theorem 1 implies that the reduction rate in expectation approaches 1 as the perturbation scale t increases. This is also consistent with our intuition since $g \rightarrow r$ when $t \rightarrow \infty$. Note that r is a modular function, which always has the highest reduction rate 1.

4.2 Performance Loss

Suppose $X_* \in \mathcal{X}_{min}$ ($X_* \in \mathcal{X}_{max}$), i.e., X_* (X^*) is an optimum of P1 (P2). Recall that X_*^p (X^*^p) is the output of A3 (A4), for P1 (P2), we define $f(X_*^p) - f(X_*)$ ($f(X^*^p) - f(X^*)$) as the *performance loss* incurred by perturbation. For P1 (P2), the following result shows that the performance loss is upper bounded by the total perturbation of the ‘‘mistakenly’’ reduced elements, which will be explained later on.

Theorem 2. *Given an irreducible submodular function $f : 2^N \mapsto \mathbb{R}$. Suppose t is the perturbation scale in A3 (A4), and R_t is the reduction rate. We have,*

$$\begin{aligned} f(X_*^p) - f(X_*) &< -r(X_t \setminus X_*) + r(X_* \setminus Y_t) < ntR_t, \\ f(X^*^p) - f(X^*) &< r(X_t \setminus X^*) - r(X^* \setminus Y_t) < ntR_t. \end{aligned}$$

Proof. We prove the maximization case. In general, we have $X^* \notin [X_t, Y_t]$, otherwise the loss is zero. Note that $X_*^p \in [X_t, Y_t]$ according to A4. So we firstly introduce an intermediate set $X^* \cup X_t \cap Y_t$, i.e., the contraction of X^* in $[X_t, Y_t]$ for our analysis. Given the fact that $f(X^* \cup X_t \cap Y_t) \leq f(X_*^p)$, if we can upper bound $f(X^*) - f(X^* \cup X_t \cap Y_t)$, then the total performance loss is also upper bounded. In A2, we have $[X_t, Y_t] \subset \dots \subset [X_1, Y_1] \subset [X_0, Y_0] = [\emptyset, N]$. By definition, $\forall 0 \leq k \leq t-1$, $X_{k+1} = X_k \cup D_k$, and $Y_{k+1} = Y_k \setminus U_k$. We have,

$$f(X^* \cup X_t) - f(X^*) \tag{2}$$

$$= \sum_{s=1, x_s \in X_t \setminus X^*}^{|X_t \setminus X^*|} f(x_s | X^* + x_1 + \dots + x_{s-1}) \tag{3}$$

$$\geq \sum_{i=0}^{t-1} \sum_{d \in D_i \setminus X^*} f(d | Y_i - d) \tag{4}$$

$$> - \sum_{i=0}^{t-1} \sum_{d \in D_i \setminus X^*} r(d) \tag{5}$$

$$= -r(X_t \setminus X^*), \tag{6}$$

where (3) is the telescopic version of (2). According to submodularity, we have (4) holds, and (5) comes from the third step of A2. Similarly, we have,

$$f(X^* \cup X_t) - f(X^* \cup X_t \cap Y_t) < -r(X^* \setminus Y_t). \tag{7}$$

Combining (6) with (7), and noting $X^* \cup X_t \cap Y_t \in [X_t, Y_t]$ and $f(X_p^*) \geq f(X^* \cup X_t \cap Y_t)$, we have,

$$f(X^*) - f(X_p^*) < r(X_t \setminus X^*) - r(X^* \setminus Y_t). \quad (8)$$

We note in (8), $X_t \setminus X^*$ is actually the set of all the elements which are not in X^* but added by A2. Similarly, $X^* \setminus Y_t$ is the set of all the elements which are in X^* but eliminated by A2. Consequently, the performance loss is upper bounded by the total perturbation value of all the mistakenly reduced elements. Since the number of all the mistakenly reduced elements is no more than the number of all the reduced elements nR_t , and the perturbation is generated in $[-t, t]$, we have $r(X_t \setminus X^*) - r(X^* \setminus Y_t) \leq ntR_t$.

For the minimization case, the proof is similar. \square

Note that in Theorem 2, the performance loss is upper bounded by the sum of random variables, which means we can obtain high probability bounds using some concentration inequalities, such as [11].

Theorem 3. (Hoeffding) *Let X_1, X_2, \dots, X_n be independent real-valued random variables such that $\forall i \in \{1, 2, \dots, n\}$, $|X_i| \leq t$. Then with probability $1 - \delta$,*

$$\sum_{i=1}^n X_i - \mathbb{E} \left[\sum_{i=1}^n X_i \right] < t\sqrt{2n \log(1/\delta)}.$$

Theorem 4. *Define $X_*^c \triangleq X_* \cup X_t \cap Y_t$, and $X_c^* \triangleq X^* \cup X_t \cap Y_t$. Denote $M_r \triangleq |X_*^c \Delta X_*|$, and $N_r \triangleq |X_c^* \Delta X^*|$, where $A \Delta B \triangleq (A \setminus B) \cup (B \setminus A)$ is the symmetric difference between the two sets A and B . Then with probability at least $1 - \delta$,*

$$f(X_*^p) - f(X_*) < t\sqrt{2M_r(n + \log(1/\delta))}, \quad (9)$$

$$f(X^*) - f(X_p^*) < t\sqrt{2N_r(n + \log(1/\delta))}. \quad (10)$$

Proof. We prove (10). Since the perturbation vector r has zero expectation value, and each element of r is independently generated. For any fixed $X \subseteq N$, according to Theorem 3, with probability at least $1 - \delta$,

$$r(X) - r(X^*) \leq t\sqrt{2|X \Delta X^*| \log(1/\delta)}. \quad (11)$$

Suppose $X^* \in [S, T]$, and define $m \triangleq |[S, T]|$. Obviously, we have $m \leq 2^n$. Hence,

$$\begin{aligned} & \Pr \left[r(X_c^*) - r(X^*) \geq t\sqrt{2N_r \log(m/\delta)} \right] \\ &= \sum_{X \in [S, T]} \Pr \left[r(X_c^*) - r(X^*) \geq t\sqrt{2|X_c^* \Delta X^*| \log\left(\frac{m}{\delta}\right)}, X_c^* = X \right] \\ &= \sum_{X \in [S, T]} \Pr \left[r(X) - r(X^*) \geq t\sqrt{2|X \Delta X^*| \log\left(\frac{m}{\delta}\right)}, X_c^* = X \right] \\ &\leq \sum_{X \in [S, T]} \Pr \left[r(X) - r(X^*) \geq t\sqrt{2|X \Delta X^*| \log\left(\frac{m}{\delta}\right)} \right] \\ &\leq \sum_{X \in [S, T]} \frac{\delta}{m} = m \frac{\delta}{m} = \delta, \end{aligned}$$

where the first equality holds by the law of total probability. The second equality holds because replacing X_c^* with X in the first expression does not change the event. The first inequality comes from dropping the event $X_c^* = X$ increases the probability. The last line results from (11) and the definition of m . Combining the above result with Theorem 2, and note $r(X_t \setminus X^*) - r(X^* \setminus Y_t) = r(X_c^*) - r(X^*)$, we have, with probability at least $1 - \delta$,

$$f(X^*) - f(X_p^*) < r(X_c^*) - r(X^*) < t\sqrt{2N_r(n + \log(1/\delta))}.$$

Using a similar method, (9) can also be proved. \square

Theorem 4 has an intuitive interpretation. Take P2 and P4 as examples, $\forall Y \subseteq N$, if $f(X^*) - f(Y)$ is large, then it is unlikely that Y is an optimum of P4. Suppose $f(X^*) - f(Y) = \sigma > 0$, then we have,

$$\begin{aligned} & \Pr [f(Y) + r(Y) \geq f(X) + r(X), \forall X \subseteq N] \\ & \leq \Pr [f(Y) + r(Y) \geq f(X^*) + r(X^*)] \\ & = \Pr [r(Y) - r(X^*) \geq \sigma]. \end{aligned}$$

Totally, the probability of Y being an optimum of P4 is upper bounded by the probability that the perturbation difference $r(Y) - r(X^*)$ can compensate the function value difference σ , where the later probability is small when σ is large.

Finally, we show that the M_r and N_r in Theorem 4, which are the numbers of mistakenly reduced elements, can also be upper bounded by functions of f and t .

Theorem 5. *Denote the total number of the mistakenly reduced elements in the first iteration of A1 and A2 as M_r^1 and N_r^1 , respectively. We have,*

$$\mathbb{E}_r[M_r^1] \leq \frac{n}{2} - \frac{F - f(X_*)}{t}, \quad (12)$$

$$\mathbb{E}_r[N_r^1] \leq \frac{n}{2} - \frac{f(X^*) - F}{t}, \quad (13)$$

where $F \triangleq \frac{1}{2}(f(\emptyset) + f(N)) \in [f(X_*), f(X^*)]$.

Proof. For (13), we calculate the total mistakenly reduced element number in expectation in the first iteration of A2. According to the definition of symmetric difference, $N_r^1 = |X^* \setminus Y_1| + |X_1 \setminus X^*|$.

$\forall i \in X^*$, $f(i|\emptyset) \geq f(i|X^* - i) = f(X^*) - f(X^* - i) \geq 0$. And $i \in X^* \setminus Y_1$ iff $f(i|\emptyset) + t < 0$. Similarly, $\forall j \notin X^*$, $f(j|N - j) \leq f(j|X^*) = f(X^* + j) - f(X^*) \leq 0$. And $j \in X_1 \setminus X^*$ iff $f(j|N - j) + t > 0$. Thus we have,

$$\begin{aligned} \mathbb{E}_r[N_r^1] &= \mathbb{E}_r|X^* \setminus Y_1| + \mathbb{E}_r|X_1 \setminus X^*| \\ &= \sum_{i \in X^*} \frac{t - f(i|\emptyset)}{2t} + \sum_{j \notin X^*} \frac{t + f(j|N - j)}{2t} \\ &= \frac{n}{2} - \frac{1}{2t} \left[\sum_{i \in X^*} f(i|\emptyset) - \sum_{j \notin X^*} f(j|N - j) \right] \\ &\leq \frac{n}{2} - \frac{f(X^*) - f(\emptyset) + f(X^*) - f(N)}{2t}. \end{aligned}$$

For (12), the proof is similar. \square

Using similar methods we can obtain the following results, which recover Theorem 5 as a special case.

Theorem 6. *Denote the total number of the mistakenly reduced elements in the k th iteration as M_r^k , and N_r^k , respectively. We have,*

$$\begin{aligned}\mathbb{E}_r[M_r^k] &\leq \frac{n_{k-1}}{2} - \frac{F_{k-1} - f(X_*)}{t}, \\ \mathbb{E}_r[N_r^k] &\leq \frac{n_{k-1}}{2} - \frac{f(X^*) - F_{k-1}}{t},\end{aligned}$$

where $n_{k-1} \triangleq |Y_{k-1} \setminus X_{k-1}|$, and $F_{k-1} \triangleq \frac{1}{2}(f(X_{k-1}) + f(Y_{k-1})) \in [f(X_*), f(X^*)]$.

Theorem 5 implies that the expected number of mistakenly reduced elements in the first iteration will approach $\frac{n}{2}$ as the perturbation scale t increases. This is consistent with the intuition. Let $t \rightarrow \infty$, then $g \rightarrow r$. Each element will be randomly selected to be added or eliminated with probability $\frac{1}{2}$, so the expected number of mistakenly reduced elements is $\frac{n}{2}$.

Remark 1. *When t is large enough, most elements will be reduced in the first iteration, i.e., $N_r \approx N_r^1$. Let $t = \frac{2(f(X^*) - F)}{n(1-2\varepsilon)}$, where $\varepsilon > 0$, by (13) we have $\mathbb{E}_r[N_r] \approx \mathbb{E}_r[N_r^1] \leq \varepsilon n$, which indicates that if $t = \frac{2(f(X^*) - F)}{n(1-2\varepsilon)}$ is a large enough perturbation scale, then the number of mistakenly reduced elements can be desirably upper bounded.*

Remark 2. *Suppose t is large enough and $N_r \approx N_r^1$. With the result of Theorem 2, we have*

$$f(X^*) - f(X_p^*) < tN_r \approx \frac{nt}{2} - (f(X^*) - F).$$

Let $t = \frac{2[(1+\delta)f(X^*) - F]}{n}$ where $\delta > 0$, we have $f(X_p^*) > (1 - \delta)f(X^*)$ from above. This means if there is some relationship between the optimum and the perturbation ($t = \frac{2[(1+\delta)f(X^*) - F]}{n}$ is a large perturbation scale), then the previous performance loss results can be transformed into approximation ratios.

5 EXPERIMENTAL RESULTS

For reducible submodular functions, by incorporating reduction into optimization methods, favorable performance has been achieved [6, 8, 12, 18]. In our experiments, we mainly focus on (nearly) irreducible submodular functions, as listed below.

Subset Selection Function. The objective function [17, 12] is irreducible. Given $M \in \mathbb{R}_+^{n \times n}$, $f(X) \triangleq \sum_{i \in N} \sum_{j \in X} M_{ij} - \lambda \sum_{i,j \in X} M_{ij}$, where $\lambda \in [0.5, 1]$. We set $n = 100$, $\lambda = 0.7$, and randomly generate symmetric matrix M in $(0, 1)^{n \times n}$, and set $M_{ii} = 1$, $\forall i \in N$.

Mutual Information Function. Given n random vectors X_1, X_2, \dots, X_n , define $h(X)$ as the entropy of random variables $\{X_i | i \in N\}$, which is a highly reducible submodular function. The symmetrization [1] of h leads to the mutual information $f(X) \triangleq h(X) + h(N \setminus X)$, which is irreducible. We set $n = 100$, and randomly generate $\{X_i | i = 1, 2, \dots, n\}$.

Log-Determinant Function. Given a positive definite matrix $K \in \mathbb{S}_{++}^n$, the determinant [16] is log-submodular. The symmetrization of log-determinant is $f(X) \triangleq \log \det(K_X) + \log \det(K_{N \setminus X})$, where $K_X \triangleq [K_{ij}]_{i,j \in X}$, $\forall X \subseteq N$. We set $n = 100$. We randomly generate n data points and compute the $n \times n$ similarity matrix as the positive definite matrix K .

Negative Half-Products Function. The objective [2] is $f(X) \triangleq c(X) - \sum_{i,j \in X, i < j} a(i)b(j)$, where a, b, c are non-negative vectors. When c is not non-negative, f can be highly reducible [18]. Here c is non-negative, and f is nearly irreducible. The reduction rate of A1 (A2) is about 1%. We set $n = 100$, and randomly generate a, b in $(0.1, 0.5)^n$ and c in $(1, 5)^n$.

5.1 Perturbation Scale

In Theorem 1, we lower bound the expectation of reduction rate using the expectation of reduction rate after the first iteration of A1 (A2). It is recently reported that for reducible submodular minimization, this bound is often not tight in practice [12]. Given that our method is actually transforming irreducible functions to reducible ones, it is reasonable to borrow experience from reducible cases. We conjecture that relatively small reduction rates after the first iteration would be sufficient for high reduction rates after the last iteration, thereby we only need to choose small perturbation scales to obtain desirable reducibility gains. We empirically verify the conjecture as shown in Figure 1. Appropriate perturbation scales t are chosen so that the reduction rates after the last iteration are changing from 0 to nearly 1. Given a certain perturbation scale, we repeatedly generate r for 10 times and record the average reduction rates of A2 after iteration 1-4 and the last iteration. We observe that A2 terminates within 10 iterations for all objective functions.

We defer similar results for minimization to the supplementary material. As conjectured, we learn from Figure 1 that the gap between the average reduction rates after the first iteration and the last iteration is always large in practice. Hence, we can choose t to get appropriate reduction rates in expectation (e.g., 0.3) after the first iteration, so as to obtain potentially high final reduction rates. Although we can empiri-

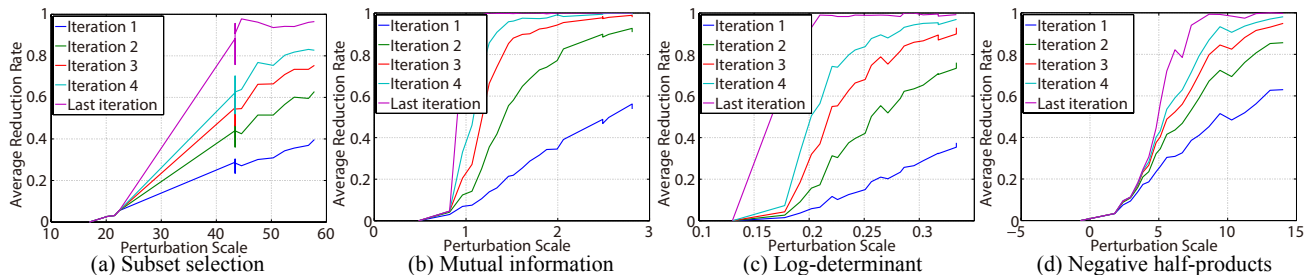


Figure 1: Average Reduction Rates of Maximization

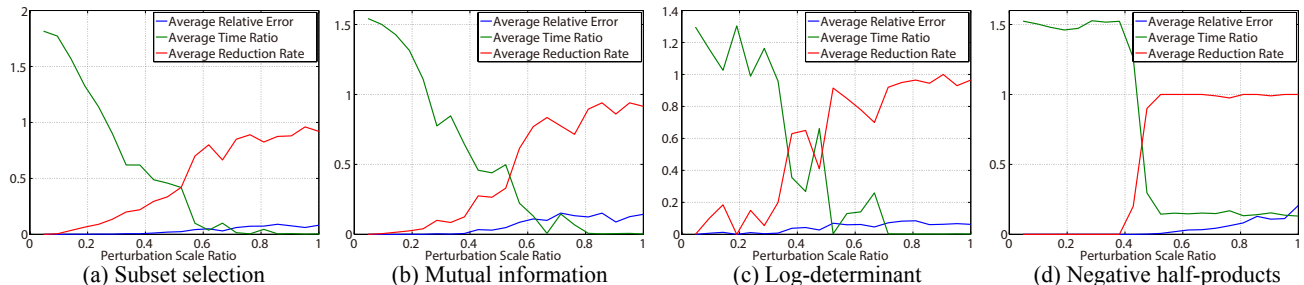


Figure 2: Maximization Results Using Branch-and-Bound Method [9] (Exact Solver)

cally utilize the gap of reducibility gain to choose relative small perturbation scales, we would like to point out that theoretically determining the reduction rates in expectation after the last iteration given certain perturbation scales is still an open problem.

5.2 Optimization Results

We implement our method using SFO toolbox [15].

For maximization, we compare A4 with both exact and approximate methods, as exact methods usually cannot terminate in acceptable time with larger input scales. Denote the outputs of the proposed A4 and the existing method as X^p and X^e , respectively. Also denote the running time as T_p and T_e . We measure the performance loss using *relative error*, which is defined as $E_r \triangleq \frac{|f(X^e) - f(X^p)|}{|f(X^e)|}$. When X^e is exact, $1 - E_r$ is the approximation ratio. We measure the reducibility gain using both the reduction rate and the time ratio T_p/T_e . Small time ratios and relative errors indicate large reducibility gains and small performance losses, respectively.

We employ the branch-and-bound method [9] as the exact solver. Since it has exponential time complexity, we reset $n = 20$ so that it terminates within acceptable time. The results are shown in Figure 2. For comparison, we normalize the perturbation scale as follows. We define $M\{f, [S, T]\} \triangleq \max_{i \in T \setminus S} \max\{f(i|S), -f(i|T - i)\}$, and define the *perturbation scale ratio* as $P(t) \triangleq \frac{t-m}{M-m}$. We change

the perturbation scale t in $[m, M]$ by varying $P(t)$ in $[0, 1]$. We then randomly generate 10 cases for each objective function and record the average relative errors, average reduction rates, and average time ratios for each perturbation scale ratio. Figure 3 shows the results compared with the random bi-directional greedy method [3], which is used as the approximate solver. Note that n is set to 100. For each case, we firstly run A2 once, and then run the random method 5 times on both the original and the reduced lattice, and record the best solutions.

According to Figure 2 and Figure 3, when the perturbation scale ratio is smaller than 0.3, the time ratio is larger than 1. This is because the small reducibility gain cannot make the combination methods more efficient than before. As the perturbation scale ratio increases, the reduction rate increases and the time ratio decreases as expected. Meanwhile, the relative error increases gently when $P(t)$ increases, indicating that there exist useful intervals, in which the perturbation scales can lead to large reducibility gains and small performance losses.

For minimization, since the subset selection and the mutual information function have trivial zero optimal values, *i.e.*, $f(X_*) = f(\emptyset) = 0$, we use the later two as objective functions. We employ the Fujishige-Wolfe minimum-norm point algorithm [7] as the exact solver. All the settings are the same as those of maximization. The results of minimization are shown in Figure 4. We note that for negative half-products function, the

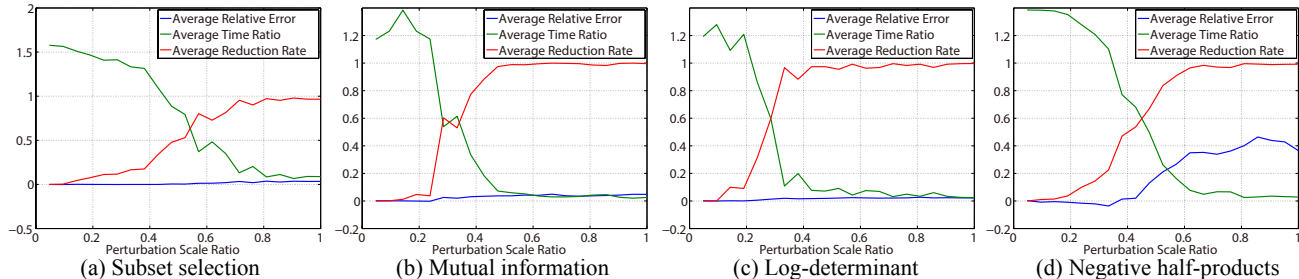


Figure 3: Maximization Results Using Random Bi-directional Greedy [3] (Approximate Solver)

useful interval of perturbation scales is smaller than those of other functions. According to Remark 2, as the marginal gains are relatively large compared to the optimal value, it is inappropriate to choose large perturbation scales in this case.

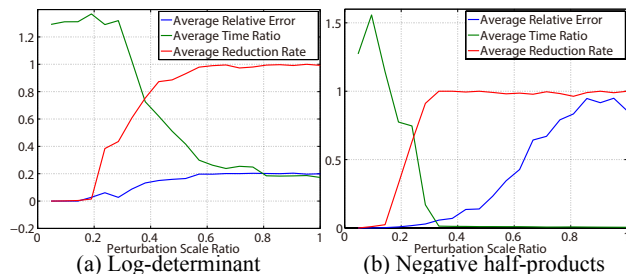


Figure 4: Minimization Results

6 RELATED WORK

In this section, we review some existing works related to solution space reduction for submodular optimization. For P1, Fujishige [6] firstly proves $\mathcal{X}_{\min} \subseteq [A, B]$, where $A = \{i \in N \mid f(i|\emptyset) < 0\}$ and $B = \{j \in N \mid f(j|N - j) \leq 0\}$. Note that actually $[A, B] = [X_1, Y_1]$, which is the working lattice of A1 after its first iteration. Recently, Iyer et al. [12] propose the discrete Majorization-Minimization (MMin) framework for P1. They prove that by choosing appropriate supergradients, MMin is identical with A1. For P2, Golden-gorin [8] proposes the Preliminary Preservation Algorithm (PPA), which is identical with A2. For general cases, Mei et al. [18] prove that the two algorithms work for quasi-submodular functions. Beyond unconstrained problems, for uniform matroid constrained monotone submodular function optimization, Wei et al. [24] propose similar pruning method in which the reduced ground set contains all the original solutions of the *greedy algorithm*.

7 CONCLUSIONS

In this paper, we introduce the reducibility of submodularity, which can improve the efficiency of submodular optimization methods. We then propose the perturbation-reduction framework, and demonstrate its advantages theoretically and empirically. We analyze the reducibility gain and performance loss given perturbation scales. Experimental results show that there exists practically useful intervals, and choosing perturbation scales from them enables us to significantly accelerate the existing methods with only small performance loss. For the future work, we would like to study the reducibility of submodular functions in constrained problems.

Acknowledgements

Jincheng Mei would like to thank Csaba Szepesvári for fixing the proof of Theorem 4. Bao-Liang Lu was supported in part by the grants from the National Natural Science Foundation of China (Grant No.61272248), the National Basic Research Program of China (Grant No.2013CB329401), and the Major Basic Research Program of Shanghai Science and Technology Committee (15JC1400103).

Asterisk indicates the corresponding author.

References

- [1] Francis Bach. Learning with submodular functions: A convex optimization perspective. *Foundations and Trends Machine Learning*, 6(2-3):145–373, 2013.
- [2] Endre Boros and Peter L Hammer. Pseudo-boolean optimization. *Discrete Applied Mathematics*, 123(1):155–225, 2002.
- [3] Niv Buchbinder, Michael Feldman, Joseph Naor, and Roy Schwartz. A tight linear time (1/2)-approximation for unconstrained submodular maximization. In *IEEE Annual Symposium on Foundations of Computer Science*, pages 649–658, 2012.
- [4] Michele Conforti and Gérard Cornuéjols. Submodular set functions, matroids and the greedy algorithm

- tight worst-case bounds and some generalizations of the rado-edmonds theorem. *Discrete Applied Mathematics*, 7(3):251–274, 1984.
- [5] Abhimanyu Das and David Kempe. Submodular meets spectral: Greedy algorithms for subset selection, sparse approximation and dictionary selection. In *International Conference on Machine Learning*, pages 1057–1064, 2011.
- [6] Satoru Fujishige. *Submodular Functions and Optimization*, volume 58. Elsevier, 2005.
- [7] Satoru Fujishige and Shiguo Isotani. A submodular function minimization algorithm based on the minimum-norm base. *Pacific Journal of Optimization*, 7(1):3–17, 2011.
- [8] Boris Goldengorin. Maximization of submodular functions: Theory and enumeration algorithms. *European Journal of Operational Research*, 198(1):102–112, 2009.
- [9] Boris Goldengorin, Gerard Sierksma, Gert A Tijssen, and Michael Tso. The data-correcting algorithm for the minimization of supermodular functions. *Management Science*, 45(11):1539–1551, 1999.
- [10] Daniel Golovin and Andreas Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *Journal of Artificial Intelligence Research*, 42:427–486, 2011.
- [11] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963.
- [12] Rishabh Iyer, Stefanie Jegelka, and Jeff Bilmes. Fast semidifferential-based submodular function optimization. In *International Conference on Machine Learning*, pages 855–863, 2013.
- [13] Rishabh Iyer, Stefanie Jegelka, and Jeff A Bilmes. Curvature and optimal algorithms for learning and minimizing submodular functions. In *Advances in Neural Information Processing Systems*, pages 2742–2750, 2013.
- [14] Stefanie Jegelka, Francis Bach, and Suvrit Sra. Reflection methods for user-friendly submodular optimization. In *Advances in Neural Information Processing Systems*, pages 1313–1321, 2013.
- [15] Andreas Krause. Sfo a toolbox for submodular function optimization. *The Journal of Machine Learning Research*, 11:1141–1144, 2010.
- [16] Alex Kulesza and Ben Taskar. Determinantal point processes for machine learning. *Foundations and Trends in Machine Learning*, 5(2-3):123–286, 2012.
- [17] Hui Lin and Jeff Bilmes. How to select a good training-data subset for transcription submodular active selection for sequences. In *Conference of the International Speech*, pages 2859–2862, 2009.
- [18] Jincheng Mei, Kang Zhao, and Bao-Liang Lu. On unconstrained quasi-submodular function optimization. In *AAAI Conference on Artificial Intelligence*, pages 1191–1197, 2015.
- [19] Baharan Mirzasoleiman, Ashwinkumar Badanidiyuru, Amin Karbasi, Vondrák Jan, and Andreas Krause. Lazier than lazy greedy. In *AAAI Conference on Artificial Intelligence*, pages 1812–1818, 2015.
- [20] Baharan Mirzasoleiman, Amin Karbasi, Rik Sarkar, and Andreas Krause. Distributed submodular maximization: Identifying representative elements in massive data. In *Advances in Neural Information Processing Systems*, pages 2049–2057, 2013.
- [21] Mukund Narasimhan, Nebojsa Jojic, and Jeff A Bilmes. Q-clustering. In *Advances in Neural Information Processing Systems*, pages 979–986, 2005.
- [22] James B Orlin. A faster strongly polynomial time algorithm for submodular function minimization. *Mathematical Programming*, 118(2):237–251, 2009.
- [23] Xinghao Pan, Stefanie Jegelka, Joseph E Gonzalez, Joseph K Bradley, and Michael I Jordan. Parallel double greedy submodular maximization. In *Advances in Neural Information Processing Systems*, pages 118–126, 2014.
- [24] Kai Wei, Rishabh Iyer, and Jeff Bilmes. Fast multi-stage submodular maximization. In *International Conference on Machine Learning*, pages 1494–1502, 2014.