
Fast Maximization of Non-Submodular, Monotonic Functions on the Integer Lattice

Alan Kuhnle¹ J. David Smith¹ Victoria G. Crawford¹ My T. Thai¹

Abstract

The optimization of submodular functions on the integer lattice has received much attention recently, but the objective functions of many applications are non-submodular. We provide two approximation algorithms for maximizing a non-submodular function on the integer lattice subject to a cardinality constraint; these are the first algorithms for this purpose that have polynomial query complexity. We propose a general framework for influence maximization on the integer lattice that generalizes prior works on this topic, and we demonstrate the efficiency of our algorithms in this context.

1. Introduction

As a natural extension of finite sets S (equivalently, $\{0, 1\}^S$), optimization of discrete functions on the integer lattice \mathbb{N}^S has received attention recently (Alon et al., 2012; Demaine et al., 2014; Soma & Yoshida, 2015). As an example, consider the placement of sensors in a water network (Krause et al., 2008a); in the set version, each sensor takes a value in $\{0, 1\}$, which corresponds to whether the sensor was placed. In the lattice version (Soma & Yoshida, 2015), each sensor has a power level in $\{0, \dots, b\} \subseteq \mathbb{N}$, to which the sensitivity of the sensor is correlated. As a second example, consider the influence maximization problem (Kempe et al., 2003); instead of the binary seeding of a user, the lattice version enables partial incentives or discounts to be used (Demaine et al., 2014).

Although many results from the optimization of submodular set functions have been generalized to the integer lattice (Soma & Yoshida, 2015; 2016; Ene & Nguyen, 2016), many objective functions arising from applications are not submodular (Bian et al., 2017b; Lin et al., 2017; Das &

Kempe, 2011; Horel & Singer, 2016). In this work, we consider maximization subject to a cardinality constraint (MCC), where the function f to be maximized may be non-submodular. Let $k \in \mathbb{N}$ (the budget), $\mathbf{b} \in (\mathbb{N} \cup \{\infty\})^S$ (the box), and let $f : \{\mathbf{x} \in \mathbb{N}^S : \mathbf{x} \leq \mathbf{b}\} \rightarrow \mathbb{R}^+$ (the objective) be a non-negative and monotonic¹ function with $f(\mathbf{0}) = 0$. Then determine

$$\max_{\|\mathbf{w}\|_1 \leq k} f(\mathbf{w}), \quad (\text{MCC})$$

where $\mathbf{w} = (\mathbf{w}_s)_{s \in S} \in \mathbb{N}^S$, $\|\mathbf{w}\|_1 = \sum_{s \in S} |\mathbf{w}_s|$.

Since the integer lattice may be represented as a multiset of size $k|S|$, one may use results for Problem MCC with non-submodular set functions. In particular, the tight ratio $\frac{1}{\alpha}(1 - e^{-\alpha\gamma_s})$ of the standard greedy algorithm by Bian et al. (2017b), where α, γ_s are discussed below, applies with the lattice adaptation of the standard greedy algorithm (StandardGreedy) given in Alg. 1. However, this approach requires $\Omega(|S|k)$ queries of f , which is not polynomial in the input² size $O(|S| \log k)$. Even for applications with set functions, $\Omega(|S|k)$ queries may be prohibitive, and researchers (Leskovec et al., 2007; Mirzasoleiman et al., 2015; Badanidiyuru & Vondrák, 2014) have sought ways to speed up the StandardGreedy algorithm. Unfortunately, these approaches rely upon the submodularity of f , and there has been no analogous effort for non-submodular functions.

To quantify the non-submodularity of a lattice function f , we generalize the following quantities defined for set functions to the lattice: (1) the diminishing-return (DR) ratio γ_d of f (Lehmann et al., 2006), (2) the weak DR ratio γ_s of f (Das & Kempe, 2011), and (3) the generalized curvature α of f (Bian et al., 2017b). Our main contributions are:

- To speed up StandardGreedy (Alg. 1), we adapt the threshold greedy framework of Badanidiyuru & Vondrák (2014) to non-submodular functions; this yields an algorithm (ThresholdGreedy, Alg. 2) with approximation ratio $(1 - e^{-\gamma_d\gamma_s} - \eta)$, for

¹for all $\mathbf{v} \leq \mathbf{w}$ (coordinate-wise), $f(\mathbf{v}) \leq f(\mathbf{w})$

²The input is considered to be the vector \mathbf{b} of length $n = |S|$ and the number k represented in $\log k$ bits (w.l.o.g. each component of \mathbf{b} is at most k); the function is regarded as an oracle and hence does not contribute to input size.

¹University of Florida, Gainesville, Florida. Correspondence to: Alan Kuhnle <kuhnle@ufl.edu>, My T. Thai <mythai@ufl.edu>.

Algorithm 1 StandardGreedy

```

1: Input:  $f \in \mathcal{F}_{\mathbf{b}}, k \in \mathbb{N}, \mathbf{b} \in \mathbb{N}^S$ 
2: Output:  $\mathbf{g} \in \mathbb{N}^S$ 
3:  $\mathbf{g} \leftarrow \mathbf{0}$ 
4: for  $i = 1$  to  $k$  do
5:    $\mathbf{g} \leftarrow \mathbf{g} + \arg \max_{s \in S: \mathbf{g} + \mathbf{s} \leq \mathbf{b}} \delta_s(\mathbf{g})$ 
6: return  $\mathbf{g}$ 

```

any $\eta > 0$, the first approximation algorithm with polynomial query complexity for Problem MCC on the lattice. The query complexity of the StandardGreedy algorithm is improved from $\Omega(nk)$ to $O(n \log k \log_{\kappa}(\varepsilon^2/k))$, where $\kappa, \varepsilon \in (0, 1)$ are parameters of ThresholdGreedy.

- We introduce the novel approximation algorithm FastGreedy, which combines elements of StandardGreedy and ThresholdGreedy to improve the performance ratio to $(1 - e^{-\beta^* \gamma_s} - \eta)$, where β^* is at least γ_d and in many cases³ is determined by the algorithm. Furthermore, FastGreedy exploits the non-submodularity of the function to decrease its runtime in practice without sacrificing its performance guarantee, while maintaining the same worst-case query complexity as ThresholdGreedy up to a constant factor.
- To demonstrate our algorithms, we introduce a general budget allocation problem for viral marketing, which unifies submodular *influence maximization* (IM) under the independent cascade model (Kempe et al., 2003) with the non-submodular *boosting problem* (Lin et al., 2017) and in addition allows partial incentives. We prove a lower bound on the DR and weak DR ratios for this unified framework, and we experimentally validate our proposed algorithms in this setting.

2. Related Work

The study of optimization of submodular set functions is too extensive to give a comprehensive overview. On the integer lattice, there have been many efforts to maximize submodular functions, e.g. Soma & Yoshida (2017); Bian et al. (2017a); Gottschalk & Peis (2016). To the best of our knowledge, we are the first to study the optimization of non-submodular functions on the integer lattice. In the following discussion, we primarily restrict our attention to the maximization of monotonic, submodular lattice functions subject to a cardinality constraint and the maximization of non-submodular set functions.

³When the solution \mathbf{g} returned by FastGreedy satisfies $\|\mathbf{g}\|_1 = k$. Otherwise, an upper bound on β^* is returned.

Reduction of Ene & Nguyen (2016). Ene & Nguyen (2016) have given a polynomial-time reduction from the lattice to a set that enables unified translation of submodular optimization strategies to DR-submodular (i.e. DR ratio $\gamma_d = 1$, see Section 3) functions on the integer lattice. Since this translation is designed for DR-submodular functions, it does not give a polynomial-time algorithm for Problem MCC when f is non-submodular. Specifically, for the case of maximization subject to a cardinality constraint, Ene & Nguyen (2016) rely upon the threshold greedy algorithm for submodular set functions (Badanidiyuru & Vondrák, 2014), which does not work for non-submodular functions without modifications such as the ones in our paper.

Threshold Greedy and Lattice Optimization. To speed up the StandardGreedy for submodular set functions, Badanidiyuru & Vondrák (2014) introduced the threshold greedy framework, which speeds up the StandardGreedy algorithm for maximizing submodular set functions under cardinality constraint from $O(nk)$ function evaluations to $O(\frac{n}{\varepsilon} \log \frac{n}{\varepsilon})$, and it maintains the approximation ratio $(1 - 1/e - \varepsilon)$, for $\varepsilon > 0$. Soma & Yoshida (2016) adapted the threshold approach for efficiently maximizing DR-submodular functions on the integer lattice and provided $(1 - 1/e - \varepsilon)$ -approximation algorithms. Other adaptations of the threshold approach of Badanidiyuru & Vondrák (2014) to the integer lattice include (Ene & Nguyen, 2016; Soma & Yoshida, 2015). Elenberg et al. (2017) recently extended a similar algorithm (Sieve Streaming of Badanidiyuru et al. (2014)) to non-submodular functions in a somewhat similar fashion. The key difference is Sieve Streaming makes one pass over the data using many geometrically spaced thresholds in parallel, whereas the approach of Badanidiyuru & Vondrák (2014) makes many sequential passes at geometrically spaced thresholds.

Our ThresholdGreedy algorithm is an adaptation of the algorithm of Soma & Yoshida (2016) for DR-submodular maximization to non-submodular functions. The non-submodularity requires new analysis, in the following specific ways: (1) during the binary search phase, we cannot guarantee that we find the maximum number of copies whose average gain exceeds the threshold τ ; hence, we must settle for any number of copies whose average gain exceeds τ , while ensuring that the gain of adding one additional copy falls belows τ . (2) To prove the performance ratio, we require a combination of the DR ratio γ_d and the weak DR ratio γ_s .

The very recent work of Qian et al. (2018) considers the same problem as in our paper, and they define the same DR ratio as ours; their lattice submodularity ratio is different from the weak DR ratio, however. The algorithm (POMS) in Qian et al. (2018) is based upon a Pareto optimization

technique, which is substantially different from our approach; our emphasis is on developing an approximation algorithm with runtime logarithmic in k , while POMS has running time $\Omega(k^2n)$.

Optimization of Non-Submodular Set Functions. For non-submodular set functions, the weak DR ratio γ_s was introduced by Das & Kempe (2011) under the name submodularity ratio; we generalize γ_s to lattice functions in Section 3, and we show the DR ratio $\gamma_d \leq \gamma_s$. Bian et al. (2017b) introduced *generalized curvature* α of a set function, an analogous concept to the DR ratio as we discuss in Section 3. Bian et al. (2017b) extended the analysis of Conforti & Cornuéjols (1984) to non-submodular set functions; together with the weak DR ratio γ_s , they proved Standard-Greedy has tight approximation ratio $\frac{1}{\alpha}(1 - e^{-\gamma_s\alpha})$ under cardinality constraint. The DR ratio γ_d is introduced as inverse curvature in Bogunovic et al. (2018), wherein robust maximization of set functions is considered; the DR ratio has also been introduced by Lehmann et al. (2006); Qian et al. (2018).

Many other notions of non-submodular set functions have been introduced (Krause et al., 2008b; Horel & Singer, 2016; Borodin et al., 2014; Feige & Izsak, 2013). For a comprehensive discussion of the relation of these and additional notions to the weak DR ratio γ_s , we refer the reader to Bian et al. (2017b).

3. Non-Submodularity on the Lattice

In this section, we define the lattice versions of DR ratio γ_d , weak DR ratio γ_s , and generalized curvature α , which are used in the approximation ratios proved in Section 4.

Notations. For each $s \in S$, let \mathbf{s} be the unit vector with 1 in the coordinate corresponding to s , and 0 elsewhere. We write $\delta_{\mathbf{w}}(\mathbf{v}) = f(\mathbf{v} + \mathbf{w}) - f(\mathbf{v})$ for $\mathbf{v}, \mathbf{w} \in \mathbb{N}^S$. Given a box in the integer lattice $\mathbf{b} \in \mathbb{N}^S$, let the set of all non-negative, monotonic lattice functions with $f(\mathbf{0}) = 0$, and domain $\{\mathbf{x} \in \mathbb{N}^S : \mathbf{x} \leq \mathbf{b}\}$ be denoted $\mathcal{F}_{\mathbf{b}}$. It is often useful to think of a vector $\mathbf{v} \in \mathbb{N}^S$ as a multi-set containing \mathbf{v}_s copies of $s \in S$, where \mathbf{v}_s is the value of \mathbf{v} 's coordinate corresponding to s . We use the notation $\{\mathbf{v}\}$ to represent the multiset corresponding to the vector \mathbf{v} . Finally, we define $\mathbf{v} \vee \mathbf{w}$ and $\mathbf{v} \wedge \mathbf{w}$ for $\mathbf{v}, \mathbf{w} \in \mathbb{N}^S$ to be the vector with the coordinate-wise maximum and minimum respectively. Rather than an algorithm taking an explicit description of the function f as input, we consider the function f as an oracle and measure the complexity of an algorithm in terms of the number of oracle calls or queries.

We begin with the related concepts of DR ratio and generalized curvature.

Definition 1. Let $f \in \mathcal{F}_{\mathbf{b}}$. The *diminishing-return (DR)*

ratio of f , $\gamma_d(f)$, is the maximum value in $[0, 1]$ such that for any $s \in S$, and for all $\mathbf{v} \leq \mathbf{w}$ such that $\mathbf{w} + \mathbf{s} \leq \mathbf{b}$, $\gamma_d(f)\delta_{\mathbf{s}}(\mathbf{w}) \leq \delta_{\mathbf{s}}(\mathbf{v})$.

Definition 2. Let $f \in \mathcal{F}_{\mathbf{b}}$. The *generalized curvature of f* , $\alpha(f)$, is the minimum value in $[0, 1]$ such that for any $s \in S$, and for all $\mathbf{v} \leq \mathbf{w}$ such that $\mathbf{w} + \mathbf{s} \leq \mathbf{b}$, $\delta_{\mathbf{s}}(\mathbf{w}) \geq (1 - \alpha(f))\delta_{\mathbf{s}}(\mathbf{v})$.

The DR ratio extends the notion of DR-submodularity of Soma & Yoshida (2015), which is obtained as the special case $\gamma_d = 1$. Generalized curvature for set functions was introduced in Bian et al. (2017b). Notice that α results in lower bounds on the marginal gain of \mathbf{s} to a vector \mathbf{w} , while γ_d results in upper bounds on the same quantity: $(1 - \alpha)\delta_{\mathbf{s}}(\mathbf{v}) \leq \delta_{\mathbf{s}}(\mathbf{w}) \leq \frac{1}{\gamma_d}\delta_{\mathbf{s}}(\mathbf{v})$, whenever $\mathbf{v} \leq \mathbf{w}$ and the above expressions are defined. Next, we generalize the weak DR ratio of Das & Kempe (2011) to the integer lattice.

Definition 3. Let $f \in \mathcal{F}_{\mathbf{b}}$. The *weak DR ratio of f* , $\gamma_s(f)$, is the maximum value in $[0, 1]$ such that for all \mathbf{v}, \mathbf{w} , such that $\mathbf{v} \leq \mathbf{w}$, $\gamma_s(f)(f(\mathbf{w}) - f(\mathbf{v})) \leq \sum_{s \in \{\mathbf{w} - \mathbf{v}\}} \delta_{\mathbf{s}}(\mathbf{v})$.

The next proposition, proved in Appendix C, shows the relationship between DR ratio and weak DR ratio.

Proposition 1. For all $f \in \mathcal{F}_{\mathbf{b}}$, $\gamma_d(f) \leq \gamma_s(f)$. The function f is DR submodular iff $\gamma_d(f) = \gamma_s(f) = 1$.

In the rest of this work, we will parameterize functions by the non-submodularity ratios defined above and partition functions into the sets $\mathcal{F}_{\mathbf{b}}^{\gamma_d, \gamma_s, \alpha} = \{f \in \mathcal{F}_{\mathbf{b}} : \gamma_d(f) = \gamma_d, \gamma_s(f) = \gamma_s, \alpha(f) = \alpha\}$.

Greedy versions. In the proofs of this paper, the full power of the parameters defined above is not required. It suffices to consider restricted versions, where the maximization is taken over only those vectors which appear in the ratio proofs. We define these greedy versions in Appendix B and include more discussion in Remark 1 of Section 4.1.

4. Algorithms

4.1. The ThresholdGreedy Algorithm

In this section, we present the algorithm ThresholdGreedy (Alg. 2) to approximate Problem MCC with ratio $1 - e^{-\gamma_g \gamma_s} - \eta$ with polynomial query complexity. Appendix D contains the proofs of all lemmas, claims, and omitted details from this section.

Description. ThresholdGreedy operates by considering decreasing thresholds for the marginal gain in its outer **for** loop; for each threshold τ , the algorithm adds on line 2 elements whose marginal gain exceeds τ as described be-

low. The parameter $\kappa \in (0, 1)$ determines the stepsize between successive thresholds; the algorithm continues until the budget k is met (line 2) or the threshold is below a minimum value dependent on the parameter $\varepsilon \in (0, 1)$.

Intuitively, the goal of the threshold approach (Badani-diyuru & Vondrák, 2014) for submodular set functions is as follows. At each threshold (*i.e.*, iteration of the outer **for** loop), add all elements whose marginal gain exceeds τ to the solution \mathbf{g} . On the lattice, adding all copies of $s \in S$ whose average gain exceeds τ on line 2 would require the addition of the maximum multiple l s such that the average marginal gain exceeds τ :

$$\delta_{ls}(\mathbf{g}) \geq l\tau, \quad (\text{P1})$$

as in the threshold algorithm of Soma & Yoshida (2016) for DR-submodular functions, in which the maximum l is identified by binary search. However, since f is not DR-submodular, it is not always the case that $\delta_s(\mathbf{g} + ls) \geq \delta_s(\mathbf{g} + (l+1)s)$, for each l . For this reason, we cannot find the maximum such l by binary search. Furthermore, even if we found the maximum l for each $s \in S$, we could not guarantee that all elements of marginal gain at least τ were added due to the non-submodularity of f : an element whose gain is less than τ when considered in the inner **for** loop might have gain greater than τ after additional elements are added to the solution.

ThresholdGreedy more conservatively ensures that the number l chosen for each $s \in S$ satisfies both (P1) and

$$\delta_s(\mathbf{g} + ls) < \tau, \quad (\text{P2})$$

but it is not necessarily the maximum such l .

Pivot. Any l satisfying both (P1) and (P2) is termed a *pivot*⁴ with respect to \mathbf{g}, s, τ . Perhaps surprisingly, a valid pivot can be found with binary search in $O(\log \mathbf{b}_{max}) = O(\log k)$ function queries, where $\mathbf{b}_{max} = \max_{s \in S} \mathbf{b}_s$; discussion of BinarySearchPivot and proof of this results is provided in Appendix D, Lemma 2. By finding a pivot for each $s \in S$, ThresholdGreedy does not attempt to add all elements exceeding the marginal gain of threshold τ ; instead, ThresholdGreedy maintains the following property at each threshold.

Property 1. Let \mathbf{g}_τ be the solution of ThresholdGreedy immediately after the iteration of the outer **for** loop corresponding to threshold τ . Then for each $s \in S$, there exists $\mathbf{h} \leq \mathbf{g}_\tau$ such that $\delta_s(\mathbf{h}) < \tau$.

⁴For convenience, we also define the maximum value of l , $l_{max} = \min\{\mathbf{b}_s - \mathbf{g}_s, k - \|\mathbf{g}\|_1\}$ to be a pivot if l_{max} satisfies (P1) only, and set $\delta_s(\mathbf{g} + l_{max}s) = 0$, so that all pivots satisfy both properties.

Algorithm 2 ThresholdGreedy

```

1: Input:  $f \in \mathcal{F}_{\mathbf{b}}, k \in \mathbb{N}, \kappa, \varepsilon \in (0, 1)$ .
2: Output:  $\mathbf{g} \in \mathbb{N}^S$ 
3:  $\mathbf{g} \leftarrow \mathbf{0}, M \leftarrow \max_{s \in S} f(s)$ .
4: for  $(\tau = M; \tau \geq \frac{\kappa \varepsilon^2 M}{k}; \tau \leftarrow \kappa \tau)$  do
5:   for  $s \in S$  do
6:      $l \leftarrow \text{BinarySearchPivot}(f, \mathbf{g}, \mathbf{b}, s, k, \tau)$ 
7:      $\mathbf{g} \leftarrow \mathbf{g} + ls$ 
8:     if  $\|\mathbf{g}\|_1 = k$  then
9:       return  $\mathbf{g}$ 
10: return  $\mathbf{g}$ 
    
```

Performance ratios. Next, we present the main result of this section, the performance guarantee involving the DR and weak DR ratios. Observe that the query complexity of ThresholdGreedy is polynomial in the input size $O(n \log k)$.

Theorem 1. Let an instance of Problem MCC be given, with $f \in \mathcal{F}_{\mathbf{b}}^{\gamma_d, \gamma_s, \alpha}$. If \mathbf{g} is the solution returned by ThresholdGreedy and Ω is an optimal solution to this instance, then

$$f(\mathbf{g}) \geq (1 - e^{-\kappa \gamma_d \gamma_s} - \varepsilon) f(\Omega).$$

The query complexity of ThresholdGreedy is $O(n \log k \log_{\kappa}(\varepsilon^2/k))$.

If $\eta > 0$ is given, the assignment $\kappa = (1 - \eta/2), \varepsilon = \eta/2$ yields performance ratio at least $1 - e^{-\gamma_d \gamma_s} - \eta$.

Proof. If $\gamma_d < \varepsilon$, the ratio holds trivially; so assume $\gamma_d \geq \varepsilon$. The proof of the following claim requires an application of the DR ratio.

Claim 1. Let \mathbf{g} be produced by a modified version of ThresholdGreedy that continues until $\|\mathbf{g}\|_1 = k$. If we show $f(\mathbf{g}) \geq (1 - e^{-\kappa \gamma_d \gamma_s})f(\Omega)$, the results follows.

Thus, for the rest of the proof let \mathbf{g} be as described in Claim 1. Let \mathbf{g}^t be the value of \mathbf{g} after the t th execution of line 2 of ThresholdGreedy. Let l^t be the t th pivot, such that $\mathbf{g}^t = \mathbf{g}^{t-1} + l^t s^t$. The next claim lower bounds the marginal gain in terms of the DR ratio and the previous threshold.

Claim 2. For each $s \in \{\Omega - \mathbf{g}^{t-1} \wedge \Omega\}$,

$$l^t \gamma_d \kappa \delta_s(\mathbf{g}^{t-1}) \leq f(\mathbf{g}^t) - f(\mathbf{g}^{t-1}).$$

Proof. Let τ be the threshold at which $l^t s^t$ is added to \mathbf{g}^{t-1} ; let $s \in \{\Omega - \mathbf{g}^{t-1} \wedge \Omega\}$. If τ is the first threshold, $\gamma_d \delta_s(\mathbf{g}^{t-1}) \leq \delta_s(\mathbf{0}) \leq \tau < \frac{\tau}{\kappa}$. If τ is not the first threshold, $\tau' = \tau/\kappa$ is the previous threshold value of the previous iteration of the outer **for** loop. By Property 1, there

exists $\mathbf{h} \leq \mathbf{g}_{\tau'} \leq \mathbf{g}^{t-1}$, such that $\delta_s(\mathbf{h}) < \tau'$. By the definition of DR ratio, $\gamma_d \delta_s(\mathbf{g}^{t-1}) \leq \delta_s(\mathbf{h}) < \tau' = \tau/\kappa$.

In either case, by the fact that property (P1) of a pivot holds for l^t , we have

$$f(\mathbf{g}^t) - f(\mathbf{g}^{t-1}) \geq l^t \tau \geq l^t \gamma_d \kappa \delta_s(\mathbf{g}^{t-1}). \quad \square$$

Since $|\Omega| \leq k$, we have by Claim 2

$$\begin{aligned} f(\mathbf{g}^t) - f(\mathbf{g}^{t-1}) &\geq \frac{l^t \gamma_d \kappa}{k} \sum_{s \in \{\Omega - (\mathbf{g}^{t-1} \wedge \Omega)\}} \delta_s(\mathbf{g}^{t-1}) \\ &= \frac{l^t \gamma_d \kappa}{k} \sum_{s \in \{\Omega \vee \mathbf{g}^{t-1} - \mathbf{g}^{t-1}\}} \delta_s(\mathbf{g}^{t-1}) \\ &\geq \frac{l^t \gamma_d \gamma_s \kappa}{k} (f(\Omega \vee \mathbf{g}^{t-1}) - f(\mathbf{g}^{t-1})) \\ &\geq \frac{l^t \gamma_d \gamma_s \kappa}{k} (f(\Omega) - f(\mathbf{g}^{t-1})), \end{aligned}$$

where the equality follows from the lattice identity $\mathbf{v} \vee \mathbf{w} - \mathbf{v} = \mathbf{w} - \mathbf{v} \wedge \mathbf{w}$ for all $\mathbf{v}, \mathbf{w} \in \mathbb{N}^S$, the second inequality is by definition of the weak DR ratio, and the third inequality is from monotonicity. From here, we obtain $f(\mathbf{g}) \geq \left(1 - \prod_{t=1}^T \left(1 - \frac{l^t \gamma_d \gamma_s \kappa}{k}\right)\right) f(\Omega)$, from which the hypothesis of Claim 1 follows.

Query complexity. The **for** loop on line 2 (Alg. 2) iterates at most $\log_{\kappa} \varepsilon^2/k$ times; each iteration requires $O(n \log k)$ queries, by Lemma 2. \square

For additional speedup, the inner **for** loop of FastGreedy may be parallelized, which divides the factor of n in the query complexity by the number of threads but worsens the performance ratio; in addition to γ_d, γ_s , the generalized curvature α is required in the proof.

Corollary 1. *If the inner **for** loop of ThresholdGreedy is parallelized, the performance ratio becomes $1 - e^{-(1-\alpha)\gamma_d \gamma_s} - \eta$, for $\eta > 0$.*

Remark 1. A careful analysis of the usage of γ_d, γ_s in the proof of Theorem 1 shows that the full power of the definitions of these quantities is not required. Rather, it is sufficient to consider ThresholdGreedy versions of these parameters, as defined in Appendix B. In the same way, we also have FastGreedy version of γ_s based upon the proof of Theorem 2. The FastGreedy version of the DR ratio is an integral part of how the algorithm works and is calculated directly by the algorithm, as we discuss in the next section.

4.2. The FastGreedy Algorithm

The proof of the performance ratio of ThresholdGreedy requires both the submodularity ratio γ_s and the DR ratio γ_d . In this section, we provide an algorithm (FastGreedy, Alg.

3) that achieves ratio $1 - e^{-\beta^* \gamma_s} - \eta$, with factor $\beta^* \geq \gamma_d$ that it can determine during its execution. Appendix E provides proofs for all lemmas, claims, and omitted details.

Description. FastGreedy employs a threshold framework analogous to ThresholdGreedy. Each iteration of the outer **while** loop of FastGreedy is analogous to an iteration of the outer **for** loop in ThresholdGreedy, in which elements are added whose marginal gain exceeds a threshold. FastGreedy employs BinarySearchPivot to find pivots for each $s \in S$ for each threshold value τ . Finally, the parameter ε determines a minimum threshold value.

As its threshold, FastGreedy uses $\tau = \beta \kappa m$, where m is the maximum marginal gain found on line 3, parameter κ is the intended stepsize between thresholds as in ThresholdGreedy, and β is an upper bound on the DR ratio γ_d , as described below. This choice of τ has the following advantages over the approach of ThresholdGreedy: (1) since the threshold is related to the maximum marginal gain m , the theoretical performance ratio is improved; (2) the use of β to lower the threshold ensures the same⁵ worst-case query complexity as ThresholdGreedy and leads to substantial reduction of the number of queries in practice, as we demonstrate in Section 6.

FastGreedy DR ratio β^* . If FastGreedy is modified⁶ to continue until $\|\mathbf{g}\|_1 = k$, let the final, smallest value β^* of β be termed the *FastGreedy DR ratio* on the instance. The FastGreedy DR ratio β^* is at least the DR ratio γ_d of the function, up to the parameter δ :

Lemma 1. Let parameters $\kappa, \delta, \varepsilon \in (0, 1)$ be given. Throughout the execution of FastGreedy on an instance of Problem MCC with $f \in \mathcal{F}_{\mathbf{b}}^{\gamma_d, \gamma_s}$, $\beta \geq \gamma_d \delta$. Since ε can be arbitrarily small, $\beta^* \geq \gamma_d \delta$.

Proof. Initially, $\beta = 1$; it decreases by a factor of $\delta \in (0, 1)$ at most once per iteration of the **while** loop. Suppose $\beta \leq \gamma_d$ for some iteration i of the **while** loop, and let \mathbf{g} have the value assigned immediately after iteration i , m have the value assigned after line 3 of iteration i . Since a valid pivot was found for each $s \in S$ during iteration i , by property (P2) there exists $\mathbf{g}^s \leq \mathbf{g}$, $\delta_s(\mathbf{g}^s) < \beta \kappa m \leq \gamma_d \kappa m$. Hence $\delta_s(\mathbf{g}) \leq \kappa m$, by the definition of DR ratio. In iteration $i + 1$, m' has the value of m from iteration i , so the value of m computed during iteration $i + 1$ is at most $\kappa m'$, and β does not decrease during iteration $i + 1$. \square

Performance ratio. Next, we present the main result of this section. In contrast to ThresholdGreedy, the factor of

⁵Up to a constant factor, which depends on γ_d .

⁶This modification can be accomplished by setting ε to ensure the condition on line 3 is always true on this instance.

Algorithm 3 FastGreedy

```

1: Input:  $f \in \mathcal{F}_{\mathbf{b}}, k \in \mathbb{N}, \kappa, \delta, \varepsilon \in (0, 1)$ .
2: Output:  $\mathbf{g} \in \mathbb{N}^S$ 
3:  $\mathbf{g} \leftarrow 0, M \leftarrow \max_{s \in S} f(\mathbf{s}), m \leftarrow M, m' \leftarrow M/\kappa,$ 
    $\beta \leftarrow 1$ 
4: while  $m \geq M\varepsilon^2/k$  do
5:    $m \leftarrow \max_{s \in S} \delta_s(\mathbf{g})$ 
6:   if  $m > \kappa m'$  then
7:      $\beta \leftarrow \beta\delta$ 
8:      $m' \leftarrow m$ 
9:      $\tau \leftarrow \beta\kappa m$ 
10:  for  $s \in S$  do
11:     $l \leftarrow \text{BinarySearchPivot}(f, \mathbf{g}, \mathbf{b}, s, k, \tau)$ 
12:     $\mathbf{g} \leftarrow \mathbf{g} + l\mathbf{s}$ 
13:    if  $\|\mathbf{g}\|_1 = k$  then
14:      return  $\mathbf{g}$ 
15: return  $\mathbf{g}$ 

```

γ_d in the performance ratio has been replaced with β^* ; at the termination of the algorithm, the value of β^* is an output of FastGreedy if the solution \mathbf{g} satisfies $\|\mathbf{g}\|_1 = k$. In any case, by Lemma 1, the performance ratio is at worst the same as that of ThresholdGreedy.

Theorem 2. *Let an instance of Problem MCC be given, with $f \in \mathcal{F}_{\mathbf{b}}^{\gamma_d, \gamma_s}$. Let \mathbf{g} be the solution returned by FastGreedy with parameters $\kappa, \delta, \varepsilon \in (0, 1)$, and let Ω be an optimal solution to this instance; also, suppose $\gamma_d \geq \varepsilon$. Let β^* be the FastGreedy DR ratio on this instance. Then,*

$$f(\mathbf{g}) \geq \left(1 - e^{-\kappa\beta^*\gamma_s} - \varepsilon\right) f(\Omega)$$

The worst-case query complexity of FastGreedy is $O((\log_{\delta}(\gamma_d) \log_{\kappa}(\gamma_d) + \log_{\kappa} \varepsilon^2/k) n \log k)$.

If $\eta > 0$ is given, the assignment $\kappa = (1 - \eta/2)$, $\varepsilon = \eta/2$ yields performance ratio at least $1 - e^{-\beta^*\gamma_s} - \eta$.

Proof of query complexity. The performance ratio is proved in Appendix E. Let m'_1, \dots, m'_K be the sequence of m' values in the order considered by the algorithm. By Lemma 1, $m'_j > \kappa m'_{j-1}$ at most $\Gamma = \log_{\delta} \gamma_d$ times; label each such index j an *uptick*, and let j_1, \dots, j_l be the indices of each uptick in order of their appearance. Also, let k_i be the first index after j_i such that $m'_{k_i} \leq \kappa m'_{j_i-1}$, for each $i \in \{1, \dots, l\}$.

Next, we will iteratively delete from the sequence of m' values. Initially, let $\ell = l$ be the last uptick in the sequence; delete all terms $m'_{j_\ell}, \dots, m'_{k_\ell-1}$ from the m' sequence. Set $\ell = \ell - 1$ and repeat this process until $\ell = 0$.

Claim 3. For each ℓ selected in the iterative deletion above, there are at most $\log_{\kappa} \gamma_d$ values deleted from the sequence.

By Claim 3 and the bound on the number of upticks, we have deleted at most $\log_{\kappa} \gamma_d \log_{\delta} \gamma_d$ thresholds m' from the sequence; every term in the remaining sequence satisfies $m'_j \leq \kappa m'_{j-1}$; hence, the remaining sequence contains at most $\log_{\kappa} \varepsilon^2/k$ terms, by its initial and terminal values. The query complexity follows from the number of queries per value of m' , which is $O(n \log k)$ by Lemma 2. \square

5. Influence Maximization: A General Framework

In this section, we provide a non-submodular framework for viral marketing on a social network that unifies the classical influence maximization (Kempe et al., 2003) with the boosting problem (Lin et al., 2017).

Overview. The goal of influence maximization is to select seed users (*i.e.* initially activated users) to maximize the expected adoption in the social network, where the total number of seeds is restricted by a budget, such that the expected adoption in the social network is maximized. The boosting problem is, given a fixed seed set S , to incentivize (*i.e.* increase the susceptibility of a user to the influence of his friends) users within a budget such that the expected adoption with seed set S increases the most.

Our framework combines the above two scenarios with a partial incentive: an incentive (say, $x\%$ off the purchase price) increases the probability a user will purchase the product independently and increases the susceptibility of the user to the influence of his friends. Hence, our problem asks how to best allocate the budget between (partially) seeding users and boosting the influence of likely extant seeds. Both the classical influence maximization and the non-submodular boosting problem can be obtained as special cases, as shown in Appendix F.

Our model is related to the formulation of Demaine et al. (2014); however, they employ a submodular threshold-based model, while our model is inherently non-submodular due to the boosting mechanism (Lin et al., 2017). Also, GIM is related to the submodular budgeted allocation problem of Alon et al. (2012), in which the influence of an advertiser increases with the amount of budget allocated; the main difference with GIM is that we modify incoming edge weights with incentives instead of outgoing, which creates the boosting mechanism responsible for the non-submodularity.

Model. Given a social network $G = (V, E)$, and a product p , we define the following model of adoption. The allocation of budget to u is thought of as a discount towards purchasing the product; this discount increases the probability that this user will adopt or purchase the product. Furthermore, this discount increases the susceptibility of the

user to influence from its (incoming) social connections.

Formally, an incentive level \mathbf{x}_u is chosen for each user u . With independent probability $p(u, \mathbf{x}_u)$, user u initially activates or adopts the product; altogether, this creates a probabilistic initial set S of activated users. Next, through the classical Independent Cascade (IC) model⁷ of adoption, users influence their neighbors in the social network; wherein the weight $p(v, u, \mathbf{x}_u)$ for edge (v, u) is determined by the incentive level \mathbf{x}_u of user u as well as the strength of the social connection from v to u .

We write $p^\mathbf{x}(H, T)$ to denote the probability of full graph realization H and seed set T when \mathbf{x} gives the incentive levels for each user. We write $R(H, T)$ to denote the size of the reachable set from T in realization H . The expected activation in the network given a choice \mathbf{x} of incentive levels is given by $\mathbb{I}(\mathbf{x}) = \sum_{T \subseteq V} \sum_{H \subseteq G} p^\mathbf{x}(H, T) R(H, T)$, where an explicit formula for $p^\mathbf{x}(H, T)$ is given in Appendix F. Finally, let $\mathbb{A}(\mathbf{x}) = \mathbb{I}(\mathbf{x}) - \mathbb{I}(\mathbf{0})$.

Definition 4 (Generalized Influence Maximization (GIM)). Let social network $G = (V, E)$ be given, together with the mappings $i \mapsto p(u, i)$, $i \mapsto p(u, v, i)$, for all $u \in V$, $(u, v) \in E$, for each $i \in \{0, \dots, L\}$, where L is the number of incentive levels. Given budget k , determine incentive levels \mathbf{x} , with $\|\mathbf{x}\|_1 \leq k$, such that $\mathbb{A}(\mathbf{x})$ is maximized.

Bound on non-submodularity. Next, we provide a lower bound on the greedy DR ratios (see Appendix B). We emphasize that the assumption that the probability mappings as a function of incentive level be submodular does not imply the objective $\mathbb{A}(\mathbf{x})$ is DR-submodular. Theorem 3 is proved in Appendix F.

Theorem 3. *Let \mathcal{I} be an instance of GIM, with budget k . Let $c_e = \max_{(u,v) \in E, i \in L} \frac{p(u,v,i+1)}{p(u,v,i)}$, $c_n = \max_{x \in V, i \in L} \frac{p(x,i+1)}{p(x,i)}$. Suppose for all $(u,v) \in E, w \in V$, the mappings $i \mapsto p(u, v, i)$, $i \mapsto p(w, i)$ are submodular set functions. Then, the greedy DR ratios defined in Appendix B and the FastGreedy DR ratio are lower bounded by $c_e^{-k\Delta} c_n^{-k}$, where Δ is the maximum in-degree in G .*

6. Experimental Evaluation

In this section, we evaluate our proposed algorithms for the GIM problem defined in Section 5. Source code for the implementation is available at <https://gitlab.com/emallson/lace>. We evaluate our algorithms as compared with StandardGreedy; by the naive reduction of the lattice to sets in exponential time, this algorithm is equivalent to performing this reduction and running the standard greedy for sets, the performance of which for

⁷The IC model is defined in Appendix F.

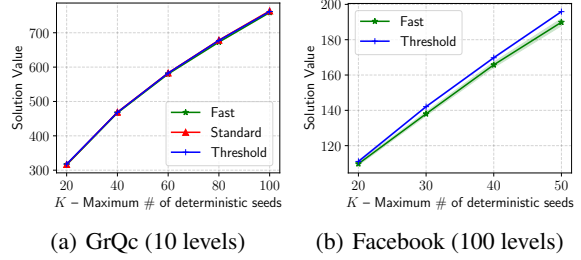


Figure 1. Activation $\mathbb{A}(\mathbf{g})$ for the solution returned by each algorithm.

non-submodular set functions was analyzed by Bian et al. (2017b).

In Section 6.1, we describe our methodology; in Section 6.2, we compare the algorithms and non-submodularity parameters. In Appendix G.1, we explore the behavior of FastGreedy as the parameters δ , κ , and ε are varied.

6.1. Methodology

Our implementation uses Monte Carlo sampling to estimate the objective value $\mathbb{A}(\mathbf{x})$, with 10 000 samples used. As a result, each function query is relatively expensive.

We evaluate on two networks taken from the SNAP dataset (Leskovec & Krevl, 2014): `ca-grqc` (“GrQc”; 15k nodes, 14.5K edges) and `facebook` (“Facebook”; 4k nodes, 176K edges). Unless otherwise specified, we use 10 repetitions per datapoint and display the mean. The width of shaded intervals is one standard deviation. Standard greedy is omitted from some figures where running time is prohibitive. Unless noted otherwise, we use default settings of $\varepsilon = 0.05$, $\delta = 0.9$, $\kappa = 0.95$. We use a uniform box constraint and assign each user the same number of incentive levels; the maximum incentive level for a user corresponds to giving the product to the user for free and hence deterministically seeds the user; we adopt linear models for the mappings $i \mapsto p(u, i)$, $i \mapsto p(u, v, i)$. We often plot versus K , which is defined as the maximum number of deterministic seeds; for example, if $k = 200$ with 10 incentive levels, then $K = 20$.

6.2. Results

In this section, we demonstrate the following: (1) our algorithms exhibit virtually identical quality of solution with StandardGreedy, (2) our algorithms query the function much fewer times, which leads to dramatic runtime improvement over StandardGreedy, (3) FastGreedy further reduces the number of queries of ThresholdGreedy while sacrificing little in solution quality, and (4) the non-submodularity parameters on a small instance are com-

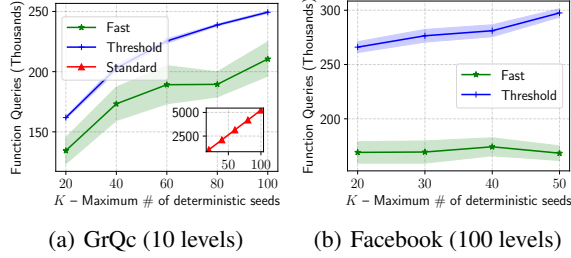


Figure 2. Total function queries on the GrQc and Facebook networks.

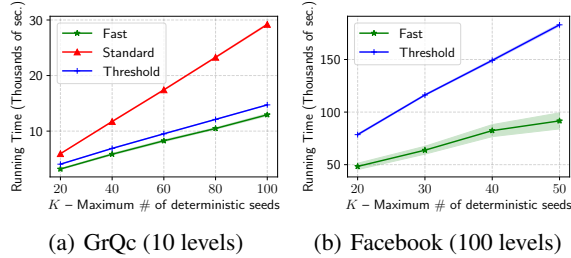


Figure 3. Runtime on the GrQc and Facebook networks with 100 levels.

puted, which provides evidence that our theoretical performance ratios are useful.

Quality of Solution In Fig. 1(a), we plot $\mathbb{A}(g)$ for the solution returned by each algorithm on the GrQc network with 10 incentive levels; the difference in quality of solution returned by the three algorithms is negligible. In Fig. 1(b), we plot the same for the Facebook network with 100 incentive levels; on Facebook, we drop StandardGreedy due to its prohibitive runtime. FastGreedy is observed to lose a small (up to 3%) factor, which we consider acceptable in light of its large runtime improvement, which we discuss next.

Number of Queries Next, we present in Fig. 2 the number of function queries⁸ each algorithm requires on the GrQc and Facebook networks. StandardGreedy required up to 20M queries on Facebook, hence it is not shown in Fig. 2(b). Both of our algorithms provide a large improvement over StandardGreedy; in particular, notice that StandardGreedy increases linearly with k , while both of the others exhibit logarithmic increase in agreement with the theoretical query complexity of each. Furthermore, FastGreedy uses at least 14.5% fewer function queries than ThresholdGreedy and up to 43% fewer as k grows.

⁸Our implementation is in terms of the marginal gain. The number of function queries shown is the number of times the marginal gain function was called.

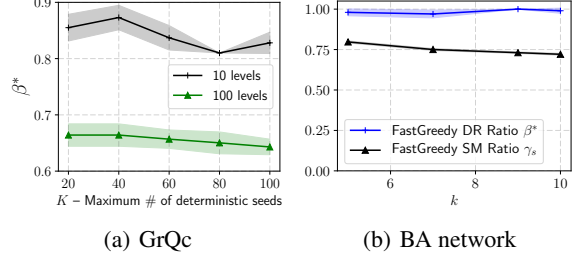


Figure 4. (a): The value of FastGreedy DR ratio β^* on the GrQc dataset. (b): FastGreedy submodularity ratio γ_s and FastGreedy DR Ratio β^* on a small, random BA network.

Non-Submodularity Parameters The value of the FastGreedy DR ratio β^* on GrQc is shown in Fig. 4(a); notice that it is relatively stable as the budget increases from $k = 20$ to 100, although there is substantial drop from 10 incentive levels to 100; this may be explained as an increase in the non-submodularity resulting from inaccurate sampling of \mathbb{A} , since it is more difficult to detect differences between the finer levels. Still, on all instances tested, $\beta^* > 0.6$, which suggests the worst-case performance ratio of FastGreedy is not far from that of StandardGreedy.

Finally, we examine the various non-submodularity parameters on a very small instance which admits their computation: a random Barabasi-Albert network with 10 nodes and 10 incentive levels. We compute the FastGreedy version of the submodularity ratio γ_s defined in Appendix B by direct enumeration and consider the FastGreedy DR ratio β^* . Results are shown in Fig. 4(b). The value of β^* is close to 1 and remains constant with increasing budget k , while the FastGreedy submodularity ratio decreases slowly with k . With β^* and the FastGreedy γ_s , we can compute the worst-case performance ratio of FastGreedy across these instances: 0.449692.

7. Conclusions

In this work, we provide two approximation algorithms for maximizing non-submodular functions with respect to a cardinality constraint on the integer lattice with polynomial query complexity. Since set functions are a special case, our work provides faster algorithms for the same problem with set functions than the standard greedy algorithm, although the performance ratio degrades from at least $1 - e^{-\gamma_s}$ to $1 - e^{-\beta^* \gamma_s}$, where β^* is the FastGreedy DR Ratio. We propose a natural application of non-submodular influence maximization, for which we lower bound the relevant non-submodularity parameters and validate our algorithms.

Acknowledgement

This work was supported in part by US NSF EFRI 1441231, CCF 1422116, and DTRA HDTRA1-14-1-0055.

References

- Alon, Noga, Gamzu, Iftah, and Tennenholtz, Moshe. Optimizing budget allocation among channels and influencers. *Proceedings of the 21st International Conference on World Wide Web (WWW)*, pp. 381–388, 2012.
- Badanidiyuru, Ashwinkumar and Vondrák, J. Fast algorithms for maximizing submodular functions. *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1497–1514, 2014.
- Badanidiyuru, Ashwinkumar, Mirzasoleiman, Baharan, Karbasi, Amin, and Krause, Andreas. Streaming Submodular Maximization: Massive Data Summarization on the Fly. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and Data Mining (KDD)*, pp. 671–680, 2014.
- Bian, An, Levy, Kfir Y., Krause, Andreas, and Buhmann, Joachim M. Non-monotone Continuous DR-submodular Maximization: Structure and Algorithms. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 486–496, 2017a.
- Bian, Andrew An, Buhmann, Joachim M., Krause, Andreas, and Tschintschek, Sebastian. Guarantees for Greedy Maximization of Non-submodular Functions with Applications. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 2017b.
- Bogunovic, Ilija, Zhao, Junyao, and Cevher, Volkan. Robust Maximization of Non-Submodular Objectives. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2018.
- Borodin, Allan, Le, Dai Tri Man, and Ye, Yuli. Weakly Submodular Functions. *arXiv preprint arXiv:1401.6697*, 2014.
- Conforti, Michele and Cornuéjols, Gérard. 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.
- Das, Abhimanyu and Kempe, David. Submodular meets Spectral: Greedy Algorithms for Subset Selection, Sparse Approximation and Dictionary Selection. *Proceedings of the 28th International Conference on Machine Learning (ICML)*, 2011.
- Demaine, Erik D., Hajiaghayi, Mohammad T., Mahini, Hamid, Malec, David L., Raghavan, S., Sawant, Anshul, and Zadimoghadam, Morteza. How to Influence People with Partial Incentives. *Proceedings of the 23rd International Conference on World Wide Web (WWW)*, pp. 937–948, 2014.
- Elenberg, Ethan R., Dimakis, Alexandros G., Feldman, Moran, and Karbasi, Amin. Streaming Weak Submodularity: Interpreting Neural Networks on the Fly. In *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- Ene, Alina and Nguyen, Huy L. A Reduction for Optimizing Lattice Submodular Functions with Diminishing Returns. *arXiv preprint arXiv:1606.08362v1*, 2016.
- Feige, Uriel and Izsak, Rani. Welfare Maximization and the Supermodular Degree. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science (ITCS)*, pp. 247–256, 2013.
- Gottschalk, Corinna and Peis, Britta. Submodular Function Maximization over Distributive and Integer Lattices. *arXiv preprint arXiv:1505:05423*, 2016.
- Horel, Thibaut and Singer, Yaron. Maximization of Approximately Submodular Functions. In *Advances in Neural Information Processing Systems (NIPS)*, 2016.
- Kempe, David, Kleinberg, Jon, and Tardos, Éva. Maximizing the spread of influence through a social network. In *Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pp. 137–146, 2003.
- Krause, Andreas, Leskovec, Jure, Guestrin, Carlos, VanBriesen, Jeanne M., and Faloutsos, Christos. Efficient sensor placement optimization for securing large water distribution networks. *Journal of Water Resources Planning and Management*, 134(6):516–526, 2008a.
- Krause, Andreas, Singh, Ajit, and Guestrin, Carlos. Near-Optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies. In *Journal of Machine Learning Research*, volume 9, pp. 235–284, 2008b.
- Lehmann, Benny, Lehmann, Daniel, and Nisan, Noam. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55(2):270–296, 2006.
- Leskovec, Jure and Krevl, Andrej. SNAP Datasets: Stanford large network dataset collection. <http://snap.stanford.edu/data>, June 2014.

- Leskovec, Jure, Krause, Andreas, Guestrin, Carlos, Faloutsos, Christos, VanBriesen, Jeanne, and Glance, Natalie. Cost-effective Outbreak Detection in Networks. In *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pp. 420–429, 2007.
- Lin, Yishi, Chen, Wei, and Lui, John C.S. Boosting information spread: An algorithmic approach. *Proceedings of the International Conference on Data Engineering (ICDE)*, pp. 883–894, 2017.
- Mirzasoleiman, Baharan, Badanidiyuru, Ashwinkumar, Karbasi, Amin, Vondrak, Jan, and Krause, Andreas. Lazier Than Lazy Greedy. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI)*, pp. 1812–1818, 2015.
- Qian, Chao, Zhang, Yibo, Tang, Ke, and Yao, Xin. On Multiset Selection with Size Constraints. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI)*, pp. 1395–1402, 2018.
- Soma, Tasuku and Yoshida, Yuichi. A Generalization of Submodular Cover via the Diminishing Return Property on the Integer Lattice. *Advances in Neural Information Processing Systems (NIPS)*, 2015.
- Soma, Tasuku and Yoshida, Yuichi. Maximizing Monotone Submodular Functions over the Integer Lattice. In *Integer Programming and Combinatorial Optimization*, pp. 325–336, Cham, 2016. Springer International Publishing.
- Soma, Tasuku and Yoshida, Yuichi. Non-monotone DR-Submodular Function Maximization. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI)*, pp. 898–904, 2017.