

Fast First-Order Methods for Monotone Strongly DR-Submodular Maximization*

Omid Sadeghi Maryam Fazel
University of Washington

Abstract

Continuous DR-submodular functions are a class of functions that satisfy the Diminishing Returns (DR) property, which implies that they are concave along non-negative directions. Existing works have studied monotone continuous DR-submodular maximization subject to a convex constraint and have proposed efficient algorithms with approximation guarantees. However, in many applications, e.g., computing the stability number of a graph and mean-field inference for probabilistic log-submodular models, the DR-submodular function has the additional property of being *strongly* concave along non-negative directions that could be utilized for obtaining faster convergence rates. In this paper, we first introduce and characterize the class of *strongly DR-submodular* functions and show how such a property implies strong concavity along non-negative directions. Then, we study L -smooth monotone strongly DR-submodular functions that have bounded curvature, and we show how to exploit such additional structure to obtain algorithms with improved approximation guarantees and faster convergence rates for the maximization problem. In particular, we propose the SDRFW algorithm that matches the provably optimal $1 - \frac{\epsilon}{c}$ approximation ratio after only $\lceil \frac{L}{\mu} \rceil$ iterations, where $c \in [0, 1]$ and $\mu \geq 0$ are the curvature and the strong DR-submodularity parameter. Furthermore, we study the Projected Gradient Ascent (PGA) method for this problem and provide a refined analysis of the algorithm with an improved $\frac{1}{1+c}$ approximation ratio (compared to $\frac{1}{2}$ in prior works) and a linear convergence rate. Given that both algorithms require knowledge of the smoothness parameter L , we provide a *novel* characterization of L for DR-submodular functions showing that in many cases, computing L could be formulated as a convex optimization problem, i.e., a geometric program, that could be solved efficiently. Experimental results illustrate and validate the efficiency and effectiveness of our algorithms.

1 Introduction

Submodular set functions are a class of discrete functions that exhibit the Diminishing Returns (DR) property. A set function F defined over the ground set V is called submodular if for all items $j \in V$ and $A \subseteq B \subseteq V \setminus \{j\}$, we have:

$$F(A \cup \{j\}) - F(A) \geq F(B \cup \{j\}) - F(B).$$

In other words, the gain of adding a particular element j to a set decreases as the set gets larger. Such functions are commonly used to quantify coverage and diversity in discrete domains. Submodular set function optimization has found many applications in machine learning such as viral marketing (Kempe et al., 2003), dictionary learning (Das and Kempe, 2011), feature selection for classification (Krause and Guestrin, 2007) and document summarization (Lin and Bilmes, 2011) to name a few.

While the DR property is mostly associated with submodular set functions, it can be defined similarly for continuous functions. A differentiable function $f : \mathcal{K} \rightarrow \mathbb{R}$, $\mathcal{K} \subseteq \mathbb{R}_+^n$, satisfies the DR property if for all $x, y \in \mathcal{K}$ such that $x_i \leq y_i \forall i \in [n]$, we have $\nabla_i f(x) \geq \nabla_i f(y)$ for all $i \in [n]$, i.e., ∇f is an order-reversing mapping. Continuous functions that satisfy the DR property are called DR-submodular. Similar to submodular set functions, continuous DR-submodular functions find applications in multiple domains such as influence and revenue maximization, MAP inference for DPP (Determinantal Point Process) and mean-field inference of probabilistic graphical models (see Bian et al. (2020, Section 6) for more applications and details). While DR-submodular functions are generally non-convex/non-concave, the DR property provides a natural structure that allows designing tractable approximation algorithms. In particular, DR-submodular functions are concave along non-negative directions (Bian et al., 2017), i.e., for all x, y such that $x_i \leq y_i \forall i \in [n]$, we have $f(y) \leq f(x) + \langle \nabla f(x), y - x \rangle$.

Monotone DR-submodular maximization subject to a convex constraint has been previously studied in the literature. Bian et al. (2017) proposed a Frank-Wolfe

*The full version of the paper can be accessed at <https://arxiv.org/abs/2111.07990>

variant that obtains a provably optimal $1 - \frac{1}{e}$ approximation guarantee at a sub-linear convergence rate. Later, Hassani et al. (2017) studied the well-known Projected Gradient Ascent (PGA) method for this problem and proved that PGA has a $\frac{1}{2}$ approximation ratio and sub-linear rate of convergence.

A number of the DR-submodular objective functions in the aforementioned applications of the continuous DR property are indeed strongly concave along non-negative directions. For instance, consider a graph $G = (V, E)$ with adjacency matrix A . Computing the stability number $s(G)$ of the graph (i.e., cardinality of the largest subset of vertices such that no two vertices in this subset are adjacent) is a well-known NP-hard combinatorial problem. This problem could be formulated as $s(G)^{-1} = \min_{x \in \Delta} x^T(A + I)x$ where $\Delta = \{x \in \mathbb{R}^{|V|} : \mathbf{1}^T x = 1, x_v \geq 0 \forall v \in V\}$ is the standard simplex and I is the identity matrix (Motzkin and Straus, 1965). We can rewrite the problem as:

$$\begin{aligned} s(G)^{-1} &= \min_{x \in \Delta} x^T(A + I)x \\ &= -\max_{x \in \Delta} x^T(-A - I)x \\ &= 2 - \max_{x \in \Delta} x^T(-A - I)x + \underbrace{2\mathbf{1}^T x}_{=2}. \end{aligned}$$

It will soon be clear that the function $x^T(-A - I)x + 2\mathbf{1}^T x$ is 2-strongly DR-submodular (see Definition 2.1 in Section 2) and thus, finding the stability number of a graph could be formulated as a convex-constrained monotone 2-strongly DR-submodular maximization problem. Another example is the mean-field inference problem for probabilistic log-submodular models (Section 6.6 of Bian et al. (2020)). Let $F : 2^V \rightarrow \mathbb{R}$ be a submodular set function. Consider distributions over subsets $S \subseteq V$ of the form $P(S) = \frac{1}{Z} \exp(F(S))$ where the normalizing quantity $Z = \sum_{S \subseteq V} \exp(F(S))$ is called the *partition function*. In general, computing the partition function involves summing over exponential number of terms and is not computationally feasible. Alternatively, one can use mean-field inference to approximate $P(S)$ by a completely factorized distribution $Q(S)$, i.e., elements $i \in V$ are picked independently and $Q(S|x) = \prod_{i \in S} x_i \prod_{j \notin S} (1 - x_j)$ where $x \in [0, 1]^V$ is the vector of marginals, by minimizing the KL divergence $\mathbf{KL}(x) = \sum_{S \subseteq V} Q(S|x) \ln \left(\frac{Q(S|x)}{P(S)} \right)$ between the two distributions. $\mathbf{KL}(x)$ can be written as:

$$\begin{aligned} \mathbf{KL}(x) &= - \sum_{S \subseteq V} F(S) \prod_{i \in S} x_i \prod_{j \notin S} (1 - x_j) \\ &\quad + \sum_{i=1}^{|V|} (x_i \ln x_i + (1 - x_i) \ln(1 - x_i)) + \ln Z. \end{aligned}$$

It will be clear later that the problem $\max_{x \in [0, 1]^V} -\mathbf{KL}(x)$ (equivalent to minimizing $\mathbf{KL}(x)$) is a 4-strongly DR-submodular maximization problem.

1.1 Contributions In this paper, we precisely characterize the class of *strongly DR-submodular* functions, i.e., DR-submodular functions that are strongly concave along non-negative directions. We consider the optimization problem $\max_{x \in \mathcal{K}} f(x)$ where \mathcal{K} is a convex set and f is a monotone, smooth and strongly DR-submodular function with bounded curvature $c_f \in [0, 1]$, and we provide algorithms with refined approximation guarantees that exploit the strong DR-submodularity structure of the objective function. Specifically, we make the following contributions:

- We propose the SDRFW algorithm in Section 3 that obtains the approximation guarantee $1 - \frac{c_f}{e}$ after $\lceil \frac{L}{\mu} \rceil$ iterations, where L is the smoothness parameter and μ is the strong DR-submodularity parameter of the function f .
 - In Section 4, we analyze PGA for our problem and we provide a refined and sharper analysis of the algorithm showing that PGA has an improved $\frac{1}{1+c_f}$ approximation guarantee (compared to $\frac{1}{2}$ in prior works) at a linear convergence rate, i.e., it only takes $\mathcal{O}(\ln(\frac{1}{\epsilon}))$ iterations to get within ϵ of $\frac{1}{1+c_f}$ OPT, where OPT denotes the optimal value of the problem.
 - We also study *online* strongly DR-submodular maximization in Section 4.1. We analyze the online counterpart of PGA, called Online Gradient Ascent (OGA), for this problem and we show how our techniques could be used to obtain improved logarithmic bounds for the $(\frac{1}{1+c})$ -regret of the algorithm.
 - Given that both SDRFW and PGA require knowledge of the smoothness parameter L , in Section 5, we provide a *novel* characterization of L for (strongly) DR-submodular functions showing that in many cases, computing L could be formulated as a convex optimization problem, i.e., a geometric program, that could be solved efficiently. Therefore, we can compute L in an efficient manner before running the algorithms.
- Finally, we test our algorithms on a class of strongly DR-submodular functions and also for the problem of computing the stability number of graphs in Section 6. A summary of our results is provided in Table 1.

2 Preliminaries

2.1 Notation The set $\{1, 2, \dots, n\}$ is denoted by $[n]$. For a vector $x \in \mathbb{R}^n$, x_i is used to denote the i -th entry of x . Similarly, for a matrix $A \in \mathbb{R}^{n \times n}$, we use A_{ij} to indicate the entry in the i -th row and j -th column of

| Paper | Setting | Approximation ratio | Convergence rate |
|--------------------------|-----------|---------------------|--|
| Bian et al. (2017) | $\mu = 0$ | $1 - \frac{1}{e}$ | sub-linear |
| Hassani et al. (2017) | $\mu = 0$ | $\frac{1}{2}$ | sub-linear |
| This work (PGA) | $\mu = 0$ | $\frac{1}{1+c_f}$ | sub-linear |
| This work (PGA) | $\mu > 0$ | $\frac{1}{1+c_f}$ | linear |
| This work (SDRFW) | $\mu > 0$ | $1 - \frac{c_f}{e}$ | after $\lceil \frac{L}{\mu} \rceil$ iterations |

Table 1: Comparison of our results and the prior works.

the matrix. The inner product of two vectors $x, y \in \mathbb{R}^n$ is denoted by either $\langle x, y \rangle$ or $x^T y$. Moreover, for two vectors $x, y \in \mathbb{R}^n$, we have $x \preceq y$ if $x_i \leq y_i$ holds for every $i \in [n]$. Also, we use $x \vee y = \max\{x, y\}$ and $x \wedge y = \min\{x, y\}$ to denote the component-wise maximum and minimum of $x, y \in \mathbb{R}^n$, i.e., for all $i \in [n]$, $[x \vee y]_i = \max\{x_i, y_i\}$ and $[x \wedge y]_i = \min\{x_i, y_i\}$ holds. $\|\cdot\|$ indicates the Euclidean norm by default. We use $\text{Proj}_{\mathcal{K}}(x)$ to denote the Euclidean projection of x onto the convex set \mathcal{K} , i.e., $\text{Proj}_{\mathcal{K}}(x) = \arg \min_{z \in \mathcal{K}} \|z - x\|$. A function $f : \mathcal{K} \rightarrow \mathbb{R}$ is called β -Lipschitz if for all $x, y \in \mathcal{K}$, we have $|f(y) - f(x)| \leq \beta \|y - x\|$. Also, f is called monotone if for all $x, y \in \mathcal{K}$ such that $x \preceq y$, $f(x) \leq f(y)$ holds. The diameter of a set \mathcal{K} is defined as $R = \max_{x, y \in \mathcal{K}} \|y - x\|$. By default, $f(\cdot)$ indicates a continuous function and $F(\cdot)$ denotes a discrete set function.

2.2 DR-submodular functions A differentiable function $f : \mathcal{K} \rightarrow \mathbb{R}$, $\mathcal{K} \subset \mathbb{R}_+^n$, is called DR-submodular if for all x, y such that $x \preceq y$, $\nabla f(x) \succeq \nabla f(y)$ holds. If f is twice differentiable, DR-submodularity is equivalent to the Hessian matrix $\nabla^2 f(x)$ being entry-wise non-positive for all $x \in \mathcal{K}$. For instance, for a submodular set function F over the ground set V , the multilinear extension $f : [0, 1]^{|V|} \rightarrow \mathbb{R}$ of F defined as $f(x) = \sum_{S \subset V} F(S) \prod_{i \in S} x_i \prod_{j \notin S} (1 - x_j)$ is DR-submodular. While DR-submodularity and concavity are equivalent for the special case of $n = 1$, DR-submodular functions are generally non-concave. Nonetheless, an important consequence of DR-submodularity is concavity along non-negative directions (Bian et al., 2017), i.e., for all x, y such that $x \preceq y$, we have $f(y) \leq f(x) + \langle \nabla f(x), y - x \rangle$.

2.3 Strongly DR-submodular functions We define the class of strongly DR-submodular functions below.

DEFINITION 2.1. For $\mu \geq 0$, we call a differentiable function $f : \mathcal{K} \rightarrow \mathbb{R}$, $\mathcal{K} \subset \mathbb{R}_+^n$, to be μ -strongly DR-submodular if any of the following equivalent properties hold:

- $f(\cdot) + \frac{\mu}{2} \|\cdot\|^2$ is DR-submodular.

- For all $x, y \in \mathcal{K}$ such that $x \preceq y$, we have

$$\nabla f(x) \succeq \nabla f(y) + \mu(y - x).$$

- If f is twice differentiable, $\nabla_{ii}^2 f(x) \leq -\mu \forall i \in [n]$ and $\nabla_{ij}^2 f(x) \leq 0 \forall i \neq j$ holds for all $x \in \mathcal{K}$.

We provide two classes of strongly DR-submodular functions below:

- **Indefinite quadratic functions.** Let $f(x) = \frac{1}{2} x^T H x + h^T x + c$ where H is a symmetric matrix. If H is entry-wise non-positive, f is a DR-submodular function and if in addition, $H_{ii} \leq -\mu$ holds for all $i \in [n]$, f is μ -strongly DR-submodular. Such quadratic utility functions have a wide range of applications. As an example, consider the problem of computing the stability number of a graph (introduced in Section 1). The Hessian of the objective function $x^T(-A - I)x + 21^T x$ is $H = -2A - 2I$. Given that $H_{ii} = -2$ and $H_{ij} \leq 0$ for all $i \neq j \in [n]$, the objective function is 2-strongly DR-submodular. In addition, price optimization with continuous prices (Ito and Fujimaki, 2016) is also a non-concave (strongly) DR-submodular quadratic optimization problem.

- **Concave functions with negative dependence.** Let $d \geq 2$. If $h_i : \mathbb{R}_+ \rightarrow \mathbb{R}$ is (strongly) concave for all $i \in [n]$ and $\theta_{i_1, \dots, i_r} \leq 0$ for all $r \in [d]$ and $(i_1, \dots, i_r) \subseteq [n]$, the following function $f : \mathbb{R}_+^n \rightarrow \mathbb{R}$ is (strongly) DR-submodular:

$$f(x) = \sum_{i=1}^n h_i(x_i) + \sum_{(i,j):i \neq j} \theta_{ij} x_i x_j + \dots \\ + \sum_{(i_1, \dots, i_d): i_r \neq i_s \forall r, s \in [d]} \theta_{i_1, \dots, i_d} x_{i_1} \dots x_{i_d}.$$

We can prove that μ -strongly DR-submodular functions are indeed μ -strongly concave along non-negative directions. This property is extensively used in the design and analysis of our algorithms.

LEMMA 2.1. If $f : \mathcal{K} \rightarrow \mathbb{R}$, $\mathcal{K} \subset \mathbb{R}_+^n$, is differentiable and μ -strongly DR-submodular, for all $x \in \mathcal{K}$ and $v \succeq 0$ or $v \preceq 0$, the following holds:

$$f(x + v) \leq f(x) + \langle \nabla f(x), v \rangle - \frac{\mu}{2} \|v\|^2.$$

Proof. Without loss of generality, assume $v \succeq 0$ holds (analysis for $v \preceq 0$ is similar). For any $z \in \mathcal{K}$, $v \succeq 0$ and $i \in [n]$, we have:

$$\begin{aligned} [\nabla^2 f(z)v]_i + \mu v_i &= \nabla_{ii}^2 f(z)v_i + \sum_{j \neq i} \nabla_{ij}^2 f(z)v_j + \mu v_i \\ &= \underbrace{(\nabla_{ii}^2 f(z) + \mu)}_{\leq 0} \underbrace{v_i}_{\geq 0} + \sum_{j \neq i} \underbrace{\nabla_{ij}^2 f(z)}_{\leq 0} \underbrace{v_j}_{\geq 0} \\ &\leq 0. \end{aligned}$$

Therefore, $\nabla^2 f(z)v + \mu v \preceq 0$ holds. We can use the mean value theorem to write:

$$\begin{aligned} f(x+v) - f(x) - \langle \nabla f(x), v \rangle &= \int_0^1 \langle \nabla f(x+tv), v \rangle dt - \langle \nabla f(x), v \rangle \\ &= \int_0^1 \langle \nabla f(x+tv) - \nabla f(x), v \rangle dt \\ &= \int_0^1 \sum_{i=1}^n t \underbrace{([\nabla^2 f(z_{i,t})v]_i + \mu v_i)}_{\leq 0} \underbrace{v_i}_{\geq 0} dt - \underbrace{\mu \int_0^1 t \langle v, v \rangle dt}_{= \frac{\mu}{2} \|v\|^2} \\ &\leq \frac{-\mu}{2} \|v\|^2, \end{aligned}$$

where $z_{i,t} \forall i \in [n]$ is in the line segment between x and $x+tv$. Thus, f is μ -strongly concave along the non-negative direction v . \square

Bian et al. (2020) defined μ -strongly DR-submodular functions to be the class of functions that are μ -strongly concave along non-negative directions. Note that our definition of μ -strong DR-submodularity is a stronger condition than μ -strong concavity along non-negative directions. For instance, μ -strongly concave functions are μ -strongly concave along any direction, but they may not even be DR-submodular.

2.4 Smooth functions A differentiable function $f: \mathcal{K} \rightarrow \mathbb{R}$, $\mathcal{K} \subset \mathbb{R}_+^n$, is called L -smooth if for all $x, y \in \mathcal{K}$, we have

$$f(y) \geq f(x) + \langle \nabla f(x), y - x \rangle - \frac{L}{2} \|y - x\|^2.$$

If f is twice differentiable, there is an equivalent definition of smoothness: f is L -smooth if $\nabla^2 f(x) \succeq -LI$ holds for all $x \in \mathcal{K}$ where I is the identity matrix. In other words, the smallest eigenvalue of the Hessian of f is uniformly lower bounded by $-L$ everywhere. Combining the definition of smooth functions and the result of Lemma 2.1, it is clear that for a μ -strongly DR-submodular and L -smooth function, $\mu \leq L$ holds.

2.5 Curvature We define the notion of curvature for monotone continuous functions below.

DEFINITION 2.2. Given a monotone differentiable function $f: \mathcal{K} \rightarrow \mathbb{R}$, $\mathcal{K} \subset \mathbb{R}_+^n$, we define the curvature of f as follows:

$$c_f = 1 - \inf_{x, y \in \mathcal{K}, i \in [n]} \frac{\nabla_i f(y)}{\nabla_i f(x)}.$$

If f is DR-submodular and $0 \in \mathcal{K}$, we have $c_f = 1 - \inf_{x \in \mathcal{K}, i \in [n]} \frac{\nabla_i f(x)}{\nabla_i f(0)}$.

It is easy to see that $c_f \in [0, 1]$ holds for all monotone f . $c_f \leq 1$ is due to monotonicity of f (i.e., ∇f being non-negative) and $c_f \geq 0$ follows from setting $x = y$ in the definition. If $c_f = 0$, f is linear and larger c_f corresponds to f being more curved.

A similar notion of curvature was introduced in Sessa et al. (2019). The definition is inspired by the curvature of submodular set functions (Conforti and Cornuéjols, 1984). In fact, if f is the multilinear extension of a monotone submodular set function F , c_f coincides with the curvature of F (Sadeghi and Fazel, 2021). Submodular set function maximization with bounded curvature has been widely studied in the literature. Conforti and Cornuéjols (1984) showed that the greedy algorithm applied to the monotone submodular set function maximization problem subject to a cardinality constraint has a $\frac{1}{\kappa}(1 - e^{-\kappa})$ approximation ratio, where κ is the curvature of the set function. More recently, Sviridenko et al. (2017) proposed two approximation algorithms for the more general problem of monotone submodular maximization subject to a matroid constraint and obtained a $1 - \frac{\kappa}{e}$ approximation ratio for these two algorithms. They also provided matching upper bounds for this problem showing that the $1 - \frac{\kappa}{e}$ approximation ratio is indeed optimal. Later on, Feldman (2021) managed to obtain the same $1 - \frac{\kappa}{e}$ approximation ratio with an algorithm that is much faster than the ones proposed by Sviridenko et al. (2017).

3 Strongly DR-submodular Frank-Wolfe (SDRFW) algorithm

In this section, we propose the SDRFW algorithm for strongly DR-submodular maximization with bounded curvature. Throughout the section, we make the further assumption that the domain set \mathcal{K} contains the origin, i.e., $0 \in \mathcal{K}$. Furthermore, we consider f to be a monotone, μ -strongly DR-submodular and L -smooth function with curvature $c_f \in [0, 1]$. Without loss of generality, we also assume that f is normalized, i.e., $f(0) = 0$. For the DR-submodular setting ($\mu = 0$), Bian et al. (2017) proposed a Frank-Wolfe variant for

solving the problem. Starting from $x_0 = 0$, their algorithm performs K Frank-Wolfe updates where at each iteration $k \in \{0, \dots, K-1\}$, it finds v_k such that $v_k = \arg \max_{x \in \mathcal{K}} \langle x, \nabla f(x_k) \rangle$, performs the update $x_{k+1} = x_k + \frac{1}{K}v_k$ and outputs x_K .

Define $g(x) = f(x) - \ell^T x$ where for all $i \in [n]$, $\ell_i = \min_x \nabla_i f(x)$. Note that similar to f , g is also a normalized, monotone, μ -strongly DR-submodular and L -smooth function.

The SDRFW algorithm is presented in Algorithm 3.1. First, note that the output of the algorithm ($x = x_K$) is the average of K points $\{v_k\}_{k=0}^{K-1}$ in the convex domain set \mathcal{K} , and therefore, $x \in \mathcal{K}$. Also, it is noteworthy that the update rule for $\{v_k\}_{k=1}^K$ can be computed efficiently in many cases. To see this, for all $k \in \{0, \dots, K-1\}$, we can equivalently write:

$$v_k = \text{Proj}_{\mathcal{K}} \left(\frac{1}{\mu} \nabla g(x_k) + \frac{1}{\mu(1 - \frac{1}{K})^{K-k-1}} \ell \right).$$

In many cases, such projection could be computed in linear time $\mathcal{O}(n)$ (Brucker, 1984; Pardalos and Kooroor, 1990), e.g., for $\mathcal{K} = \{x \in \mathbb{R}^n : 1^T x \leq 1, 0 \leq x \leq 1\}$.

Algorithm 3.1 is different from the Frank-Wolfe variant of Bian et al. (2017) in two important aspects:

1) At step $k \in \{0, \dots, K-1\}$, Algorithm 3.1 is applied to the function $(1 - \frac{1}{K})^{K-k-1} g(\cdot) + \langle \ell, \cdot \rangle$ (instead of $f(\cdot) = g(\cdot) + \langle \ell, \cdot \rangle$),

2) The linear maximization step for computing $\{v_k\}_{k=0}^{K-1}$ in the Frank-Wolfe variant of Bian et al. (2017) is replaced by a strongly concave maximization problem. Modification 1 is inspired by a similar idea in Feldman (2021) where they provided an algorithm for the setting where the objective function is the multilinear extension of a submodular set function ($\mu = 0$) and obtained a $1 - \frac{c_f}{e}$ approximation ratio for the problem. They also proved matching negative results showing that no polynomial time algorithm can perform better in terms of the approximation ratio. The same upper bound applies to our framework as well. However, the additional strong DR-submodularity of f allows for a faster convergence to $(1 - \frac{c_f}{e})\text{OPT}$. We provide the approximation guarantee of Algorithm 3.1 below.

THEOREM 3.1. *Let $f : \mathcal{K} \rightarrow \mathbb{R}$, $\mathcal{K} \subset \mathbb{R}_+^n$ and $0 \in \mathcal{K}$, be a normalized, monotone, μ -strongly DR-submodular and L -smooth function. If we set $K = \lceil \frac{L}{\mu} \rceil$, the output of Algorithm 3.1 has the following performance guarantee:*

$$f(x) \geq (1 - \frac{c_f}{e})f(x^*),$$

where $x^* = \arg \max_{x \in \mathcal{K}} f(x)$.

Proof. Define $g(x) = f(x) - \ell^T x$ where for all $i \in [n]$, $\ell_i = \min_x \nabla_i f(x)$. Note that similar to f , g is also a

ALGORITHM 3.1. SDRFW

Input: Convex set \mathcal{K} with $0 \in \mathcal{K}$, L -smooth and μ -strongly DR-submodular f , $g(\cdot) = f(\cdot) - \langle \ell, \cdot \rangle$ where for all $i \in [n]$, $\ell_i = \min_x \nabla_i f(x)$, $K > 0$.

Set $x_0 = 0$.

for $k = 0$ **to** $K - 1$ **do**

Set $v_k = \arg \max_{x \in \mathcal{K}} \langle (1 - \frac{1}{K})^{K-k-1} \nabla g(x_k) + \ell, x \rangle - \frac{\mu(1 - \frac{1}{K})^{K-k-1}}{2} \|x\|^2$.

$x_{k+1} = x_k + \frac{1}{K}v_k$.

end for

Output: $x = x_K$.

normalized, monotone, μ -strongly DR-submodular and L -smooth function. For instance, in order to verify the μ -strong concavity of g along non-negative directions (which will be used in the proof), for all $x \in \mathcal{K}$ and $v \succeq 0$ or $v \preceq 0$, we can write:

$$\begin{aligned} g(x+v) &= f(x+v) - \langle \ell, x+v \rangle \\ &\leq f(x) + \langle \nabla f(x), v \rangle - \frac{\mu}{2} \|v\|^2 - \langle \ell, x+v \rangle \\ &= \underbrace{f(x) - \langle \ell, x \rangle}_{=g(x)} + \underbrace{\langle \nabla f(x) - \ell, v \rangle}_{= \nabla g(x)} - \frac{\mu}{2} \|v\|^2 \\ &= g(x) + \langle \nabla g(x), v \rangle - \frac{\mu}{2} \|v\|^2, \end{aligned}$$

where the inequality uses the μ -strong DR-submodularity of f . Therefore, g is μ -strongly concave along non-negative directions.

Now, we move on to prove the theorem. For all $k \in \{0, \dots, K-1\}$, we can write:

$$g(x_{k+1}) \geq g(x_k) + \frac{1}{K} \langle \nabla g(x_k), v_k \rangle - \frac{L}{2K^2} \|v_k\|^2.$$

Rearranging the terms, we have:

$$(3.1) \quad g(x_{k+1}) - g(x_k) \geq \frac{1}{K} \langle \nabla g(x_k), v_k \rangle - \frac{L}{2K^2} \|v_k\|^2.$$

For all $k \in \{0, \dots, K\}$, define the function $\phi(k)$ as follows:

$$\phi(k) = (1 - \frac{1}{K})^{K-k} g(x_k) + \langle \ell, x_k \rangle.$$

For a fixed $k \in \{0, \dots, K-1\}$, we have:

$$\begin{aligned}
& K(\phi(k+1) - \phi(k)) \\
&= K\left(1 - \frac{1}{K}\right)^{K-k-1}g(x_{k+1}) - K\left(1 - \frac{1}{K}\right)^{K-k}g(x_k) + \langle \ell, v_k \rangle \\
&= \frac{\left(1 - \frac{1}{K}\right)^{K-k-1}(g(x_{k+1}) - g(x_k))}{1/K} \\
&+ \left(1 - \frac{1}{K}\right)^{K-k-1}g(x_k) + \langle \ell, v_k \rangle \\
&\stackrel{(a)}{\geq} \frac{\left(1 - \frac{1}{K}\right)^{K-k-1}\left(\frac{1}{K}\langle \nabla g(x_k), v_k \rangle - L\|v_k\|^2/2K^2\right)}{1/K} \\
&+ \left(1 - \frac{1}{K}\right)^{K-k-1}g(x_k) + \langle \ell, v_k \rangle \\
&= \left(1 - \frac{1}{K}\right)^{K-k-1}\left(\langle \nabla g(x_k), v_k \rangle - \frac{L}{2K}\|v_k\|^2\right) \\
&+ \left(1 - \frac{1}{K}\right)^{K-k-1}g(x_k) + \langle \ell, v_k \rangle \\
&= \left\langle \left(1 - \frac{1}{K}\right)^{K-k-1}\nabla g(x_k) + \ell, v_k \right\rangle \\
&- \left(1 - \frac{1}{K}\right)^{K-k-1}\frac{L}{2K}\|v_k\|^2 + \left(1 - \frac{1}{K}\right)^{K-k-1}g(x_k) \\
&\stackrel{(b)}{\geq} \left(1 - \frac{1}{K}\right)^{K-k-1}\langle \nabla g(x_k), x^* \rangle + \langle \ell, x^* \rangle \\
&+ \frac{\mu\left(1 - \frac{1}{K}\right)^{K-k-1}}{2}\|v_k\|^2 - \frac{\mu\left(1 - \frac{1}{K}\right)^{K-k-1}}{2}\|x^*\|^2 \\
&- \left(1 - \frac{1}{K}\right)^{K-k-1}\frac{L}{2K}\|v_k\|^2 + \left(1 - \frac{1}{K}\right)^{K-k-1}g(x_k),
\end{aligned}$$

where (a) follows from inequality 3.1 and (b) is due to the update rule of the SDRFW algorithm for v_k . We can use the monotonicity and strong DR-submodularity of g respectively to write:

$$\begin{aligned}
g(x^*) - g(x_k) &\leq g(x^* + x_k) - g(x_k) \\
&\leq \langle \nabla g(x_k), x^* \rangle - \frac{\mu}{2}\|x^*\|^2.
\end{aligned}$$

Putting the above two inequalities together, we have:

$$\begin{aligned}
& K(\phi(k+1) - \phi(k)) \\
&\geq \left(1 - \frac{1}{K}\right)^{K-k-1}\left(g(x^*) - g(x_k) + \frac{\mu}{2}\|x^*\|^2\right) + \langle \ell, x^* \rangle \\
&+ \frac{\mu\left(1 - \frac{1}{K}\right)^{K-k-1}}{2}\|v_k\|^2 - \frac{\mu\left(1 - \frac{1}{K}\right)^{K-k-1}}{2}\|x^*\|^2 \\
&- \left(1 - \frac{1}{K}\right)^{K-k-1}\frac{L}{2K}\|v_k\|^2 + \left(1 - \frac{1}{K}\right)^{K-k-1}g(x_k) \\
&= \left(1 - \frac{1}{K}\right)^{K-k-1}g(x^*) + \langle \ell, x^* \rangle \\
&+ \left(1 - \frac{1}{K}\right)^{K-k-1}\left(\frac{\mu}{2} - \frac{L}{2K}\right)\|v_k\|^2
\end{aligned}$$

Therefore, if we set $K = \lceil \frac{L}{\mu} \rceil$ and divide both sides by K , we obtain:

$$\phi(k+1) - \phi(k) \geq \frac{1}{K}\left(1 - \frac{1}{K}\right)^{K-k-1}g(x^*) + \frac{1}{K}\langle \ell, x^* \rangle.$$

Applying the inequality for all $k \in \{0, \dots, K-1\}$ and taking the sum, we have:

$$\begin{aligned}
f(x_K) &= g(x_K) + \langle \ell, x_K \rangle \\
&= \phi(K) - \phi(0) \\
&\geq \left(\frac{1}{K} \sum_{k=0}^{K-1} \left(1 - \frac{1}{K}\right)^{K-k-1}\right)g(x^*) + \langle \ell, x^* \rangle \\
&= \left(\frac{1}{K} \frac{1 - \left(1 - \frac{1}{K}\right)^K}{1/K}\right)g(x^*) + \langle \ell, x^* \rangle \\
&\geq \left(1 - \frac{1}{e}\right)g(x^*) + \langle \ell, x^* \rangle,
\end{aligned}$$

where the last inequality uses $\left(1 - \frac{1}{K}\right)^K \leq \frac{1}{e}$. Therefore, $f(x) \geq \left(1 - \frac{1}{e}\right)g(x^*) + \langle \ell, x^* \rangle$ holds.

Using the mean value theorem and considering that $f(0) = 0$, we have $f(x^*) = \sum_{i=1}^n \nabla_i f(tx^*)x_i^*$ where $t \in [0, 1]$. Therefore, we can use the definition of ℓ and the curvature to write:

$$\ell^T x^* = \sum_{i=1}^n \ell_i x_i^* \geq (1 - c_f) \sum_{i=1}^n \nabla_i f(tx^*)x_i^* = (1 - c_f)f(x^*).$$

Putting the above inequalities together, we have:

$$\begin{aligned}
f(x) &\geq \left(1 - \frac{1}{e}\right)g(x^*) + \ell^T x^* \\
&= \left(1 - \frac{1}{e}\right)f(x^*) - \left(1 - \frac{1}{e}\right)\ell^T x^* + \ell^T x^* \\
&= \left(1 - \frac{1}{e}\right)f(x^*) + \frac{1}{e}\ell^T x^* \\
&\geq \left(1 - \frac{1}{e}\right)f(x^*) + \frac{1 - c_f}{e}f(x^*) \\
&= \left(1 - \frac{c_f}{e}\right)f(x^*).
\end{aligned}$$

□

In comparison, the Frank-Wolfe variant of Bian et al. (2017) obtains the approximation guarantee $f(x) \geq \left(1 - \frac{1}{e}\right)f(x^*) - \frac{LR^2}{2K}$ where $R = \max_{x,y \in \mathcal{K}} \|y - x\|$ is the diameter of \mathcal{K} . Therefore, Algorithm 3.1 has an improved approximation ratio for objective functions f with curvature $c_f < 1$ and its guarantee does not deteriorate as the diameter of \mathcal{K} becomes larger. Intuitively, for K chosen large enough (i.e., $K \geq \frac{L}{\mu}$) in the analysis of Algorithm 3.1, the $-\frac{LR^2}{2K}$ term is cancelled with a positive term resulting from μ -strong DR-submodularity of the objective function.

We need to know the smoothness parameter L and the strong DR-submodularity parameter μ to set K in Algorithm 3.1. For μ -strongly DR-submodular functions, if we run Algorithm 3.1 with $\hat{\mu}$ instead of μ , the algorithm obtains the same guarantee as long as $\hat{\mu} \leq \mu$. In

ALGORITHM 3.2. PGA

Input: Convex set \mathcal{K} , $x_1 \in \mathcal{K}$, L -smooth and μ -strongly DR-submodular f , $K > 0$.

for $k = 1$ **to** K **do**

 Set $x_{k+1} = \text{Proj}_{\mathcal{K}}(x_k + \frac{1}{L}\nabla f(x_k))$.

end for

Output: $x = x_{K+1}$.

order to compute such a lower bound, one can investigate the diagonal entries of the Hessian matrix. In Section 5, we show how to compute L efficiently using convex optimization tools.

4 Projected Gradient Ascent (PGA) algorithm

In this section, we study the well-known Projected Gradient Ascent (PGA) method for strongly DR-submodular maximization with bounded curvature. The PGA algorithm is provided in Algorithm 3.2 (Nesterov, 2003). Given an initial point $x_1 \in \mathcal{K}$, PGA iteratively applies the update $x_{k+1} = \text{Proj}_{\mathcal{K}}(x_k + \frac{1}{L}\nabla f(x_k))$. In other words, at each iteration $k \in [K]$, the current iterate x_k is updated by adding $\frac{1}{L}\nabla f(x_k)$ and the resulting point is then projected back to the constraint set \mathcal{K} . The algorithm outputs the final iterate $x = x_{K+1}$. Unlike the SDRFW algorithm, PGA does not require to start from the origin and for any feasible initial point $x_1 \in \mathcal{K}$, PGA still converges to a competitive solution. However, as we will soon see in the result of Theorem 4.1, the rate of convergence depends on the distance between the initial point x_1 and the optimal point x^* . We first provide a key lemma below that is used in the analysis of Algorithm 3.2.

LEMMA 4.1. *For any $x, z \in \mathcal{K}$, if f is a non-negative monotone μ -strongly DR-submodular function with curvature c_f , we have:*

$$f(z) - (1 + c_f)f(x) \leq \langle \nabla f(x), z - x \rangle - \frac{\mu}{2}\|z - x\|^2.$$

Proof. Let $u = x \vee z$ and $w = x \wedge z$. Using the μ -strong DR-submodularity property of f , we can write:

$$\begin{aligned} f(u) - f(x) &\leq \langle \nabla f(x), u - x \rangle - \frac{\mu}{2}\|u - x\|^2, \\ f(w) - f(x) &\leq \langle \nabla f(x), w - x \rangle - \frac{\mu}{2}\|w - x\|^2. \end{aligned}$$

Taking the sum of the two inequalities and using the fact that $u + w = x + z$, we have:

$$(4.1) \quad f(u) + f(w) - 2f(x) \leq \langle \nabla f(x), z - x \rangle - \frac{\mu}{2}\|z - x\|^2.$$

Using the mean value theorem, we can write:

$$\begin{aligned} f(u) - f(z) &= \int_0^1 \langle u - z, \nabla f(z + t(u - z)) \rangle dt, \\ f(w) - f(x) &= \int_0^1 \langle w - x, \nabla f(w + t(x - w)) \rangle dt. \end{aligned}$$

Given that $x - w = u - z$ and $\nabla_i f(z + t(u - z)) \geq (1 - c_f)\nabla_i f(w + t(x - w))$ holds for all $i \in [n]$, we can bound the first inequality as follows:

$$\begin{aligned} f(u) - f(z) &\geq -(1 - c_f) \int_0^1 \langle w - x, \nabla f(w + t(x - w)) \rangle dt \\ &= -(1 - c_f)(f(w) - f(x)). \end{aligned}$$

Equivalently, we can write:

$$(4.2) \quad f(u) + f(w) \geq f(z) + (1 - c_f)f(x) + c_f f(w).$$

Combining the inequalities 4.1 and 4.2, we conclude:

$$f(z) - (1 + c_f)f(x) + c_f f(w) \leq \langle \nabla f(x), z - x \rangle - \frac{\mu}{2}\|z - x\|^2.$$

Given that $c_f \in [0, 1]$ and f is non-negative, we can drop the term $c_f f(w)$ and derive the result as stated. \square

We can now exploit the result of Lemma 4.1 to obtain the approximation guarantee of the PGA algorithm.

THEOREM 4.1. *If f is a monotone, μ -strongly DR-submodular and L -smooth function, PGA obtains the following approximation guarantee:*

$$f(x) \geq \frac{1}{1 + c_f}f(x^*) - \frac{e^{-\mu K/L}}{1 + c_f}(f(x^*) - (1 + c_f)f(x_1)).$$

Moreover, for the DR-submodular setting ($\mu = 0$), the utility of the output of the PGA algorithm is bounded as follows:

$$f(x) \geq \frac{1}{1 + c_f}f(x^*) - \frac{L}{2K(1 + c_f)}\|x_1 - x^*\|^2.$$

Hassani et al. (2017) analyzed the PGA method in the DR-submodular setting ($\mu = 0$) and proved that $f(x) \geq \frac{1}{2}f(x^*) - \frac{LR^2}{2K}$. In comparison, thanks to Lemma 4.1, we obtain an improved $\frac{1}{1 + c_f}$ approximation ratio with a similar sub-linear convergence rate in the DR-submodular setting. Moreover, if f is μ -strongly DR-submodular, Theorem 4.1 shows that the $\frac{1}{1 + c_f}$ approximation ratio could be achieved at a faster linear convergence rate. Furthermore, if $c_f < \frac{1}{e-1} \approx 0.58$, the $\frac{1}{1 + c_f}$ approximation ratio obtained in Theorem 4.1 is

ALGORITHM 4.1. OGA

Input: Convex set \mathcal{K} , $x_1 \in \mathcal{K}$, L -smooth and μ -strongly DR-submodular $\{f_t\}_{t=1}^T$, $\{\eta_t\}_{t=1}^T$.

Output: $\{x_t\}_{t=1}^T$.

for $t = 1$ **to** T **do**

 Play x_t and receive the reward $f_t(x_t)$.

 Set $x_{t+1} = \text{Proj}_{\mathcal{K}}(x_t + \eta_t \nabla f_t(x_t))$.

end for

larger than the $1 - \frac{1}{e}$ approximation ratio guaranteed by the Frank-Wolfe variant of Bian et al. (2017). However, $\frac{1}{1+c_f} \leq 1 - \frac{c_f}{e}$ always holds, i.e., the approximation ratio of the SDRFW algorithm is greater than that of the PGA algorithm.

4.1 Online setting We can also use Lemma 4.1 to obtain improved regret bounds in the online setting for the Online Gradient Ascent (OGA) algorithm. To be precise, consider the following online optimization protocol: A convex constraint set \mathcal{K} with diameter R is given. At each iteration $t \in [T]$, the online algorithm first chooses an action $x_t \in \mathcal{K}$. Upon committing to this choice, a monotone (strongly) DR-submodular function $f_t : \mathcal{K} \rightarrow \mathbb{R}$, $\mathcal{K} \subset \mathbb{R}_+^n$, is revealed and the algorithm receives the reward $f_t(x_t)$. The goal is to maximize the total obtained reward or equivalently minimize the α -regret $\alpha\text{-}R_T = \alpha \max_{x \in \mathcal{K}} \sum_{t=1}^T f_t(x) - \sum_{t=1}^T f_t(x_t)$, i.e., the difference between the cumulative reward of the algorithm and the α approximation to that of the best fixed decision in hindsight where $\alpha \in (0, 1]$.

The Online Gradient Ascent (OGA) algorithm is provided in Algorithm 4.1. OGA is the online counterpart of the PGA algorithm for the offline setting. Starting from an arbitrary initial point $x_1 \in \mathcal{K}$, for all $t \in [T-1]$, OGA uses the update $x_{t+1} = \text{Proj}_{\mathcal{K}}(x_t + \eta_t \nabla f_t(x_t))$ to obtain the next iterate x_{t+1} , where $\eta_t > 0$ is the step size. Chen et al. (2018) analyzed the OGA algorithm in the DR-submodular setting ($\mu = 0$) and provided $\mathcal{O}(\sqrt{T})$ bounds for the $\frac{1}{2}$ -regret of the algorithm. Using Lemma 4.1, we can obtain improved $\mathcal{O}(\sqrt{T})$ and $\mathcal{O}(\ln T)$ ($\frac{1}{1+c}$)-regret bounds in the DR-submodular and strongly DR-submodular settings respectively where $c = \max_{t \in [T]} c_{f_t}$. This result is stated in the theorem below.

THEOREM 4.2. *Assume that the functions $\{f_t\}_{t=1}^T$ are all monotone, β -Lipschitz and μ -strongly DR-submodular. If for all $t \in [T]$, we set $\eta_t = \frac{1}{\mu t}$, the OGA algorithm has the following ($\frac{1}{1+c}$)-regret bound.*

$$\frac{1}{1+c} \sum_{t=1}^T f_t(x^*) - \sum_{t=1}^T f_t(x_t) \leq \frac{\beta^2}{2\mu(1+c)} (1 + \ln T),$$

where $x^* = \arg \max_{x \in \mathcal{K}} \sum_{t=1}^T f_t(x)$. Moreover, in the

DR-submodular setting ($\mu = 0$), if we set $\eta_t = \eta = \frac{R}{\beta\sqrt{T}} \forall t \in [T]$, the ($\frac{1}{1+c}$)-regret of the algorithm is bounded as follows:

$$\frac{1}{1+c} \sum_{t=1}^T f_t(x^*) - \sum_{t=1}^T f_t(x_t) \leq \frac{R\beta}{1+c} \sqrt{T}.$$

For the online monotone DR-submodular maximization problem with bounded curvature, the only prior study was done by Harvey et al. (2020) where the authors proposed an algorithm for the special setting where the DR-submodular functions $\{f_t\}_{t=1}^T$ are the multilinear extensions of corresponding submodular set functions $\{F_t\}_{t=1}^T$ and they showed that the algorithm obtains an $\mathcal{O}(\sqrt{T})$ ($1 - \frac{c}{e} - \epsilon$)-regret bound with $\frac{n^2}{e}$ projections per iteration. In comparison, while the approximation ratio in Theorem 4.2 is slightly worse, Algorithm 4.1 performs *only a single projection per step* (hence a significantly lower computational complexity) and its logarithmic regret bound is superior in the strongly DR-submodular setting.

5 Computing the smoothness parameter

Both the SDRFW and PGA algorithms require knowledge of the smoothness parameter L of the objective function to be implemented. In this section, we show how computing L of a twice differentiable μ -strongly DR-submodular objective function f ($\mu \geq 0$) could be formulated as a convex optimization problem that could be solved efficiently prior to running the algorithms. Given that our technique applies to the DR-submodular setting (i.e., $\mu = 0$) as well, the results of this section are useful for the proposed algorithms in prior works for DR-submodular maximization. In particular, while Hassani et al. (2017) chose the step size of (stochastic) PGA as a function of the smoothness parameter L to obtain the theoretical approximation guarantees in their work, they mentioned that estimating L is difficult in general and poses a challenge for implementation. Therefore, they suggested an alternative adaptive step size rule (as a function of the iteration number) with no theoretical performance guarantees in their experiments. This section precisely addresses the aforementioned challenge.

As we defined smoothness earlier in Section 2, we need to find a constant $L > 0$ such that for all $x \in \mathcal{K}$, the smallest eigenvalue of $\nabla^2 f(x)$ is lower bounded by $-L$, i.e., $\nabla^2 f(x) \succeq -LI$. This is equivalent to finding $L > 0$ such that the largest eigenvalue of $-\nabla^2 f$ is uniformly upper bounded by L everywhere. Given that f is μ -strongly DR-submodular, $-\nabla^2 f$ is an element-wise non-negative matrix. The Perron-Frobenius theorem (Horn and Johnson, 2012, Theorem 8.4.4) states that if $-\nabla^2 f$

is irreducible, i.e., the matrix $(I - \nabla^2 f)^{n-1}$ is element-wise positive, $-\nabla^2 f$ has a positive real eigenvalue λ_{pf} equal to its spectral radius (which is the largest magnitude of its eigenvalues) and an entry-wise positive eigenvector $v_{pf} \succ 0$ corresponding to λ_{pf} and therefore, we can set $L = \lambda_{pf}$. In order to check irreducibility of the symmetric matrix $-\nabla^2 f$, we can associate with it an undirected graph G with n vertices labeled $\{1, \dots, n\}$ where there is an edge between vertices i and j if $-\nabla_{ij}^2 f = -\nabla_{ji}^2 f \geq 0$. $-\nabla^2 f$ is irreducible if and only if its associated graph G is connected. According to a result in the theory of non-negative matrices, the Perron-Frobenius eigenvalue is the solution of the following optimization problem:

$$(5.1) \quad \begin{aligned} & \text{minimize} && \lambda \\ & \text{subject to} && \sum_{j=1}^n -\nabla_{ij}^2 f(x) v_j \leq \lambda v_i \quad \forall i \in [n], \end{aligned}$$

where the variables are $x \succ 0$, $v \succ 0$ and $\lambda > 0$. We show how the above problem could be transformed to a convex optimization problem in many cases.

A function $h : \mathbb{R}_{++}^n \rightarrow \mathbb{R}$ defined as $h(x) = cx_1^{a_1} \dots x_n^{a_n}$, where $c > 0$ and $a_i \in \mathbb{R} \forall i \in [n]$, is called a monomial. A sum of m monomials, i.e., a function of the form $h(x) = \sum_{s=1}^m c_s x_1^{a_{1s}} \dots x_n^{a_{ns}}$ where $c_s > 0 \forall s \in [m]$ and $a_{is} \in \mathbb{R} \forall i \in [n], s \in [m]$, is called a posynomial. Consider the following optimization problem with variable $x \in \mathbb{R}_{++}^n$:

$$\begin{aligned} & \text{minimize} && h_0(x) \\ & \text{subject to} && h_{1i}(x) \leq 1 \quad \forall i \in [n] \\ & && h_{2j}(x) = 1 \quad \forall j \in [p]. \end{aligned}$$

If $h_0, h_{11}, h_{12}, \dots, h_{1n}$ are posynomials and h_{21}, \dots, h_{2p} are monomials, the above problem is called a Geometric Program (GP). While GPs are not generally convex optimization problems, they can be transformed to convex problems (see Boyd and Vandenberghe (2004, Section 4.5.3) for details).

For all $i \in [n]$, we can rewrite the constraints of problem (5.1) in the following equivalent way:

$$(\lambda^{-1} v_i^{-1}) \sum_{j=1}^n -\nabla_{ij}^2 f(x) v_j \leq 1.$$

If the non-zero entries of $-\nabla^2 f(x)$ are posynomial functions of the variable x , the constraints of problem (5.1) are posynomial inequalities and therefore, the optimization problem for computing the smoothness parameter $L = \lambda_{pf}$ can be expressed as a GP.

As an example, consider the problem of computing the stability number of an undirected graph G with adjacency matrix A where $f(x) = x^T(-A - I)x + 21^T x$ (introduced in Section 1). Without loss of generality, assume that G is connected (otherwise, we can consider

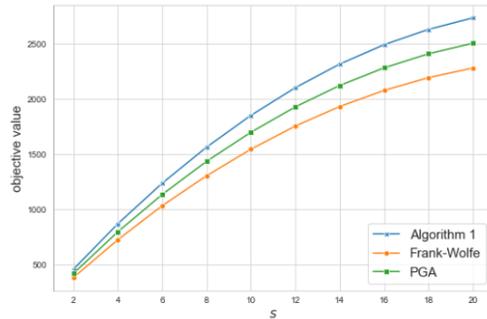


Figure 1: Comparison of algorithms for a class of strongly DR-submodular quadratic functions.

each connected component of G separately and use the fact that the stability number of a graph equals the sum of stability numbers of its connected components). Since G is connected, $-\nabla^2 f(x) = 2A + 2I$ is an entry-wise non-negative and irreducible matrix and therefore, computing the smoothness parameter L of f could be formulated as a GP.

More generally, consider the class of concave functions with negative dependence. Let $d \geq 2$. If $h_i : \mathbb{R}_+ \rightarrow \mathbb{R}$ is μ -strongly concave for all $i \in [n]$ and $\theta_{i_1, \dots, i_r} \leq 0$ for all $r \in [d]$ and $(i_1, \dots, i_r) \subseteq [n]$, the following function $f : \mathbb{R}_+^n \rightarrow \mathbb{R}$ is μ -strongly DR-submodular:

$$\begin{aligned} f(x) = & \sum_{i=1}^n h_i(x_i) + \sum_{(i,j):i \neq j} \theta_{ij} x_i x_j + \dots \\ & + \sum_{(i_1, \dots, i_d): i_r \neq i_s \forall r, s \in [d]} \theta_{i_1, \dots, i_d} x_{i_1} \dots x_{i_d}. \end{aligned}$$

It is easy to see that all off-diagonal entries of $-\nabla^2 f(x)$ are posynomials. If for all $i \in [n]$, $-h_i''(x_i)$ is a posynomial as well, the smoothness parameter of f could be computed as the solution of a GP.

6 Numerical examples

For the first experiment, we set $n = 25$ and chose $\mathcal{K} = \{x \in \mathbb{R}^n : 1^T x \leq s, 0 \preceq x \preceq 1\}$. We considered the class of indefinite quadratic functions and defined $f(x) = (\frac{1}{2}x - \mathbf{1})^T H x$, where $H \in \mathbb{R}^{n \times n}$ is a matrix

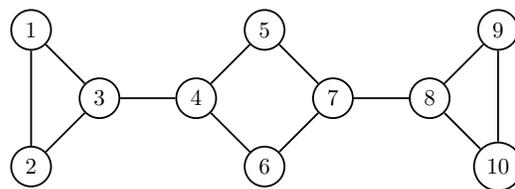


Figure 2: Graph G with $n = 10$ used in Experiment 2

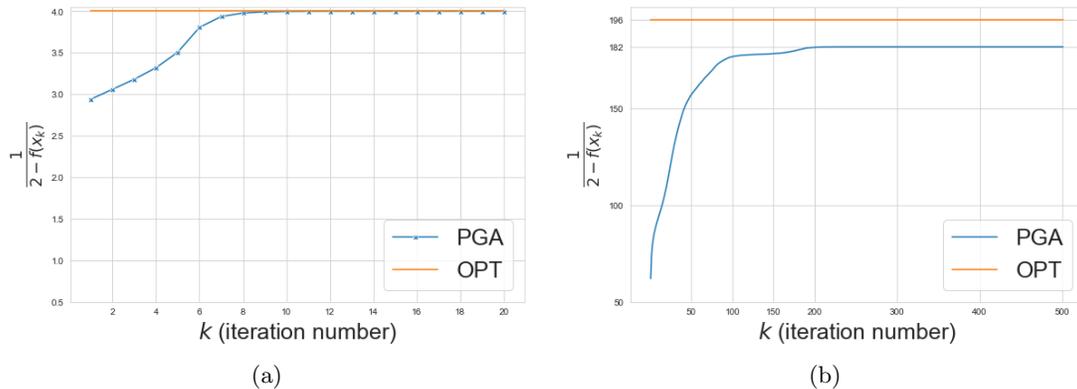


Figure 3: Computing the stability number $s(G)$ for a graph G with (a) $n = 10$ vertices and $s(G) = 4$, (b) $n = 1024$ vertices and $s(G) = 196$.

whose entries are uniformly distributed in the range $[-10, -5]$. Therefore, $\mu = 5$ in this setting. We computed L using the technique described in Section 5 and set $K = \lceil \frac{L}{\mu} \rceil$. We ran Algorithm 3.1, Frank-Wolfe variant of Bian et al. (2017) and PGA for K iterations using different values of s in the range $[2, 20]$ (using $x_1 = 0$ for PGA) and plotted the utility of the output of all three algorithms in each setting. The plot is depicted in Figure 1. This plot shows that Algorithm 3.1 obtains slightly higher utilities, followed by PGA and the Frank-Wolfe variant of Bian et al. (2017).

As the second experiment, we studied the problem of computing the stability number $s(G)$ for two graphs. The first graph is provided in Figure 2. It is easy to see that $s(G) = 4$ for this graph (e.g., vertices $\{3, 5, 6, 8\}$ form a maximum stable set of G). As it was mentioned earlier in Section 1, we set $\mathcal{K} = \{x \in \mathbb{R}^n : 1^T x = 1, x_i \geq 0 \forall i \in [n]\}$ where $n = 10$. Also, we defined $f(x) = x^T(-A - I)x + 21^T x$ and using the formula $s(G)^{-1} = 2 - \max_{x \in \mathcal{K}} f(x)$, we set $\frac{1}{2-f(x)}$ as the estimate of the stability number at $x \in \mathcal{K}$. Since all the diagonal entries of $\nabla^2 f(\cdot)$ are equal to -2 , we have $\mu = 2$. We also computed L using the method presented in Section 5. We ran PGA for this problem and plotted $\frac{1}{2-f(x_k)}$ versus the iteration number k in Figure 3(a). As the plot shows, PGA converges to the optimal value 4 after only 10 iterations. As the second example, we considered a graph with $n = 1024$ vertices from <https://oeis.org/A265032/a265032.html> that contains a collection of graph instances commonly used in coding theory. Using an algorithm with a running time of 200 hours, Niskanen and Östergård (2003) managed to show that the stability number of this graph is 196. The performance of PGA for this problem is plotted in Figure 3(b). The algorithm converges to the value 182 as the estimate of the stability number and

therefore, the performance of PGA for this problem is significantly better than the approximation guarantee proved in Theorem 4.1. Note that the domain in the second experiment does not contain the origin and thus, Algorithm 3.1 is not applicable to this problem.

7 Conclusion and future work

In this paper, we considered the class of monotone, smooth and strongly DR-submodular functions with bounded curvature and we proposed a number of first-order gradient methods for this problem along with their approximation guarantees and convergence rates.

This work could be extended in a number of directions. Throughout this paper, we assumed that we have access to the exact gradient of the objective function. However, in many applications, it is difficult to compute the gradient exactly, but an unbiased estimate of the gradient can be easily obtained. It is interesting to study stochastic gradient methods for this setting and provide performance guarantees. Secondly, we introduced the class of strongly DR-submodular functions with respect to the norm $\|\cdot\|_2$. As we showed in Lemma 2.1, strongly DR-submodular functions are strongly concave along non-negative directions with respect to $\|\cdot\|_2$. This definition could be easily extended to other norms as well. For instance, a twice differentiable function f is μ -strongly DR-submodular with respect to the norm $\|\cdot\|_1$ if *all* entries of the Hessian $\nabla^2 f(x)$ are upper bounded by $-\mu$ for all x . Similarly, while the smoothness of f over the domain \mathcal{K} was defined with respect to the Euclidean norm, in some cases, f and \mathcal{K} are not well-behaved in the $\|\cdot\|_2$ norm and L scales with the ambient dimension n leading to slower convergence rates for our proposed algorithms in large-scale applications. In such cases, one can study mirror ascent methods that are designed to adapt to smoothness in general norms.

References

- Andrew An Bian, Baharan Mirzasoleiman, Joachim Buhmann, and Andreas Krause. Guaranteed non-convex optimization: Submodular maximization over continuous domains. In *Artificial Intelligence and Statistics*, pages 111–120. PMLR, 2017.
- Yatao Bian, Joachim M Buhmann, and Andreas Krause. Continuous submodular function maximization. *arXiv preprint arXiv:2006.13474*, 2020.
- Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- Peter Brucker. An $O(n)$ algorithm for quadratic knapsack problems. *Operations Research Letters*, 3(3):163–166, 1984.
- Lin Chen, Hamed Hassani, and Amin Karbasi. Online continuous submodular maximization. In *International Conference on Artificial Intelligence and Statistics*, pages 1896–1905. PMLR, 2018.
- 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.
- Abhimanyu Das and David Kempe. Submodular meets spectral: Greedy algorithms for subset selection, sparse approximation and dictionary selection. *arXiv preprint arXiv:1102.3975*, 2011.
- Moran Feldman. Guess free maximization of submodular and linear sums. *Algorithmica*, 83(3):853–878, 2021.
- Nicholas Harvey, Christopher Liaw, and Tasuku Soma. Improved algorithms for online submodular maximization via first-order regret bounds. *Advances in Neural Information Processing Systems*, 33, 2020.
- Hamed Hassani, Mahdi Soltanolkotabi, and Amin Karbasi. Gradient methods for submodular maximization. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 5843–5853, 2017.
- Roger A Horn and Charles R Johnson. *Matrix analysis*. Cambridge university press, 2012.
- Shinji Ito and Ryohei Fujimaki. Large-scale price optimization via network flow. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 3855–3863. Curran Associates, Inc., 2016.
- David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146, 2003.
- Andreas Krause and Carlos Guestrin. Near-optimal observation selection using submodular functions. In *AAAI*, volume 7, pages 1650–1654, 2007.
- Hui Lin and Jeff Bilmes. A class of submodular functions for document summarization. In *Proceedings of the 49th annual meeting of the association for computational linguistics: human language technologies*, pages 510–520, 2011.
- Theodore S Motzkin and Ernst G Straus. Maxima for graphs and a new proof of a theorem of turán. *Canadian Journal of Mathematics*, 17:533–540, 1965.
- Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2003.
- Sampo Niskanen and Patric RJ Östergård. Cliquer user’s guide, version 1.0, communications laboratory, helsinki university of technology, espoo. In *Tech rep, Finland, Tech Rep*. 2003.
- Panos M Pardalos and Naina Koor. An algorithm for a singly constrained class of quadratic programs subject to upper and lower bounds. *Mathematical Programming*, 46(1):321–328, 1990.
- Omid Sadeghi and Maryam Fazel. Differentially private monotone submodular maximization under matroid and knapsack constraints. In *International Conference on Artificial Intelligence and Statistics*, pages 2908–2916. PMLR, 2021.
- Pier Giuseppe Sessa, Maryam Kamgarpour, and Andreas Krause. Bounding inefficiency of equilibria in continuous actions games using submodularity and curvature. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2017–2027. PMLR, 2019.
- Maxim Sviridenko, Jan Vondrák, and Justin Ward. Optimal approximation for submodular and supermodular optimization with bounded curvature. *Mathematics of Operations Research*, 42(4):1197–1218, 2017.