
On Maximization of Weakly Modular Functions: Guarantees of Multi-stage Algorithms, Tractability, and Hardness

Shinsaku Sakaue

NTT Communication Science Laboratories

Abstract

Maximization of *non-submodular* functions appears in various scenarios, and many previous works studied it based on some measures that quantify the closeness to being submodular. On the other hand, some practical non-submodular functions are actually close to being *modular*, which has been utilized in few studies. In this paper, we study cardinality-constrained maximization of *weakly modular* functions, whose closeness to being modular is measured by *submodularity* and *supermodularity ratios*, and reveal what we can and cannot do by using the weak modularity. We first show that guarantees of multi-stage algorithms can be proved with the weak modularity, which generalize and improve some existing results, and experiments confirm their effectiveness. We then show that weakly modular maximization is *fixed-parameter tractable* under certain conditions; as a byproduct, we provide a new time-accuracy trade-off for ℓ_0 -constrained minimization. We finally prove that, even if objective functions are weakly modular, no polynomial-time algorithms can improve the existing approximation guarantee achieved by the greedy algorithm in general.

1 INTRODUCTION

We consider the following set function maximization with a cardinality constraint:

$$\text{maximize } F(S) \quad \text{subject to } |S| \leq k, \quad (1)$$
$$S \subseteq [d]$$

where $d, k \in \mathbb{Z}_{>0}$, $[d] := \{1, \dots, d\}$, and $F : 2^{[d]} \rightarrow \mathbb{R}$. We assume F to be monotone, normalized, and

weakly modular (WM), where the closeness to being modular is represented with *submodularity ratio* (SBR) $\gamma \in [0, 1]$ and *supermodularity ratio* (SPR) $\beta \in [0, 1]$; (see, Section 1.2 for precise definitions). We say F is *weakly submodular* (*weakly supermodular*) if its SBR (SPR) is lower bounded. The larger SBR and SPR are, the closer F is to being submodular and supermodular, respectively, and F is modular if $\gamma = \beta = 1$.

Many previous studies on non-submodular maximization are based on some measures that quantify the deviation from being submodular (Elenberg et al., 2018; Qian and Singer, 2019), and SBR is one of the most prevalent among such measures. As regards weakly submodular maximization, Das and Kempe (2018) proved a well-known $(1 - e^{-\gamma})$ -approximation guarantee of the greedy algorithm (**Greedy**).

When it comes to practical non-submodular maximization instances, it can be effective to employ additional measures other than those quantifying the distance to being submodular. Bian et al. (2017) considered a class such that F has bounded SBR and *curvature* $\alpha \in [0, 1]$, and they proved a $\frac{1}{\alpha}(1 - e^{-\alpha\gamma})$ -approximation guarantee of **Greedy**. Namely, an improved approximation guarantee is possible if $\alpha < 1$. Unfortunately, however, $\alpha = 1$ occurs quite naturally in many applications as discussed in Section 2 (see also (Soma and Yoshida, 2018)), which motivates us to consider a wider class of non-submodular maximization that can capture the structures of various practical problems.

Weakly modular maximization (WMM) forms a wider class than that of (Bian et al., 2017). In fact, SPR β and curvature α always satisfy $\beta \geq 1 - \alpha$ (Bogunovic et al., 2018); i.e., SPR β can be bounded even if $\alpha = 1$. As shown in Section 2, various problems including feature selection (Das and Kempe, 2018) and production planning (Bian et al., 2017) strictly belong to WMM; that is, F has bounded SPR β even though $\alpha = 1$ in general. This fact suggests the importance of studying WMM. However, few previous works have studied problem (1) by utilizing the weak modularity, and so WMM remains to be studied.

1.1 Our Contribution

Our first contribution provides guarantees of efficient algorithms for WMM. As shown in Section 2, WMM can model various continuous optimization problems including ℓ_0 -constrained minimization and linear programming (LP) with a cardinality constraint. Given such WMM instances, the evaluation of objective functions involves solving optimization subproblems, which is often so costly that even standard **Greedy** becomes impractical. To overcome this hardship, we consider using *multi-stage* algorithms for WMM.

Guarantees of Multi-stage Algorithms In Section 3, we show that the multi-stage approach is effective for WMM instances; with this approach, we accelerate greedy-style algorithms by adding multiple elements in each iteration, instead of a single element. The only existing study that proved guarantees of multi-stage algorithms is (Wei et al., 2014); their result requires the submodularity, and the approximation ratio is expressed as $\frac{1}{\alpha}(1 - e^{-\alpha(1-\alpha)})$ in general, which becomes 0 if curvature α is equal to 1. Our guarantee of the multi-stage greedy algorithm (**Multi-Greedy**) for WMM is advantageous relative to the previous result in two aspects: it can be applied to WM functions, which are generally non-submodular, and it can yield positive approximation ratios even if $\alpha = 1$ as long as SBR and SPR are bounded. Our result also includes the $(1 - e^{-\gamma})$ -approximation guarantee of (Das and Kempe, 2018) as a special case. We then focus on ℓ_0 -constrained minimization and prove a guarantee of the multi-stage orthogonal matching pursuit (**Multi-OMP**), which can achieve a better approximation ratio than **Multi-Greedy**. Surprisingly, our result matches that of standard **OMP** (Elenberg et al., 2018), while **Multi-OMP** can run faster than **OMP**. Our result also improves that of the latest feature selection algorithm (Qian and Singer, 2019). Experiments show that the multi-stage approach successfully accelerates **Greedy** and **OMP** at the cost of a slight decline in solution quality.

Our second and their contributions, presented in Section 4, are related to theoretical properties of WMM. These contributions are important for revealing what we can and cannot do with the weak modularity.

Fixed-parameter Tractability In Section 4.1, we show that ϵ -error solutions for WMM can be obtained with a randomized *fixed-parameter tractable* (FPT) algorithm, whose computation cost depends arbitrarily on certain inputs including SBR γ , SPR β , sparsity k , and ϵ , but it is polynomial in d . The FPT algorithm was first proposed by Skowron (2017), but its guarantee was proved only for a special case of monotone submodular maximization. As a byproduct, we provide

a time-accuracy trade-off for ℓ_0 -constrained minimization, which is contrasted with the existing sparsity-accuracy trade-off (Shalev-Shwartz et al., 2010).

Hardness of Improving Approximation Ratio

As mentioned before, if curvature α is bounded by a constant smaller than 1, the $\frac{1}{\alpha}(1 - e^{-\alpha\gamma})$ -approximation guarantee of (Bian et al., 2017) improves the approximation ratio, $1 - e^{-\gamma}$, of (Das and Kempe, 2018). When it comes to WMM, not curvature α but SPR β ($\geq 1 - \alpha$) is bounded. Given this background, the following question arises: Can we improve the approximation ratio, $1 - e^{-\gamma}$, if SPR β is bounded by a constant, instead of curvature α . In Section 4.2, we show that it is generally impossible. More precisely, we prove that, even if $\gamma = 1$ and $\beta \geq 1/2$, no polynomial-time algorithms can improve the $(1 - e^{-1})$ -approximation guarantee in general in the value oracle model; i.e., bounded SPR β does not always help us to improve the ratio, $1 - e^{-\gamma}$. This result clarifies the theoretical gap between SPR β and curvature α .

1.2 Notation and Definitions

Given any $F : 2^{[d]} \rightarrow \mathbb{R}$, we define $F(\mathsf{T} \mid \mathsf{S}) := F(\mathsf{S} \cup \mathsf{T}) - F(\mathsf{S})$ for any $\mathsf{S}, \mathsf{T} \subseteq [d]$. All the set functions considered in this paper are monotone ($F(\mathsf{T} \mid \mathsf{S}) \geq 0$, $\forall \mathsf{S}, \mathsf{T} \subseteq [d]$) and normalized ($F(\emptyset) = 0$). We say F is submodular (supermodular) if $F(j \mid \mathsf{S}) \geq F(j \mid \mathsf{T})$ ($F(j \mid \mathsf{S}) \leq F(j \mid \mathsf{T})$) holds for any $\mathsf{S} \subseteq \mathsf{T}$ and $j \notin \mathsf{T}$. We assume that F can be evaluated in polynomial time w.r.t. d (or $\text{poly}(d)$ time). Given any $\mathsf{S} \subseteq [d]$ and $\mathbf{x} \in \mathbb{R}^{[d]}$, whose j -th entry \mathbf{x}_j is associated with $j \in [d]$, $\mathbf{x}_{\mathsf{S}} \in \mathbb{R}^{\mathsf{S}}$ denotes the restriction of \mathbf{x} to S . We define the support of \mathbf{x} as $\text{supp}(\mathbf{x}) := \{j \in [d] \mid \mathbf{x}_j \neq 0\}$.

SBR and SPR Given any monotone $F : 2^{[d]} \rightarrow \mathbb{R}$, $\mathsf{U} \subseteq [d]$, and $s \in \mathbb{Z}_{>0}$, we define SBR $\gamma_{\mathsf{U},s}$ and SPR $\beta_{\mathsf{U},s}$ as the largest scalars that satisfy

$$\gamma_{\mathsf{U},s} F(\mathsf{S} \mid \mathsf{L}) \leq \sum_{j \in \mathsf{S}} F(j \mid \mathsf{L}) \quad \text{and} \\ \sum_{j \in \mathsf{S}} F(j \mid \mathsf{L}) \leq \beta_{\mathsf{U},s}^{-1} F(\mathsf{S} \mid \mathsf{L}),$$

respectively, for any disjoint $\mathsf{L}, \mathsf{S} \subseteq [d]$ such that $\mathsf{L} \subseteq \mathsf{U}$ and $|\mathsf{S}| \leq s$. We say F is $(\gamma_{\mathsf{U}_1, s_1}, \beta_{\mathsf{U}_2, s_2})$ -WM if F has bounded $\gamma_{\mathsf{U}_1, s_1}$ and $\beta_{\mathsf{U}_2, s_2}$. Note that $\gamma_{\mathsf{U}', s'} \geq \gamma_{\mathsf{U}, s}$ and $\beta_{\mathsf{U}', s'} \geq \beta_{\mathsf{U}, s}$ hold for any $\mathsf{U}' \subseteq \mathsf{U}$ and $s' \leq s$. We can confirm that $\gamma_{\mathsf{U}, s} \in [0, 1]$ and $\beta_{\mathsf{U}, s} \in [1/s, 1]$ hold for any U and s . We define $\gamma_{s', s} := \min_{|\mathsf{U}| \leq s'} \gamma_{\mathsf{U}, s}$ and $\beta_{s', s} := \min_{|\mathsf{U}| \leq s'} \beta_{\mathsf{U}, s}$; we sometimes use $\gamma_s := \gamma_{s, s}$ and $\beta_s := \beta_{s, s}$. We have $\gamma_d = 1$ ($\beta_d = 1$) iff F is submodular (supermodular).

Curvature Given monotone $F : 2^{[d]} \rightarrow \mathbb{R}$, its curvature $\alpha \in [0, 1]$ is defined as the smallest scalar that

satisfies

$$F(j \mid S \setminus \{j\} \cup M) \geq (1 - \alpha)F(j \mid S \setminus \{j\})$$

for any $S, M \subseteq [d]$ and $j \in S \setminus M$. We have $\beta_{U,s} \geq 1 - \alpha$ for any U and s (see, (Bogunovic et al., 2018)).

Restricted Strong Convexity and Restricted Smoothness When studying ℓ_0 -constrained minimization algorithms, the restricted strong convexity (RSC) and restricted smoothness (RSM) of loss function $l : \mathbb{R}^d \rightarrow \mathbb{R}$ is often used (Jain et al., 2014; Elenberg et al., 2018; Yuan et al., 2018). We assume l to be differentiable. Given any fixed $s_1, s_2 \in \mathbb{Z}_{>0}$, we say l is μ_{s_1, s_2} -RSC and ν_{s_1, s_2} -RSM if it satisfies

$$\begin{aligned} l(\mathbf{y}) &\geq l(\mathbf{x}) + \langle \nabla l(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{\mu_{s_1, s_2}}{2} \|\mathbf{y} - \mathbf{x}\|_2^2 \quad \text{and} \\ l(\mathbf{y}) &\leq l(\mathbf{x}) + \langle \nabla l(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{\nu_{s_1, s_2}}{2} \|\mathbf{y} - \mathbf{x}\|_2^2, \end{aligned}$$

respectively, for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ such that $\|\mathbf{x}\|_0 \leq s_1$, $\|\mathbf{y}\|_0 \leq s_1$, and $\|\mathbf{x} - \mathbf{y}\|_0 \leq s_2$. If l is quadratic, the above inequalities reduce to those of the well-known restricted isometric property (RIP) condition (Candès et al., 2006). We let $\mu_s := \mu_{s,s}$ and $\nu_s := \nu_{s,s}$. We define the restricted condition number as $\kappa_s := \nu_s / \mu_s$. Typically, l with a smaller κ_s value is easier to deal with. If l is μ_d -RSC and ν_d -RSM, we abbreviate the subscript and say l is μ -strongly convex (μ -SC) and ν -smooth (ν -SM); we call $\kappa := \nu / \mu$ a condition number.

1.3 Related Work

For the case where F is submodular, Nemhauser et al. (1978) proved the $(1 - e^{-1})$ -approximation guarantee of **Greedy**. Nemhauser and Wolsey (1978) proved that no polynomial-time algorithms can improve this guarantee in the value oracle model, and Feige (1998) proved the NP-hardness for the case of *Max k-cover*. As regards tractability, Skowron (2017) developed a randomized FPT approximation algorithm for maximization of monotone submodular functions with a special property called *p-separability*. Unlike our results, those results hold only for monotone submodular maximization.

When it comes to non-submodular maximization, various notions have been introduced to obtain theoretical guarantees (Krause and Cevher, 2010; Feige and Izsak, 2013; Horel and Singer, 2016; Wang et al., 2016; Zhou and Spanos, 2016). Das and Kempe (2018) proposed SBR, one of the most prevalent notion used in many studies (Hu et al., 2016; Elenberg et al., 2017; Khanna et al., 2017a,b; Chen et al., 2018; Qian and Singer, 2019), and they proved that **Greedy** outputs solution $S \subseteq [d]$ with a $(1 - e^{-\gamma_{S,k}})$ -approximation guarantee. Harshaw et al. (2019) proved that no polynomial-time algorithms can improve the $(1 - e^{-\gamma^d})$ -approximation

guarantee in general for every $\gamma_d \in (0, 1]$ value. This result is different from our hardness result since they do not assume SPR to be bounded by a constant; this difference is critical since bounded SPR could make the problem easier.

The definition of SPR that we use was introduced by Bogunovic et al. (2018). Other SPR-like notions have been used in the context of minimization problems (Takeda et al., 2013; Liberty and Sviridenko, 2017), but those are different from SPR, which quantifies the deviation from being supermodular in the context of maximization problems.

Curvature α (Conforti and Cornuéjols, 1984; Bian et al., 2017) is also used in many studies (Iyer et al., 2013; Sviridenko et al., 2015; Bai and Bilmes, 2018). Its value is, however, often pessimistic (i.e., $\alpha \approx 1$) as pointed out by Soma and Yoshida (2018), and to bound the curvature value is more demanding than to bound SPR. Hence our results obtained with SPR are different from existing guarantees that use curvature; although those results can sometimes be improved by using *greedy curvature* $\alpha_G \leq \alpha$ (Bian et al., 2017), no lower bounds of α_G for WM functions have been proved.

We remark that our work is different from some previous studies on set functions that are close to being modular. As mentioned before, Bian et al. (2017) studied the case where the curvature and SBR are bounded, and they proved that **Greedy** finds solution S with a $\frac{1}{\alpha}(1 - e^{-\alpha\gamma_{S,k}})$ -approximation guarantee. They also proved that **Greedy** cannot improve this guarantee. Unlike this result, our hardness result considers every polynomial-time algorithm. Bogunovic et al. (2018) considered the case where SBR and SPR are bounded. However, they are interested in obtaining guarantees for robust maximization, not for the standard cardinality-constrained maximization (1), which is of our interest. Chierichetti et al. (2015) defined the approximate modularity as the ℓ_∞ -distance to being modular, which is different from the weak modularity.

Wei et al. (2014) provided the curvature-dependent approximation guarantees of multi-stage algorithms for submodular maximization. Marsousi et al. (2013) applied **Multi-OMP** to a special case of ℓ_0 -constrained minimization where the loss function l is quadratic, but its theoretical guarantee has not been proved.

The idea of adding multiple elements in each round is also considered in the context of parallel algorithms (Balkanski and Singer, 2018). Qian and Singer (2019) recently developed a parallel approximation algorithm that runs in $O(\ln d)$ rounds for ℓ_0 -constrained minimization. Surprisingly, thanks to the use of the weak modularity, we can show that **Multi-OMP** with only one round achieves a better approximation ratio.

2 APPLICATIONS

We motivate to study WMM by presenting its applications. For each application, we present lower bounds of SBR and SPR. We also provide an example such that $\alpha = 1$ for one of the applications, and such examples for the other applications are presented in Appendix A.

ℓ_0 -constrained Minimization Given a differentiable loss function $l : \mathbb{R}^d \rightarrow \mathbb{R}$, we consider the ℓ_0 -constrained minimization problem: $\min_{\|\mathbf{x}\|_0 \leq k} l(\mathbf{x})$. It is generally NP-hard (Natarajan, 1995) and appears in many practical scenarios: feature selection (Das and Kempe, 2018) and M-estimation (Jain et al., 2014). The problem can be rewritten as in (1) with $F(\mathbf{S}) = l(0) - \min_{\text{supp}(\mathbf{x}) \subseteq \mathbf{S}} l(\mathbf{x})$, which has SBR $\gamma_{\mathbf{U},s} \geq \mu_{|\mathbf{U}|+s}/\nu_{|\mathbf{U}|+1,1} \geq 1/\kappa_{|\mathbf{U}|+s}$ (Elenberg et al., 2018) and SPR $\beta_{\mathbf{U},s} \geq \mu_{|\mathbf{U}|+1}/\nu_{|\mathbf{U}|+s,s} \geq 1/\kappa_{|\mathbf{U}|+s}$ (Appendix A.1); the later bound improves an existing result, $\beta_{\mathbf{U},s} \geq \mu/\nu = 1/\kappa$, of (Bogunovic et al., 2018). The evaluation of $F(\mathbf{S})$ involves solving $\min_{\text{supp}(\mathbf{x}) \subseteq \mathbf{S}} l(\mathbf{x})$. If l is quadratic, we can solve it by computing a pseudo-inverse matrix. Given a more general l , we can use iterative methods (e.g., (Shalev-Shwartz and Zhang, 2016)) to solve the minimization problem.

LP with a Cardinality Constraint We consider the following constrained LP that models optimal production planning problem (Bian et al., 2017). Given a set of d items and k production lines, we design a production plan so that the total profit is maximized; i.e., we aim to solve $\max_{\mathbf{x} \in \mathcal{P}, \|\mathbf{x}\|_0 \leq k} \mathbf{c}^\top \mathbf{x}$, where $\mathbf{c} \in \mathbb{R}^d$ and $\mathcal{P} \subseteq \mathbb{R}^d$ represent the profit of each item and a polytope specified by continuous constraints (e.g., upper bounds on the total quantities of materials), respectively. This problem can be reformulated as in (1) with $F(\mathbf{S}) := \max_{\mathbf{x} \in \mathcal{P}} \mathbf{c}_S^\top \mathbf{x}_S$. As in (Bian et al., 2017), SBR $\gamma_{\mathbf{U},s}$ of F is lower bounded by some $\gamma_0 > 0$ for any \mathbf{U} and s under the *non-degeneracy* assumption. Furthermore, from the definition of SPR, we have $\beta_{\mathbf{U},s} \geq 1/s$; although the lower bound, $1/s$, can be small if $s \approx d$, this is not always the case. For example, in the guarantee of **Multi-Greedy** (Theorem 1), s is a controllable parameter, b_{\max} ; i.e., $\beta_{\mathbf{U},s} \geq 1/b_{\max}$ holds.

Coverage Maximization Submodular functions sometimes have bounded SPR $\beta_{\mathbf{U},s}$, and such functions can be seen as special WM functions such that $\gamma_{\mathbf{U},s} = 1$ for any \mathbf{U} and s . One such example is the coverage function. Let V be a finite set and $w_v \geq 0$ ($v \in V$). We define d groups $I_1, \dots, I_d \subseteq V$, and we let $I_S := \bigcup_{j \in S} I_j$ for any $S \subseteq [d]$. The coverage function is defined as $F(\mathbf{S}) := \sum_{v \in I_S} w_v$, which is submodular and used in many scenarios including document summarization (Lin and Bilmes, 2011) and itemset mining (Kumar

et al., 2015). Given $s \in \mathbb{Z}_{>0}$, we assume that any collection of up to s groups covers every $v \in V$ at most c_s times; i.e., $c_s := \max_{v \in V, |S| \leq s} |\{j \in S \mid v \in I_j\}|$. Note that $c_s \leq s$ always holds. In this case, SPR $\beta_{\mathbf{U},s}$ of F is lower bounded by $1/c_s$ as proved in Appendix A.3.

Example with Unbounded Curvature We provide an example of LP with a cardinality constraint such that $\alpha = 1$. Let $d = 2$, $\mathbf{x} = (x_1, x_2)^\top$, and $\epsilon \in (0, 1)$. We consider a set function defined as $F(\mathbf{S}) = \max_{\text{supp}(\mathbf{x}) \subseteq \mathbf{S}} \{x_1 + \epsilon x_2 \mid x_1 + x_2 \leq 1, x_1 \geq 0, x_2 \geq 0\}$ for any $\mathbf{S} \subseteq [d]$; i.e., $\mathbf{c} = (1, \epsilon)^\top$ and $\mathcal{P} = \{\mathbf{x} \in \mathbb{R}^2 \mid x_1 + x_2 \leq 1, x_1 \geq 0, x_2 \geq 0\}$. From the definitions of SBR and SPR, we can confirm that $\gamma_{\mathbf{U},s} = 1$ and $\beta_{\mathbf{U},s} \geq \frac{1}{1+\epsilon}$ hold for any \mathbf{U} and s . On the other hand, we have $\alpha = 1$ since $F(\{2\} \mid \{1\}) \geq (1 - \alpha)F(\{1\})$ must hold for $F(\{2\} \mid \{1\}) = 0$ and $F(\{1\}) = 1$.

3 MULTI-STAGE ALGORITHMS

We study multi-stage algorithms for WMM. Let \mathbf{S}^* and \mathbf{x}^* be target solutions for WMM and ℓ_0 -constrained minimization, respectively, and $k^* := |\mathbf{S}^*| = \|\mathbf{x}^*\|_0$; note that we allow k^* to be different from k . As a warm-up, we first discuss two simple algorithms:

Single-stage Algorithm We compute $F(j)$ for $j \in [d]$ and let $\mathbf{S} = \text{argmax}_{\mathbf{S}' : |\mathbf{S}'| \leq k} \sum_{j \in \mathbf{S}'} F(j)$. The algorithm requires to evaluate F only d times, and it can find optimal solutions if F is modular. However, its approximation ratio becomes poor if F lacks the modularity. We consider a coverage maximization instance with $d = 2k$ and $V = \{v_1, \dots, v_{2k}\}$. Let $w_{v_j} = 1$ and $I_j = \{v_j\}$ for $j = 1, \dots, k$, and let $w_{v_j} = \epsilon \ll 1$ and $I_j = \{v_1, v_j\}$ for $j = k+1, \dots, 2k$. In this case, if $\epsilon > 0$ is sufficiently small, the approximation ratio achieved by the single-stage algorithm is $\frac{1+k\epsilon}{k+\epsilon} = O(1/d)$.

Greedy Algorithm Starting from $\mathbf{S} = \emptyset$, **Greedy** iteratively adds $\text{argmax}_{j \notin \mathbf{S}} F(j \mid \mathbf{S})$ to \mathbf{S} and outputs \mathbf{S} after k iterations. Given F with SBR $\gamma_{\mathbf{S},k^*}$, **Greedy** achieves a $(1 - \exp(-\gamma_{\mathbf{S},k^*}))$ -approximation guarantee. **Greedy** is, however, often costly due to the sequential evaluation of F , particularly when the evaluation of F involves solving optimization problems. For example, in the case of ℓ_0 -constrained minimization, **Greedy** solves convex minimization problems $\Theta(dk)$ times.

Namely, while the single-stage algorithm can efficiently find optimal solutions if F is modular, **Greedy** can achieve better guarantees for non-modular F at the cost of more computational effort. In the case of WMM, since F is close to being modular, we can expect that an intermediate of the above two algorithms works well. The multi-stage approach provides such an intermediate. As in Algorithm 1, we perform m ($\leq k$) iterations

Algorithm 1 Multi-stage algorithm

```

1:  $U \leftarrow [d], S \leftarrow \emptyset$ 
2: for  $i = 1, \dots, m$  do
3:    $B_i \leftarrow \operatorname{argmax}_{B \subseteq U: |B| \leq b_i} G_S(B)$ 
4:    $S \leftarrow S \cup B_i$ 
5:    $U \leftarrow U \setminus B_i$ 
6: return  $S$ 
    
```

to obtain a solution. In each i -th iteration, we choose subset $B_i \subseteq [d]$ of size at most b_i so that it maximizes a *surrogate function*, G_S , where S is the current solution. To obtain fast multi-stage algorithms, G_S should be efficiently evaluated and maximized. Below we design G_S for **Multi-Greedy** and **Multi-OMP**, and we present their theoretical guarantees. We then experimentally evaluate the multi-stage algorithms.

3.1 Theoretical Guarantees

Let $S_i = B_1 \cup \dots \cup B_i$ for $i \in [m]$ and $S_0 = \emptyset$. We first present a guarantee of **Multi-Greedy** for WMM. We then focus on ℓ_0 -constrained minimization and prove a guarantee of **Multi-OMP**. As detailed below, our results generalize and improve some existing results, which emphasizes that to utilize the weak modularity is effective for obtaining strong theoretical results. The proofs of the theorems are presented in Appendix B.1.

3.1.1 Multi-Greedy

Multi-Greedy uses $G_S(B) = \sum_{j \in B} F(j | S)$ as a surrogate function. Therefore, **Multi-Greedy** evaluates F $\Theta(dm)$ times. We can show that **Multi-Greedy** enjoys the following approximation guarantee:

Theorem 1. *Let b_{\max} be an integer satisfying $1 \leq b_{\max} \leq k^*$. Set b_1, \dots, b_m so as to satisfy $b_i \in [b_{\max}]$ for $i \in [m]$ and $\sum_{i \in [m]} b_i = k$. If S is the solution obtained with **Multi-Greedy** and F is $(\gamma_{S, k^*}, \beta_{S, b_{\max}})$ -WM, we have*

$$\begin{aligned} F(S) &\geq \left(1 - \prod_{i=1}^m \left(1 - \gamma_{S_{i-1}, k^*} \beta_{S_{i-1}, b_i} \frac{b_i}{k^*} \right) \right) F(S^*) \\ &\geq \left(1 - \exp \left(-\gamma_{S, k^*} \beta_{S, b_{\max}} \frac{k}{k^*} \right) \right) F(S^*). \end{aligned}$$

Note that, if we set $b_{\max} = 1$, this result recovers the $(1 - e^{-\gamma_{S, k}})$ -approximation of **Greedy** (Das and Kempe, 2018) since $\beta_{S, 1} = 1$. **Multi-Greedy** with $m = 1$ is the single-stage algorithm, which is studied as the oblivious algorithm in the field of ℓ_0 -constrained minimization. Elenberg et al. (2018) proved that its approximation ratio is at least $\max\{\frac{1}{k} \kappa_k^{-1}, \frac{3}{4} \kappa_k^{-2}, \kappa_k^{-3}\}$. Note that Theorem 1 improves this result since, if $b_1 = k = k^*$, the approximation ratio becomes $\max\{\frac{1}{k} \kappa_k^{-1}, \kappa_k^{-2}\}$

from the lower bounds of SBR and SPR (Section 2) and $\beta_{\emptyset, k} \geq 1/k$. More generally, for $m \geq 1$, **Multi-Greedy** achieves a $1 - \exp(\kappa_{2k}^{-1} \kappa_{k+b_{\max}}^{-1})$ -approximation. Below we show that a stronger guarantee for ℓ_0 -constrained minimization can be obtained by using **Multi-OMP**.

3.1.2 Multi-OMP

We then focus on ℓ_0 -constrained minimization; i.e., we assume $F(S) = l(0) - \min_{\operatorname{supp}(\mathbf{x}) \subseteq S} l(\mathbf{x})$ ($\forall S \subseteq [d]$). We let $\mathbf{b}^{(S)} := \operatorname{argmin}_{\operatorname{supp}(\mathbf{x}') \subseteq S} l(\mathbf{x}')$ for any $S \subseteq [d]$. **Multi-OMP** uses $G_S(B) = \sum_{j \in B} |\nabla l(\mathbf{b}^{(S)})_j|^2$ as a surrogate function. Thus, it requires to compute the gradient and to solve convex minimization problems m times. To prove the guarantee of **Multi-OMP**, we use the following lemma, which, roughly speaking, connects the decrease in l to the increase in F .

Lemma 1. *For any disjoint $A, B \subseteq [d]$, if $l(\cdot)$ is $\mu_{|A \cup B|}$ -RSC and $\nu_{|B|, |B \setminus A|}$ -RSM, we have*

$$\frac{\|\nabla l(\mathbf{b}^{(A)})_B\|_2^2}{2\nu_{|B|, |B \setminus A|}} \leq F(B | A) \leq \frac{\|\nabla l(\mathbf{b}^{(A)})_B\|_2^2}{2\mu_{|A \cup B|}}.$$

A special case of the lemma is implicitly used in (Elenberg et al., 2018). In Appendix A.1, we provide a slightly stronger version of the lemma, which we use for proving the guarantee of **Multi-OMP**. By using the lemma, we can employ the technique used when proving Theorem 1, which leads to the following result:

Theorem 2. *Set b_1, \dots, b_m as in Theorem 1. If l is μ_{k+k^*} -RSC and $\nu_{k, b_{\max}}$ -RSM, then **Multi-OMP** outputs solution S such that $\mathbf{x} = \operatorname{argmin}_{\operatorname{supp}(\mathbf{x}') \subseteq S} l(\mathbf{x}')$ satisfies*

$$\begin{aligned} l(\mathbf{x}) &\leq l(\mathbf{x}^*) + \prod_{i=1}^m \left(1 - \frac{\mu_{|S_{i-1} \cup S^*|} b_i}{\nu_{|S_i|, |B_i|} k^*} \right) (l(0) - l(\mathbf{x}^*)) \\ &\leq l(\mathbf{x}^*) + \exp \left(-\frac{\mu_{k+k^*} k}{\nu_{k, b_{\max}} k^*} \right) (l(0) - l(\mathbf{x}^*)) \\ &\leq l(\mathbf{x}^*) + \exp \left(-\frac{1}{\kappa_{k+k^*}} \frac{k}{k^*} \right) (l(0) - l(\mathbf{x}^*)). \end{aligned}$$

Note that, if $k = k^*$, Theorem 2 gives a $(1 - \exp(\kappa_{2k}^{-1}))$ -approximation guarantee, which improves the aforementioned guarantee of **Multi-Greedy**. Interestingly, the approximation ratio matches those of **OMP** and **Greedy** (Elenberg et al., 2018). Namely, the use of the multi-stage approach does not degrade the theoretical guarantee. If we let $b_1 = k = k^*$, **Multi-OMP** with only one round achieves a κ_{2k}^{-1} -approximation guarantee; this improves the existing $(1 - \exp(-\kappa_{2k}^{-4}))$ -approximation guarantee with $O(\ln d)$ rounds, recently proved by Qian and Singer (2019), in terms of both the approximation ratio and the computation complexity.




Figure 1: Results of ℓ_0 -constrained minimization with various m values. Top (bottom) figures present those of regression (classification) instances. Running times and loss function values are shown with semi-log plots. Each curve and error band indicate the average and standard deviation, respectively, calculated over 100 instances.

3.2 Experiments

We evaluate the multi-stage algorithms via experiments with two kinds instances: ℓ_0 -constrained minimization and LP with a cardinality constraint. We use Python3 to implement the algorithms, and we conduct experiments on a 64-bit macOS machine with 3.3GHz Intel Core i7 CPUs and 16 GB RAM. All the algorithms considered below can be accelerated via randomization (Li et al., 2016; Khanna et al., 2017b), but to simplify the comparisons we here do not employ such techniques.

3.2.1 ℓ_0 -constrained Minimization

We use two instances with the real-world dataset available at PMLB (Olson et al., 2017). The first is a sparse regression instance with the square loss, $l(\mathbf{x}) = \frac{1}{2n} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2$, where $\mathbf{A} \in \mathbb{R}^{n \times d}$ and $\mathbf{y} \in \mathbb{R}^n$ are obtained from “satellite_image” dataset. We use the 1st and 2nd order polynomial features; as a result, we have $d = 666$ features and a sample of size $N = 6435$. We set $k = 100$. The second is a sparse classification instance. We use the regularized logistic loss, $l(\mathbf{x}) = \frac{1}{n} \sum_{i \in [n]} \ln(1 + \exp(-\mathbf{y}_i(\mathbf{A}\mathbf{x})_i)) + \frac{\lambda}{2} \|\mathbf{x}\|_2^2$, where $\mathbf{A} \in \mathbb{R}^{n \times d}$ and $\mathbf{y} \in \mathbb{R}^n$ are obtained from “hill_valley_with_noise” dataset. The dataset has $d = 100$ features and a sample of size $N = 1212$. We let $\lambda = 0.01$ and $k = 20$. For each instance, we randomly split the sample into training and test data of sizes $\lfloor N/2 \rfloor$ and $\lfloor N/2 \rfloor$, respectively; we thus create 100 random instances. We consider multi-stage algorithms with various numbers of iterations, $m = k, 0.9k, \dots, 0.1k$ ($m = k$ corresponds to standard Greedy/OMP). We set $b_1, \dots, b_{k-m \lfloor k/m \rfloor}$ at $\lfloor k/m \rfloor$ and the rest at $\lfloor k/m \rfloor$. We use two baselines based on the projected gradient method: iterative hard thresholding (IHT) (Jain et al., 2014) and hard thresholding pursuit

(HTP) (Yuan et al., 2018). We continue their iterations until the decrease in $l(\cdot)$ value becomes smaller than 10^{-5} . We evaluate the algorithms with running times, loss function values, R^2 scores (for regression), and mean accuracy (for classification); the last two are defined by the corresponding scikit-learn score functions.

Figure 1 summarizes the results. We see that the multi-stage algorithms speed up as m decreases; in particular, Multi-OMP becomes as fast as HTP. In the regression instances, multi-stage algorithms achieve better loss function values and R^2 scores than the baselines. Other than for Multi-OMP with $m = 10$, the decrease in m has negligible effects on loss function values and R^2 scores. In the classification instances, the loss function values of the multi-stage algorithms increase as m decreases, but they are smaller on average than those of IHT and HTP. The multi-stage algorithms also achieve better mean accuracy than the baselines. To conclude, by using the multi-stage approach, Greedy and OMP can become faster while outperforming the baselines. When addressing large-scale instances in practice, it would be effective to try multi-stage algorithms with a small m and increase it until an acceptable solution is obtained.

As regards solution quality, the gap between the greedy-style algorithms (Multi-Greedy and Multi-OMP) and the baselines (IHT and HTP) can partially be explained in terms of the restricted condition number. For example, IHT requires $k \geq \Omega(\kappa_{2k+k^*}^2 \ln \epsilon^{-1})$ to achieve ϵ -errors (Jain et al., 2014), while Multi-OMP requires $k \geq \Omega(\kappa_{k+k^*} \ln \epsilon^{-1})$ as implied in Theorem 2. This suggests that greedy-style algorithms can be more resistant to being ill-conditioned (or a large restricted condition number), which is often the case with real-world instances; hence the better performance of the greedy-style algorithms. Appendix B.2 presents further experiments with well- and ill-conditioned instances.




Figure 2: Results of LP with a cardinality constraint. Each curve (error band) indicates the average (standard deviation) calculated over 100 instances.

3.3 LP with a Cardinality Constraint

We use synthetic optimal production planning instances. We let $\mathcal{P} = \{\mathbf{x} \in \mathbb{R}^d \mid \mathbf{A}\mathbf{x} \leq \mathbf{b}, \mathbf{0} \leq \mathbf{x} \leq \mathbf{1}\}$. Each entry of $\mathbf{A} \in \mathbb{R}^{m \times d}$ and $\mathbf{c} \in \mathbb{R}^d$ is drawn from the uniform distribution on $[0, 1]$. We set $d = 50$, $m = 100$, and $\mathbf{b} = 0.5k \times \mathbf{1}$. We consider various sparsities $k = 5, 10, \dots, 50$; for each k , we randomly generate 100 instances as above. We consider Multi-Greedy with $m = 2$ and $m = 5$, denoted by Multi-Greedy-2 and Multi-Greedy-5, respectively. As baselines, we employ Greedy and Random, which chooses k elements from $[d]$ uniformly at random.

Figure 2 shows the results. We see that Multi-Greedy algorithms run far faster than Greedy, and they achieve almost the same objective values as those of Greedy. Namely, for optimal production planning instances, the multi-stage strategy can accelerate Greedy at a very slight sacrifice of solution quality.

4 THEORETICAL PROPERTIES

We study theoretical properties of WMM: In Section 4.1 we show that WMM is fixed-parameter tractable (FPT) under certain conditions, and in Section 4.2 we prove that no polynomial-time algorithms can improve the $(1 - e^{-\gamma_5 \cdot k})$ -approximation guarantee in general even if SBR and SPR are bounded by some constants.

4.1 Fixed-parameter Tractability

Here we discuss the computation cost of solving WMM almost optimally. If we are to find an optimal solution for WMM, a naive approach is exhaustive search; i.e., we examine $F(\mathbf{S})$ for all $\mathbf{S} \subseteq [d]$ of size k . This, however, incurs $\Omega(d^k)$ computation cost, which becomes too large as the instance size, d , increases. Taking this into account, the following question arises: Can we solve WMM (almost) optimally without requiring an $\Omega(d^k)$ computation cost? To answer this, we use the parameterized complexity framework (Cygan et al., 2015). We regard a part of the input as a fixed parameter(s), which is denoted by \mathbf{p} and does not include the instance size, d . An algorithm is said to be FPT if it

Algorithm 2 Randomized FPT algorithm

- 1: Execute `SingleRun()` T times and return the best solution.
 - 2: **function** `SingleRun()`
 - 3: $\mathbf{S}_0 \leftarrow \emptyset$
 - 4: **for** $i = 1, \dots, k$ **do**
 - 5: Choose $j \in [d] \setminus \mathbf{S}_{i-1}$ randomly with probability $\propto F(j \mid \mathbf{S}_{i-1})$.
 - 6: $\mathbf{S}_i \leftarrow \mathbf{S}_{i-1} \cup \{j\}$
 - 7: **return** \mathbf{S}_k
-

runs in $g(\mathbf{p}) \times \text{poly}(d)$ time, where g is a computable function of \mathbf{p} . Note that, if k is a fixed parameter, algorithms that require $\Omega(d^k)$ time, including exhaustive search, are not FPT. Here, regarding k as a part of the fixed parameters, we show that ϵ -error solutions for WMM can be computed with a randomized FPT algorithm (Algorithm 2), which was originally developed by Skowron (2017) for a special case of monotone submodular maximization. Algorithm 2 performs `SingleRun()`, a randomized variant of Greedy, T times and returns the best solution. We can show that it enjoys the following guarantee for WMM:

Theorem 3. *Assume F to be $(\gamma_k, \beta_{k,d})$ -WM. Let \mathbf{S}^* be an optimal solution for problem (1) and $\tilde{F} := F([d]) - F(\mathbf{S}^*)$. For any $\epsilon > 0$, if*

$$T \geq \left\lceil \left(\frac{1}{\gamma_k \beta_{k,d}} \cdot \frac{\tilde{F} + \epsilon}{\epsilon} \right)^k \ln \delta^{-1} \right\rceil,$$

then Algorithm 2 returns solution \mathbf{S} satisfying $F(\mathbf{S}) \geq F(\mathbf{S}^) - \epsilon$ with a probability of at least $1 - \delta$.*

The key to proving Theorem 3 is the fact that the probability of choosing $j \in \mathbf{S}^*$ in each iteration can be lower bounded thanks to the weak modularity. We present the proof in Appendix C.

Since F can be evaluated in $\text{poly}(d)$ time as assumed in Section 1.2, Algorithm 2 is FPT if we regard $\mathbf{p} := (k, \gamma_k, \beta_{k,d}, \tilde{F}, \epsilon, \delta)$ as fixed parameters. Note that, since $\tilde{F} \leq F([d])$, a sufficiently large T can be computed once we obtain lower bounds of SBR and SPR, which are available for various applications as in Section 2.

While Algorithm 2 is not so practical, Theorem 3 is beneficial for studying the tractability of WMM instances. In particular, we can obtain an interesting corollary related to ℓ_0 -constrained minimization from the theorem. Let $\mathbf{x}^* := \arg\min_{\|\mathbf{x}\|_0 \leq k} l(\mathbf{x})$ be an optimal solution. As shown by Shalev-Shwartz et al. (2010), Greedy can find \mathbf{x} such that $l(\mathbf{x}) \leq l(\mathbf{x}^*) + \epsilon$ if \mathbf{x} is allowed to have $\Omega(\kappa \ln \epsilon^{-1})$ non-zeros; i.e., there is a trade-off between sparsity $\|\mathbf{x}\|_0$ and accuracy ϵ . In practice, however, \mathbf{x} is not always allowed to have sufficiently

many non-zeros. For instance, when performing feature selection for medical analysis, the number of features used for predicting a patient’s status is limited since to use many features requires the patient to undergo many medical tests, which is a considerable burden. Hence, to reveal whether we can solve ℓ_0 -constrained minimization almost optimally with limited sparsity $k \geq \|\mathbf{x}\|_0$ is an important research subject. The following corollary, which is obtained from Theorem 3 and the lower bounds of SBR and SPR (Section 2), implies that it is possible at the cost of FPT computation time; i.e, there is a time–accuracy trade-off.

Corollary 3.a. *Let $F(\mathbf{S}) = l(0) - \min_{\text{supp}(\mathbf{x}) \subseteq \mathbf{S}} l(\mathbf{x})$ for any $\mathbf{S} \subseteq [d]$ and assume l to be μ_{2k} -RSC, μ_{k+1} -RSC, $\nu_{k+1,1}$ -RSM, and ν_d -RSM. Let $\tilde{l} := l(\mathbf{x}^*) - \min_{\mathbf{x} \in \mathbb{R}^{[d]}} l(\mathbf{x})$. If Algorithm 2 runs with*

$$T \geq \left\lceil \left(\frac{\nu_{k+1,1}}{\mu_{2k}} \cdot \frac{\nu_d}{\mu_{k+1}} \cdot \frac{\tilde{l} + \epsilon}{\epsilon} \right)^k \ln \delta^{-1} \right\rceil$$

and outputs \mathbf{S} , then $\mathbf{x} = \text{argmin}_{\text{supp}(\mathbf{x}') \subseteq \mathbf{S}} l(\mathbf{x}')$ satisfies $l(\mathbf{x}) \leq l(\mathbf{x}^*) + \epsilon$ with a probability of at least $1 - \delta$.

Namely, if we take $\mathbf{p} := (k, \mu_{2k}, \mu_{k+1}, \nu_{k+1,1}, \nu_d, \tilde{l}, \epsilon, \delta)$ to be fixed parameters, ϵ -error solutions can be computed in FPT time with a high probability. Note that, unlike the aforementioned guarantee of **Greedy**, Corollary 3.a does not require $\|\mathbf{x}\|_0$ to be sufficiently large.

4.2 Hardness Result

We here prove the following hardness of improving the $(1 - e^{-\gamma_{\mathbf{S},k}})$ -approximation guarantee for WMM:

Theorem 4. *Even if F has SBR $\gamma_k = 1$ and SPR $\beta_k \geq 1/2 - o(1)$, no algorithms that evaluate F only on polynomially many subsets can achieve an approximation guarantee that exceeds $1 - e^{-1} = 1 - e^{-\gamma_k}$ for problem (1) in general.*

Note that the significance of Theorem 4 comes from SPR β_k that can be bounded by a universal constant: When curvature α , which satisfies $\beta_k \geq 1 - \alpha$, is upper bounded by a universal constant smaller than 1, then a strictly improved approximation ratio, $\frac{1}{\alpha}(1 - e^{-\alpha\gamma_{\mathbf{S},k}})$, can be obtained thanks to (Bian et al., 2017). Namely, Theorem 4 reveals a non-trivial theoretical gap between SPR β_k and curvature α . Below we describe a proof sketch, and the full proof is presented in Appendix D.1.

Proof sketch. We make a WM function that is hard to maximize approximately. As with the proof of (Nemhauser and Wolsey, 1978), given unknown subset \mathbf{M} of size k , we show that to achieve an approximation guarantee that exceeds $1 - e^{-\gamma_k}$ is at least as hard as to find \mathbf{S} such that $|\mathbf{S} \cap \mathbf{M}| > r$ and $|\mathbf{S}| \leq p_r^k := 2k - r + 1$,

where $r > 0$ is any fixed integer; this cannot be solved via polynomially many queries. To this end, we use F that satisfies the following conditions: $F(\mathbf{S})$ value depends on $|\mathbf{S}|$ and $|\mathbf{S} \cap \mathbf{M}|$ for any $\mathbf{S} \subseteq [d]$ and only on $|\mathbf{S}|$ if $|\mathbf{S} \cap \mathbf{M}| \leq r$ or $|\mathbf{S}| > p_r^k$, which, roughly speaking, means that the information about F values is useless. By using such function F , we can obtain the hardness result. The main difficulty remained in the above proof is to show that F is WM. In particular, obtaining $\beta_k \geq 1/2 - o(1)$ is the most challenging part. To prove this, we first rewrite SPR as $\beta_k = \min_{\mathbf{L}, \mathbf{S} \subseteq [d]} \left\{ \frac{F(\mathbf{S}|\mathbf{L})}{\sum_{j \in \mathbf{S}} F(j|\mathbf{L})} \mid \mathbf{L} \cap \mathbf{S} = \emptyset, |\mathbf{L}| \leq k, |\mathbf{S}| \leq k \right\}$, where we regard $0/0 = 1$. Then, by carefully designing F and using the fact that $F(\mathbf{S})$ depends only on $|\mathbf{S}|$ and $|\mathbf{S} \cap \mathbf{M}|$, we can lower bound β_k by the minimum value of some function with only three variables, and the minimum value can be proved to be at least $\frac{1}{2} - \frac{1}{2} \cdot \frac{r-1}{2k-r+1}$ (see, Lemma A.4 in Appendix D.1). By letting k increase with d and setting d at a sufficiently large value, we obtain the lower bound on β_k . \square

Given solution \mathbf{S} of **Greedy**, we always have $\gamma_{\mathbf{S},k} \geq \gamma_k$. Therefore, Theorem 4 implies that, even if $\beta_k (\leq \beta_{\mathbf{S},k})$ is lower bounded by a value that can be arbitrarily close to $1/2$, no polynomial-time algorithms can improve the $(1 - e^{-\gamma_{\mathbf{S},k}})$ -approximation guarantee in general.

We remark that it may be possible to improve the approximation ratio for some easier subclasses of WMM; for example, if $\gamma_d (\leq \gamma_k)$ and $\beta_d (\leq \beta_k)$ are bounded, we may be able to obtain a better ratio than $1 - e^{-\gamma_d}$ by using β_d . We discuss this topic in Appendix D.2. We also remark that Theorem 4 does not contradict the FPT result (Theorem 3) for the following reason: Theorem 4 is proved by using sparsity k that increases with d , and we cannot regard such a k as a fixed parameter.

5 CONCLUSION

We studied WMM, a class of non-submodular maximization that can model various practical problems. We proved guarantees of multi-stage algorithms, which generalize and improve some existing results, and confirmed their effectiveness via experiments. We then proved the fixed-parameter tractability of WMM, which yields the time–accuracy trade-off for ℓ_0 -constrained minimization as a byproduct, and the hardness of improving the $(1 - e^{-\gamma_{\mathbf{S},k}})$ -approximation guarantee.

Recent studies (Khanna et al., 2017b; Qian et al., 2018) provided various techniques for accelerating greedy algorithms, and greedy-style methods for many different problem settings have also been studied (Bogunovic et al., 2018; Fujii and Soma, 2018). It will be interesting to study how to incorporate the multi-stage approach into those methods for further acceleration.

Acknowledgements

The author is grateful to Kaito Fujii and anonymous reviewers for providing valuable comments.

References

- Bai, W. and Bilmes, J. (2018). Greed is still good: Maximizing monotone Submodular+Supermodular (BP) functions. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80, pages 304–313. PMLR.
- Balkanski, E. and Singer, Y. (2018). The adaptive complexity of maximizing a submodular function. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1138–1151. ACM.
- Bian, A. A., Buhmann, J. M., Krause, A., and Tschitschek, S. (2017). Guarantees for greedy maximization of non-submodular functions with applications. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 498–507. PMLR.
- Bogunovic, I., Zhao, J., and Cevher, V. (2018). Robust maximization of non-submodular objectives. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pages 890–899. PMLR.
- Calinescu, G., Chekuri, C., Pál, M., and Vondrák, J. (2011). Maximizing a monotone submodular function subject to a matroid constraint. *SIAM J. Comput.*, 40(6):1740–1766.
- Candès, E. J., Romberg, J. K., and Tao, T. (2006). Stable signal recovery from incomplete and inaccurate measurements. *Comm. Pure Appl. Math.*, 59(8):1207–1223.
- Chen, L., Feldman, M., and Karbasi, A. (2018). Weakly submodular maximization beyond cardinality constraints: Does randomization help greedy? In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 804–813. PMLR.
- Chierichetti, F., Das, A., Dasgupta, A., and Kumar, R. (2015). Approximate modularity. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 1143–1162.
- Conforti, M. and Cornuéjols, G. (1984). Submodular set functions, matroids and the greedy algorithm: Tight worst-case bounds and some generalizations of the Rado-Edmonds theorem. *Discrete Appl. Math.*, 7(3):251–274.
- Cygan, M., Fomin, F. V., Kowalik, L., Lokshtanov, D., Marx, D., Pilipczuk, M., Pilipczuk, M., and Saurabh, S. (2015). *Parameterized Algorithms*. Springer Publishing Company, Incorporated, 1st edition.
- Das, A. and Kempe, D. (2018). Approximate submodularity and its applications: Subset selection, sparse approximation and dictionary selection. *J. Mach. Learn. Res.*, 19(3):1–34.
- Elenberg, E. R., Dimakis, A. G., Feldman, M., and Karbasi, A. (2017). Streaming weak submodularity: Interpreting neural networks on the fly. In *Advances in Neural Information Processing Systems 30*, pages 4044–4054. Curran Associates, Inc.
- Elenberg, E. R., Khanna, R., Dimakis, A. G., and Negahban, S. (2018). Restricted strong convexity implies weak submodularity. *Ann. Statist.*, 46(6B):3539–3568.
- Feige, U. (1998). A threshold of $\ln n$ for approximating set cover. *J. ACM*, 45(4):634–652.
- Feige, U. and Izsak, R. (2013). Welfare maximization and the supermodular degree. In *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science*, pages 247–256. ACM.
- Fujii, K. and Soma, T. (2018). Fast greedy algorithms for dictionary selection with generalized sparsity constraints. In *Advances in Neural Information Processing Systems 31*, pages 4749–4758. Curran Associates, Inc.
- Harshaw, C., Feldman, M., Ward, J., and Karbasi, A. (2019). Submodular maximization beyond non-negativity: Guarantees, fast algorithms, and applications. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pages 2634–2643. PMLR.
- Horel, T. and Singer, Y. (2016). Maximization of approximately submodular functions. In *Advances in Neural Information Processing Systems 29*, pages 3045–3053. Curran Associates, Inc.
- Hu, H., Grubb, A., Bagnell, J. A., and Hebert, M. (2016). Efficient feature group sequencing for anytime linear prediction. In *Proceedings of the 32nd Conference on Uncertainty in Artificial Intelligence*, pages 279–288. AUAI Press.
- Iyer, R. K., Jegelka, S., and Bilmes, J. A. (2013). Curvature and optimal algorithms for learning and minimizing submodular functions. In *Advances in Neural Information Processing Systems 26*, pages 2742–2750. Curran Associates, Inc.
- Jain, P., Tewari, A., and Kar, P. (2014). On iterative hard thresholding methods for high-dimensional M-estimation. In *Advances in Neural Information Processing Systems 27*, pages 685–693. Curran Associates, Inc.

- Khanna, R., Elenberg, E. R., Dimakis, A. G., Ghosh, J., and Negahban, S. (2017a). On approximation guarantees for greedy low rank optimization. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 1837–1846. PMLR.
- Khanna, R., Elenberg, E. R., Dimakis, A. G., Negahban, S., and Ghosh, J. (2017b). Scalable greedy feature selection via weak submodularity. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54, pages 1560–1568. PMLR.
- Krause, A. and Cevher, V. (2010). Submodular dictionary selection for sparse representation. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, pages 567–574. Omnipress.
- Kumar, R., Moseley, B., Vassilvitskii, S., and Vattani, A. (2015). Fast greedy algorithms in mapreduce and streaming. *ACM Trans. Parallel Comput.*, 2(3):14:1–14:22.
- Li, X., Zhao, T., Arora, R., Liu, H., and Haupt, J. (2016). Stochastic variance reduced optimization for nonconvex sparse learning. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48, pages 917–925. PMLR.
- Liberty, E. and Sviridenko, M. (2017). Greedy minimization of weakly supermodular set functions. In *Proceedings of Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, volume 81, pages 19:1–19:11. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- Lin, H. and Bilmes, J. A. (2011). 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. Association for Computational Linguistics.
- Marsousi, M., Abhari, K., Babyn, P., and Alirezaie, J. (2013). MULTI-STAGE OMP sparse coding using local matching pursuit atoms selection. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 1783–1787.
- Natarajan, B. K. (1995). Sparse approximate solutions to linear systems. *SIAM J. Optim.*, 24(2):227–234.
- Nemhauser, G. L. and Wolsey, L. A. (1978). Best algorithms for approximating the maximum of a submodular set function. *Math. Oper. Res.*, 3(3):177–188.
- Nemhauser, G. L., Wolsey, L. A., and Fisher, M. L. (1978). An analysis of approximations for maximizing submodular set functions-I. *Math. Program.*, 14(1):265–294.
- Olson, R. S., La Cava, W., Orzechowski, P., Urbanowicz, R. J., and Moore, J. H. (2017). PMLB: A large benchmark suite for machine learning evaluation and comparison. *BioData Min.*, 10(1):36.
- Qian, C., Yu, Y., and Tang, K. (2018). Approximation guarantees of stochastic greedy algorithms for subset selection. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 1478–1484. International Joint Conferences on Artificial Intelligence Organization.
- Qian, S. and Singer, Y. (2019). Fast parallel algorithms for statistical subset selection problems. In *Advances in Neural Information Processing Systems 32*, pages 5073–5082. Curran Associates, Inc.
- Shalev-Shwartz, S., Srebro, N., and Zhang, T. (2010). Trading accuracy for sparsity in optimization problems with sparsity constraints. *SIAM J. Optim.*, 20(6):2807–2832.
- Shalev-Shwartz, S. and Zhang, T. (2016). Accelerated proximal stochastic dual coordinate ascent for regularized loss minimization. *Math. Program.*, 155(1):105–145.
- Skowron, P. (2017). FPT approximation schemes for maximizing submodular functions. *Inform. Comput.*, 257:65–78.
- Soma, T. and Yoshida, Y. (2018). A new approximation guarantee for monotone submodular function maximization via discrete convexity. In *Proceedings of the 45th International Colloquium on Automata, Languages, and Programming*, volume 107, pages 99:1–99:14. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- Sviridenko, M., Vondrák, J., and Ward, J. (2015). Optimal approximation for submodular and supermodular optimization with bounded curvature. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1134–1148. SIAM.
- Takeda, A., Niranjana, M., Gotoh, J., and Kawahara, Y. (2013). Simultaneous pursuit of out-of-sample performance and sparsity in index tracking portfolios. *Comput. Manag. Sci.*, 10(1):21–49.
- Wang, Z., Moran, B., Wang, X., and Pan, Q. (2016). Approximation for maximizing monotone non-decreasing set functions with a greedy method. *J. Comb. Optim.*, 31(1):29–43.
- Wei, K., Iyer, R., and Bilmes, J. (2014). Fast multi-stage submodular maximization. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32, pages 1494–1502. PMLR.
- Yuan, X.-T., Li, P., and Zhang, T. (2018). Gradient hard thresholding pursuit. *J. Mach. Learn. Res.*, 18(166):1–43.

Zhou, Y. and Spanos, C. J. (2016). Causal meets sub-modular: Subset selection with directed information. In *Advances in Neural Information Processing Systems 29*, pages 2649–2657. Curran Associates, Inc.