Adaptive Budget Allocation for Maximizing Influence of Advertisements

Daisuke Hatano, Takuro Fukunaga, Ken-ichi Kawarabayashi
National Institute of Informatics, Japan
JST, ERATO, Kawarabayashi Large Graph Project, Japan
{hatano, takuro, k_keniti}@nii.ac.jp

Abstract

The budget allocation problem is an optimization problem arising from advertising planning. In the problem, an advertiser has limited budgets to allocate across media, and seeks to optimize the allocation such that the largest fraction of customers can be influenced. It is known that this problem admits a $(1-1/e)$-approximation algorithm. However, no previous studies on this problem considered adjusting the allocation adaptively based upon the effect of the past campaigns, which is a usual strategy in the real setting. Our main contribution in this paper is to analyze adaptive strategies for the budget allocation problem. We define a greedy strategy, referred to as the insensitive policy, and then give a provable performance guarantee. This result is obtained by extending the adaptive submodularity, which is a concept studied in the context of active learning and stochastic optimization, to the functions over an integer lattice.

1 Introduction

Suppose an advertiser wishes to maximize the influence on customers, but has limited budgets to allocate across media (e.g., webpages, television, or newspapers). The main question, called the *budget allocation problem*, is how to select media considering budget constraints such that the largest fraction of customers can be influenced; that is, how can a budget achieve the maximum reach?

The main difficulty behind the budget allocation problem is the complex dynamic of influence from media to customers. This dynamic has been investigated in the framework of the *influence maximization problem*, which was first introduced by Domingos and Richardson [2001; 2002]. A seminal work by Kempe, Kleinberg, and Tardos [2003] formulated the influence maximization problem in the framework of *submodularity*. The submodularity concept represents a certain *diminishing marginal return property* in discrete settings. Kempe et al. showed that the expected number of customers influenced by media is represented by a submodular set function. On the basis of this observation, they proved an approximation guarantee of a polynomial-time greedy algorithm for the influence maximization problem.

These studies on the influence maximization problem motivated the work of Alon, Gamzu, and Tennenholtz [2012], who formulated the budget allocation problem in the bipartite influence model as another combinatorial optimization problem and provided a provable approximation algorithm. There was difficulty in expressing their problem setting using submodular functions because submodularity is usually defined for combinations of objects whereas budget allocations are assignments of budgets to media. However, Soma et al. [2014] showed that the problem setting of Alon et al. can also be expressed in the framework of submodularity. They utilized submodularity functions over an integer lattice, which are more general than submodular set functions.

Despite these developments, the previous studies on the budget allocation problem have a crucial limitation. In their settings, advertisers have to assign their entire budget at once at the beginning of the process. However, in reality, advertisers routinely adjust their strategy when they see changes in the dynamic or when something unexpected happens. For example, in the US presidential campaign of 2012, both Obama and Romney spent half a billion dollars for TV ads [The Washington Post, 2012]. In particular, they invested huge amounts of money in "swing" states. For these states, both campaigns changed their strategy for TV ads every day, according to their polls (i.e., either gaining momentum or not). In this case, momentum changed frequently, and the dynamic was a deciding factor in their strategy. Hence, both the campaigns changed their strategy *adaptively* every day. In this paper, we are motivated by this observation. We aim to consider adaptive strategies to address the budget allocation problem.

Adaptivity has been already considered in the framework of submodularity. Golovin and Krause [2011b] defined a concept of *adaptive submodularity*, and showed that a greedy adaptive algorithm has a theoretical approximation guarantee if the objective function is adaptive monotone submodular. After their initial work, numerous studies further investigated algorithms for optimization problems with adaptive submodular functions [Golovin and Krause, 2011a; Gabillon *et al.*, 2013; 2014; Gotovos *et al.*, 2015], as well as their applications [Golovin *et al.*, 2010; Chen and Krause, 2013; Chen *et al.*, 2014; Deshpande *et al.*, 2014; Krause *et al.*, 2014; Chen *et al.*, 2015]. Because the model of Golovin and Krause contains the adaptive setting of Kempe et al. [2003], adaptive strategies have been already analyzed in the influence maximization problem. However, adaptive strategies for the budget allocation problem have not been captured by their model. Thus, we need to formulate a new concept of adaptive submodularity that can model the budget allocation problem.

1.1 Contributions

In this paper, we consider adaptive strategies in the budget allocation problem in the bipartite influence model introduced by Alon et al. [2012]. To this end, we define adaptive submodularity of functions over integer lattices, which is a new concept that extends both the adaptive submodularity given by Golovin and Krause [2011b] and submodularity over integer lattice used in Soma et al. [2014] (see Section 3.3). This concept captures the objective function in the adaptive version of the budget allocation problem. Hence we obtain a good adaptive strategy for the budget allocation problem by designing a strategy for maximizing an adaptive submodular function over an integer lattice.

In many variants of the submodular maximization problems, the greedy algorithms achieve good performance both in practice and in theory. Thus we analyze the performance of greedy adaptive algorithms for maximizing adaptive monotone submodular functions over integer lattices. For our problem, a greedy strategy repeats allocating a certain amount of budget to a medium so that the increase of the influence per allocated budget is maximized. In our setting, the strategy is given a new feedback when it allocates a unit amount of budget. It is natural to update the strategy each time a new feedback is given. We call such a strategy *sensitive* greedy strategy. On the other hand, an *insensitive* greedy strategy ignores feedbacks until a certain proportion of a budget has been allocated to a media. Surprisingly, we can show both theoretically and empirically that several typical sensitive greedy strategies are inferior even to the non-adaptive algorithms. Our proposal algorithms are sort of insensitive greedy strategies.

More specifically, we present the following two variations of the insensitive greedy algorithms.

- An algorithm outputs a budget allocation that achieves $(1-1/e)$-approximation. That is, its expected objective value is at least $(1-1/e)$ times that achieved by arbitrary adaptive algorithms. The allocation may violate the budget constraints by a factor of at most two; however, its expected cost is at most the given budget upper limit. (see Theorem 3)
- Another algorithm outputs a budget allocation of an approximation ratio $(e-1)/(2e)$. The allocation is guaranteed to satisfy the budget constraints. (see Theorem 4)

Alon et al. [2012] showed that a non-adaptive greedy algorithm achieves $(1-1/e)$-approximation for the budget allocation problem. We note that our guarantee on the first insensitive policy is superior to the one obtained by Alon et al. although their approximation ratios match. This is because our algorithm is compared with adaptive algorithms whereas Alon et al. compared only non-adaptive algorithms. So if the optimal adaptive algorithm is strictly better than the non-adaptive one (which is quite often the case), then our guarantee is better. Indeed, there is an instance in which our algorithms are better than any non-adaptive algorithms by more than 58%; see at the end of Section 4.

Let us explain why the former variation of our insensitive greedy algorithm violates the budget constraints. Our budget constraint corresponds to *knapsack constraints* in the submodular function maximization. For the non-adaptive setting of maximizing submodular functions subject to the knapsack constraints, the $(1-1/e)$-approximation is achieved only by combining a greedy algorithm with a partial enumeration of solutions in the initial step. However, in the adaptive setting, this partial enumeration is not permitted; hence, it is difficult to achieve $(1-1/e)$-approximation. Therefore, we propose violating the budget constraints by a factor of at most two. The similar approach was adopted in Golovin and Krause [2011b] for the adaptive maximization of submodular set functions.

1.2 Organization

The rest of this paper is organized as follows. Section 2 introduces the budget allocation problem in the bipartite influence model and the submodular functions over an integer lattice. Section 3 formulates our problem setting and defines adaptive submodularity over an integer lattice. Section 4 analyzes the adaptive greedy algorithms. Section 5 compares performance of the algorithms through computational experiments. Section 6 concludes the paper.

2 Budget allocation problem and submodular functions

In this section, we introduce the budget allocation problem proposed by Alon et al. [2012], and slightly extended by Some et al. [2014]. We also explain a relationship with the submodular functions over an integer lattice.

2.1 Bipartite influence model of the budget allocation problem

Let $\mathbb{Z}_+$ and $\mathbb{R}_+$ be the sets of non-negative integers and real numbers, respectively. For a finite set V, let $\mathbb{Z}_+^V$ denote the set of non-negative integer vectors, where each component is indexed by an element in V. For vectors $x, y \in \mathbb{Z}_+^V$, we write $x \leq y$ if $x(v) \leq y(v)$ for all $v \in V$. For an integer $i \in \mathbb{Z}_+$, let $[i]$ denote $\{0, 1, \ldots, i\}$.

We consider a bipartite graph $(V, U; E)$, where V is the set of media, U is the set of customers, and E is the set of edges between V and U. We are given a budget $k \in \mathbb{R}_+$, a cost function $c\colon V \times \mathbb{Z}_+ \to \mathbb{R}_+$, and a vector $b \in \mathbb{Z}_+^V$ representing the numbers of slots of each media. In addition, for each edge $vu \in E$ that joins nodes $v \in V$ and $u \in U$, we are given a probability function $q_{vu}\colon [b(v)] \to [0,1]$. We assume that a media v has $b(v)$ slots in total. It costs $\sum_{j \in [i]} c(v, j)$ to buy i slots of v. We allocate a total budget of k to the media. For notational convenience, we let $c(x)$ denote $\sum_{v \in V} \sum_{j \in [x(v)]} c(v, j)$ for $x \in \mathbb{Z}_+^V$. The allocation can be represented by a vector $x \in \mathbb{Z}_+^V$ such that $x \leq b$ and $c(x) \leq k$; if $x(v) = i$, it represents that we buy i slots of media v.

Let $N(v)$ denote the set of neighbors of a node v in the bipartite graph. If $x(v)$ slots of media v are bought, v attempts to influence each customer $u \in N(v)$ $x(v)$ times. For $i \in [b(v)]$ and $u \in N(v)$, $q_{vu}(i)$ represents the probability that the i-th trial of media v to influence customer u succeeds. Here, we assume that each trial is independent. $g(x)$ is defined as the expected number of influenced customers when the budget allocation is x. That is,

$$g(x) := \sum_{u \in U} \left(1 - \prod_{v \in N(u)} \prod_{i=1}^{x(v)} (1 - q_{vu}(i)) \right). \quad (1)$$

The budget allocation problem in the bipartite influence model seeks an allocation $x \in \mathbb{Z}_+^V$ that maximizes $g(x)$ subject to $x \leq b$ and $c(x) \leq k$.

Remark 1. *In [Alon* et al.*, 2012], Alon et al. called this model by the source-side influence model to distinguish from another model they called the target-side influence model. Since we consider only the source-side influence model, we simply call it by the bipartite influence model. In the original definition, the costs for buying slots are not considered, i.e.,* $c(v, i) = 1$ *for all* $v \in V$ *and* $i \in [b(v)]$*. Thus our definition is more general than the original one.*

2.2 Submodular functions over an integer lattice

For two vectors $x, y \in \mathbb{Z}_+^V$, let $x \vee y$ denote the vector in $\mathbb{Z}_+^V$ defined by $(x \vee y)(v) = \max\{x(v), y(v)\}$ for $v \in V$, and $x \wedge y$ denote the vector in $\mathbb{Z}_+^V$ defined by $(x \wedge y)(v) = \min\{x(v), y(v)\}$ for $v \in V$. Let $f : \mathbb{Z}_+^V \to \mathbb{R}_+$ be a function defined on the integer lattice where each component corresponds to an element in V. The function f is *submodular* (over an integer lattice) if

$$f(x) + f(y) \geq f(x \wedge y) + f(x \vee y) \text{ for all } x, y \in \mathbb{Z}_+^V. \quad (2)$$

f is considered *monotone* if $f(x) \leq f(y)$ for any $x, y \in \mathbb{Z}_+^V$ with $x \leq y$.

For a finite set V, a set function $h : 2^V \to \mathbb{R}_+$ is called submodular if

$$h(X) + h(Y) \geq h(X \cap Y) + h(X \cup Y) \text{ for all } X, Y \in 2^V. \quad (3)$$

When x and y are restricted to vectors over a Boolean lattice (i.e., 2^V), the condition (2) is equivalent to (3) . Hence the submodularity over an integer lattice includes the concept of the submodularity for set functions.

Soma et al. [2014] showed that the budget allocation problem in the bipartite influence model can be captured by the submodular functions over an integer lattice. More concretely, they proved the following theorem.

Theorem 1 ([Soma *et al.*, 2014])**.** *The function g defined by* (1) *is monotone submodular over an integer lattice.*

For the set functions, it is known that h satisfies condition (3) if and only if $h(X \cup \{v\}) - h(X) \geq h(Y \cup \{v\}) - h(Y)$ for any $X, Y \in 2^V$ with $X \subseteq Y$ and $v \in V \setminus Y$. This property is known as the decreasing marginal gain property of submodular set functions. Soma et al. also showed that this property is extended to the monotone submodular functions over an integer lattice, which is summarized in the theorem below. For $v \in V$, let χ_v denote the vector in $\mathbb{Z}_+^V$ such that $\chi_v(v) = 1$ and $\chi_v(u) = 0$ for each $u \in V \setminus \{v\}$.

Theorem 2 ([Soma *et al.*, 2014])**.** *If* $f : \mathbb{Z}_+^V \to \mathbb{R}_+$ *is monotone submodular, then it satisfies* $f(x \vee k\chi_v) - f(x) \geq f(y \vee k\chi_v) - f(y)$ *for any* $k \in \mathbb{Z}_+$*,* $v \in V$*, and* $x, y \in \mathbb{Z}_+^V$ *with* $x \leq y$*.*

We note that the monotone submodular function f over an integer lattice does not always satisfy the *component-wise convexity* represented by $f(x + \chi_v) - f(x) \geq f(x + 2\chi_v) - f(x + \chi_v)$ for $x \in \mathbb{Z}_+^V$ and $v \in V$. Therefore, in contrast with submodular set functions, submodular functions over an integer lattice are not considered convex in the direction of χ_v.

By Theorem 1, the maximization problem of f extends the budget allocation problem in the bipartite influence model. In the submodular maximization problem, we are given a monotone submodular function $f : \mathbb{Z}_+^V \to \mathbb{R}_+$, and seek to find a solution $x \in \mathbb{Z}_+^V$ that maximizes $f(x)$. When the problem demands a knapsack constraint, $b \in \mathbb{Z}_+^V$, $k \in \mathbb{Z}_+$, and $c : V \times \mathbb{Z}_+ \to \mathbb{R}_+$ are given as inputs, and a solution $x \in \mathbb{Z}_+^V$ must satisfy $x \leq b$ and $c(x) \leq k$. The constraint is called a cardinality constraint if $c(v, i) = 1$ for all $v \in V$ and $i \in \mathbb{Z}_+$.

Soma et al. proposed a $(1 - 1/e)$-approximation algorithm for the submodular maximization problem with the knapsack constraints. If f is the function g defined by (1), the problem coincides with the budget allocation problem in the bipartite influence model. Hence the result of Soma et al. extends the algorithm given by Alon et al. [2012].

3 Adaptive submodularity over an integer lattice

Our main contribution in this paper is to analyze the adaptive strategies in the budget allocation problem. In this section, we first define the adaptive setting of the budget allocation problem in the bipartite influence model. Then, we extend it to the submodular maximization problem.

3.1 Adaptive setting of the bipartite influence model

In the adaptive setting, the inputs of the problem are same as the non-adaptive setting of the bipartite influence model; namely, a bipartite graph $(V, U; E)$, $b \in \mathbb{Z}_+^V$, $k \in \mathbb{R}_+$, $c : V \times \mathbb{Z}_+ \to \mathbb{R}_+$, and $q_{vu} : [b(v)] \to [0, 1]$ for each $vu \in E$.

A budget allocation $x \in \mathbb{Z}_+^V$ is initialized to the all-zero vector. Then, we allocate the total budget k to the media sequentially. If we buy one slot of a media $v \in V$ when $x(v) = i$, $x(v)$ is increased to $i + 1$, and the media v activates each customer $u \in N(v)$ in probability $q_{vu}(i + 1)$. In the adaptive setting, we can observe which customers in $N(v)$ got influenced by this trial immediately after increasing $x(v)$ from i to $i + 1$, and we can change the behavior in the subsequent steps based on the observation. Thus, our aim is to find a good *policy*, which describes how we behave for each observation.

This setting is natural in the marketing. When an advertising campaign is committed, the advertiser can observe how many customers are influenced (e.g., buy a product, subscribe to a service), and profile of the influenced customers can be easily obtained in most cases. The advertiser changes the strategies based on the observation. Particularly, this applies well to the internet advertising, wherein real-time marketing is widely used.

3.2 Adaptive setting of the maximization problem over an integer lattice

For the sake of generality, we discuss the adaptive policies in the budget allocation problem in the framework of the submodular functions. We extend the budget allocation problem to a maximization problem of a stochastic objective function over an integer lattice subject to a knapsack constraint. We then define the adaptive monotonicity and the adaptive submodularity of a stochastic objective function, and prove that the objective function defined from the budget allocation problem satisfies these properties.

First, let us introduce the problem setting in the general framework. Let R be a set of random variables. The range of variables in R is denoted by S. We call the value of a variable $r \in R$ by the *state* of r. For $r \in R$ and $s \in S$, let $p_r(s)$ denote the probability that the variable r is in the state s. In the problem, the states of the random variables are not observed in advance. Initially, the available information is the probabilities $p_r(s)$ for all $r \in R$ and $s \in S$.

We represent the states of all random variables in R by a function $\phi\colon R \to S$, which we refer to as *full realization*. The objective function depends on the full realization. Let $f_\phi\colon \mathbb{Z}_+^V \to \mathbb{R}_+$ denote the objective function when the realization is $\phi\colon R \to S$. Our goal is to find a *policy* that adaptively computes a solution $x \in \mathbb{Z}_+^V$ maximizing $f_\phi(x)$ under given constraints.

If a policy sets a solution to $x \in \mathbb{Z}_+^V$, it observes the state of random variables in a subset $O_{\phi,x}$ of S. Note that $O_{\phi,x}$ depends on the full realization ϕ and the solution x. $O_{\phi,x}$ is monotone with respect to x; namely, if $x, y \in \mathbb{Z}_+^V$ satisfy $x \le y$, then $O_{\phi,x} \subseteq O_{\phi,y}$ for any $\phi\colon R \to S$. The value of $f_\phi(x)$ depends only on the states of variables in $O_{\phi,x}$.

In the problem, the solution x is initialized by $x(v) := 0$ for all $v \in V$, and the policy repeats increasing a component of x by one while satisfying the constraints; i.e., it is prohibited to decrease x, and x must always satisfy $c(x) \le k$ and $x \le b$ for a given cost function $c\colon V \times \mathbb{Z}_+ \to \mathbb{R}_+$, a budget $k \in \mathbb{R}_+$, and the numbers of slots $b \in \mathbb{Z}_+^V$. When $x(v)$ is increased from $i-1$ to i, the policy observes the states of the variables in $O_{\phi,x} \setminus O_{\phi,x'}$, where x (resp., x') denotes the vector after (resp., before) the increase. The behavior of the policy in the subsequent steps depends on the observation.

Here, let π be a policy. We denote by $x_{\phi,\pi}$ the vector output by the policy π when the states of the random variables in R are represented by $\phi\colon R \to S$. Let $f_{\rm avg}(\pi) = E[f_\phi(x_{\phi,\pi})]$, where the expectation depends on the randomness of the variables in R. Hence $f_{\rm avg}(\pi)$ denotes the expected value of the objective function obtained by running the policy π. We may consider a randomized policy as π, and, in this case, the expectation also depends on the randomness of π. We measure the performance of π by $f_{\rm avg}(\pi)$.

In the budget allocation problem with the bipartite influence model, there is a random variable for each pair of $vu \in E$ and $i \in \{1, \ldots, b(v)\}$. Hence let R denote $\{(vu, i)\colon vu \in E, i \in \{1, \ldots, b(v)\}\}$ by abusing the notation. The state of each pair $(vu, i) \in R$ represents that the i-th trial of v to influence u succeeds or not. Let $S := \{\top, \bot\}$, where $\top$ and $\bot$ respectively denote "success" and "failure." Recall that we are given a probability function $q_{vu}\colon \mathbb{Z}_+ \to [0, 1]$ for each $vu \in E$ in the budget allocation problem. The state of (vu, i) is $\top$ in probability $q_{vu}(i)$. Let $\phi\colon R \to S$ be a full realization. The probability that ϕ occurs is

$$p(\phi) := \prod_{\substack{(vu,i)\in R\\ \phi(vu,i)=\top}} q_{vu}(i) \prod_{\substack{(v'u',i')\in R\\ \phi(v'u',i')=\bot}} (1 - q_{v'u'}(i')).$$

When the budget allocation is $x \in \mathbb{Z}_+^V$, we can observe the states of (vu, i) for all $vu \in E$ and $i \in \{1, \ldots, x(v)\}$; i.e., $O_{\phi,x} = \{(vu, i)\colon vu \in E, i \in \{1, \ldots, x(v)\}\}$. Notice that $O_{\phi,x}$ is monotone with respect to x. The number of influenced customers when the budget allocation is $x \in \mathbb{Z}_+^V$ is represented by

$$g_\phi(x) = |\{u \in U\colon \exists (vu, i) \in O_{\phi,x}, \phi(vu, i) = \top\}|. \quad (4)$$

Maximizing the number of influenced customers adaptively is equivalent to maximizing g_ϕ.

3.3 Adaptive monotonicity and adaptive submodularity over an integer lattice

In this subsection, we define the adaptive monotonicity and the adaptive submodularity over an integer lattice.

When the full realization is $\phi\colon R \to S$ and a policy sets the solution to $x \in \mathbb{Z}_+^V$, it observes the states of all random variables in $O_{\phi,x}$. Let $*$ represent the fact that the state of a random variable is not yet observed, and define $S^* := S \cup \{*\}$. We represent the observation available to the policy by a function $\psi\colon R \to S^*$ such that for each $\rho \in R$, $\psi(\rho)$ is the state of ρ if it is already observed, and $\psi(\rho) = *$ if it is not yet observed. We refer to such a function ψ as *(partial) realization*. The *domain* of a partial realization ψ, denoted by $\mathrm{dom}(\psi)$, indicates $\{\rho \in R\colon \psi(\rho) \ne *\}$. In other words, if ψ represents the observation when the full realization is $\phi\colon R \to S$ and the solution is $x \in \mathbb{Z}_+^V$, then $\mathrm{dom}(\psi) = O_{\phi,x}$.

Let Φ^* and Φ denote the set of all partial realizations and the set of all full realizations, respectively. We say that a realization ψ *extends* another realization ψ' if $\psi'(\rho) = \psi(\rho)$ for each $\rho \in \mathrm{dom}(\psi')$. If ψ extends ψ', we use the notation $\psi \sim \psi'$. Recall that a full realization ϕ happens in probability $p(\phi)$. If a realization $\psi \in \Phi^*$ is not full, we assume that ψ happens in probability $\sum\{p(\phi)\colon \phi \in \Phi, \phi \sim \psi\}$. We denote this probability by $p(\psi)$.

For a vector $x \in \mathbb{Z}_+^V$, let Φ_x^* denote $\{\psi \in \Phi^*\colon \exists \phi \in \Phi, \mathrm{dom}(\psi) = O_{\phi,x}\}$. For $x \in \mathbb{Z}_+^V$, $\psi \in \Phi_x^*$, $v \in V$, and $i \in [b(v)]$, define

$$\Delta(v, i \mid x, \psi) := E\left[f_\phi(x \vee i\chi_v) - f_\phi(x) \mid \phi \in \Phi, \phi \sim \psi\right].$$

In other words, $\Delta(v, i \mid x, \psi)$ is the expected gain we obtain by increasing $x(v)$ to i, conditioned that the current realization is ψ and the current solution is x.

Definition 1 (adaptive monotonicity)**.** *$f := \{f_\phi \colon \phi \in \Phi\}$ is* adaptive monotone *(with respect to distribution $p(\phi)$, $\phi \in \Phi$) if $\Delta(v, i \mid x, \psi) \geq 0$ holds for any $v \in V$, $i \in [b(v)]$, $x \in \mathbb{Z}_+^V$, and $\psi \in \Phi_x^*$ with $p(\psi) > 0$.*

Definition 2 (adaptive submodularity)**.** *$f := \{f_\phi \colon \phi \in \Phi\}$ is* adaptive submodular *(with respect to distribution $p(\phi)$, $\phi \in \Phi$) if $\Delta(v, i \mid x, \phi) \leq \Delta(v, i \mid y, \psi)$ holds for any $v \in V$, $i \in [b(v)]$, $x, y \in \mathbb{Z}_+^V$ with $x \geq y$, and $\phi \in \Phi_x^*$ and $\psi \in \Phi_y^*$ such that $\phi \sim \psi$ and $p(\phi) > 0$.*

It is not hard to observe that the objective function $g := \{g_\phi \colon \phi \in \Phi\}$ defined by (4) in the bipartite influence model is adaptive monotone submodular.

4 Adaptive greedy policies

As mentioned in Section 1, integer lattice setting introduces two types of greedy policies: sensitive and insensitive policies. In this section, we present two variations of insensitive policies. One achieves approximation factor $(1-1/e)$, which matches the best ratio for the non-adaptive setting. However, it may violate the knapsack constraint $c(x) \leq k$. We show that the vector x output by the policy always satisfies $c(x) \leq 2k$, and $E[c(x)] \leq k$. The other policy always outputs a feasible solution. However, its approximation factor is $(e-1)/(2e)$.

First, let us prove several preparatory lemmas. For two policies π and π', their *concatenate* $\pi @ \pi'$ is defined as follows. First, we run π to obtain $x_\pi \in \mathbb{Z}_+^V$. Then, we run π' from a fresh start to obtain $x_{\pi'} \in \mathbb{Z}_+^V$, ignoring the information from the observation during the run of π. $\pi @ \pi'$ outputs $x_\pi \vee x_{\pi'}$.

Lemma 1. *Function $f := \{f_\phi \mid \phi \in \Phi\}$ is adaptive monotone if and only if $f_{\rm avg}(\pi) \leq f_{\rm avg}(\pi' @ \pi)$ for all policies π and π'.*

Proof. Note that $\pi @ \pi'$ always outputs the same solutions as $\pi' @ \pi$ for any full realizations. Thus, $f_{\rm avg}(\pi @ \pi') = f_{\rm avg}(\pi' @ \pi)$.

Suppose that f is adaptive monotone. Assume that a vector $x \in \mathbb{Z}_+^V$ and a realization $\psi \in \Phi_x^*$ appears in probability $w(x, \psi)$ as a solution and a realization kept by the π'-portion of $\pi @ \pi'$ in a certain moment. If $\pi @ \pi'$ increases $x(v)$ to i for ψ, then the increase of the objective function is expected to be $\Delta(v, i \mid x, \psi)$. Therefore, $f_{\rm avg}(\pi @ \pi') - f_{\rm avg}(\pi)$ is expressed by $\sum_{x \in \mathbb{Z}_+^V} \sum_{\psi \in \Phi_x^*} w(x, \psi) \Delta(v, i \mid x, \psi)$. Since f is adaptive monotone, $\Delta(v, i \mid x, \psi) \geq 0$ holds for any x and ψ with $w(x, \psi) > 0$. Hence $f_{\rm avg}(\pi @ \pi') \geq f_{\rm avg}(\pi)$.

Suppose that $\Delta(v, i \mid x, \psi) < 0$ for some $(v, i) \in V \times \mathbb{Z}_+$, $x \in \mathbb{Z}_+^V$, and $\psi \in \Phi_x^*$ with $p(\psi) > 0$. We define policies π and π' as follows. Let y denote the vector kept by the policies. Both policies first increase $y(u)$ from 0 to $x(u)$ for each $u \in V$. If the observed state of some random variable $\rho \in R$ differs from $\psi(\rho)$, both policies terminate. If they succeed to increase all $u \in V$, π terminates, and π' terminates after increasing $y(v)$ to i. Then, we have

$$
\begin{aligned}
&f_{\rm avg}(\pi @ \pi') - f_{\rm avg}(\pi) \\
&\quad = \sum_{\substack{\phi \in \Phi \\ \phi \sim \psi}} E[f_\phi(\pi @ \pi') - f_\phi(\pi)] p(\phi) \\
&\quad = p(\psi) \sum_{\substack{\phi \in \Phi \\ \phi \sim \psi}} E[f_\phi(x \vee i\chi_v) - f_\phi(x)] p(\phi \mid \psi) \\
&\quad = p(\psi) \Delta(v, i \mid x, \psi) < 0.
\end{aligned}
$$

□

Recall that $x_{\phi,\pi}$ denotes the vector output by a policy π for a full realization ϕ. For $x \in \mathbb{Z}_+^V$, $\psi \in \Phi_x^*$ and a policy π, let

$$
\Delta(\pi \mid x, \psi) := E\left[f_\phi(x \vee x_{\phi,\pi}) - f_\phi(x) \mid \phi \in \Phi, \phi \sim \psi\right].
$$

In other words, $\Delta(\pi \mid x, \psi)$ is the expected gain we obtain when, after selecting vector x and observing ψ, we run policy π ignoring the information given from ψ. Note that the expectation depends on the randomness of realizations, and is conditioned on ψ being observed after selecting x. When π is a randomized policy, the expectation also depends on the randomness of π.

Lemma 2. *Let f be an adaptive monotone submodular function. Let $x \in \mathbb{Z}_+^V$, and $\psi \in \Phi_x^*$ be a realization with $p(\psi) > 0$, and let $\pi^\star$ be an arbitrary policy. Then,*

$$
\Delta(\pi^\star \mid x, \psi) \leq E[c(x_{\phi,\pi^\star}) \mid \phi \sim \psi] \max_{(v,i) \in V \times \mathbb{Z}_+} \frac{\Delta(v, i \mid x, \psi)}{\sum_{j \in [i]} c(v, j)}.
$$

Proof. We define a policy π as follows. Let x' be a solution vector kept by π. Starting from $x' \equiv 0$, π increases $x'(v)$ from 0 to $x(v)$ for all $v \in V$, and terminates if the observed state of a random variable ρ differs from $\psi(\rho)$. If it does not terminate after observing all random variables, it proceeds to run $\pi^\star$, while ignoring all information obtained up to this point. Note that π proceeds to run $\pi^\star$ in probability $p(\psi)$.

For each $(v, i) \in V \times \mathbb{Z}_+$, let $w(v, i)$ be the probability that π terminates with $x'(v) = i$ under the condition that the full realization extends ψ. Note that if this event happens, the $\pi^\star$-portion of π increases $x'(v)$ to i. The contribution of this operation to $\Delta(\pi^\star \mid x, \psi)$ is at most $\Delta(v, i \mid x, \psi)$ by the adaptive submodularity. Therefore, $\Delta(\pi^\star \mid x, \psi) \leq \sum_{(v,i) \in V \times \mathbb{Z}_+} w(v, i) \Delta(v, i \mid x, \psi)$. Note that $w(v, i) \in [0, 1]$ for each $(v, i) \in V \times \mathbb{Z}_+$, and $E[c(x_{\phi,\pi^\star}) \mid \phi \sim \psi] = \sum_{(v,i) \in V \times \mathbb{Z}_+} w(v, i) \sum_{j \in [i]} c(v, j)$. Therefore,

$$
\begin{aligned}
&\Delta(\pi^\star \mid x, \psi) \\
&\quad \leq \sum_{(v,i) \in V \times \mathbb{Z}_+} w(v, i) \Delta(v, i \mid x, \psi) \\
&\quad = \sum_{(v,i) \in V \times \mathbb{Z}_+} w(v, i) \Big(\sum_{j \in [i]} c(v, j)\Big) \frac{\Delta(v, i \mid x, \psi)}{\sum_{j \in [i]} c(v, j)} \\
&\quad \leq E[c(x_{\phi,\pi^\star}) \mid \phi \sim \psi] \max_{(v',i') \in V \times \mathbb{Z}_+} \frac{\Delta(v', i' \mid x, \psi)}{\sum_{j \in [i']} c(v', j)}.
\end{aligned}
$$

□

For a deterministic policy π and $i \in [k]$, let π_i denote the *truncation* of π defined as follows. Fixing a full realization ϕ, we define how π_i behaves for ϕ. Let x be the temporal solution kept by π during its run for ϕ. Consider the moment θ when $c(x)$ exceeds i. Suppose that π is increasing $x(v)$ at moment θ. Let θ_0 be the latest moment before θ at which π increases a component of x other than $x(v)$. If there is no such moment, θ_0 denotes the moment at which π begins to run. Similarly, let θ_1 be the earliest moment after θ at which π increases a component of x other than $x(v)$. If there is no such moment, θ_1 denotes the moment at which π terminates. Suppose that $c(x) = i_0$ at θ_0, and $x(v) = j_1$ and $c(x) = i_1$ at θ_1. Until θ_0, π_i behaves as π. Then, in probability $(i - i_0)/(i_1 - i_0)$, π_i increases $x(v)$ to j_1 and terminates. Otherwise, π_i terminates without increasing $x(v)$. Note that the truncation π_i outputs x such that $E[c(x)] \leq i$ for any full realization ϕ, where the expectation is over only the randomness of π_i.

We now define a policy π as follows. We assume without loss of generality that $c(b(v)\chi_v) \leq k$ holds for all $v \in V$ in the rest of this section. Starting from $x \equiv 0$, π chooses $(v, i) \in V \times \mathbb{Z}_+$ that maximizes $\Delta(v, i \mid x, \psi)/(\sum_{j \in [i]} c(v, j))$ and increases $x(v)$ to i, where ψ is the current realization. π repeats this procedure and terminates when $c(x) \geq k$ holds. Our first proposal algorithm is its truncation π_k. We describe the details of π_k in Policy 1. Notice that the behavior of π_k depends on the observation in Step 8, and hence it is an adaptive policy.

Policy 1 Bicriteria $(1 - 1/e)$-Approximation Policy

Input: a finite set V, an adaptive monotone submodular function $\{f_\phi \colon \mathbb{Z}_+^V \to \mathbb{R}_+\}_{\phi \in \Phi}$, $k \in \mathbb{R}_+$, $b \in \mathbb{Z}_+^V$, and $c \colon V \times \mathbb{Z}_+ \to \mathbb{R}_+$
Output: $x \in \mathbb{Z}_+^V$ such that $c(x) \leq 2k$ and $x \leq b$

1: $x(v) \longleftarrow 0$ for each $v \in V$
2: $\psi(\rho) \longleftarrow *$ for each $\rho \in R$
3: **while** $c(x) < k$ **do**
4: $\quad (v, i) \longleftarrow \arg\max \Delta(v, i \mid x, \psi)/\sum_{j \in [i]} c(v, j)$ where the maximization is over all $(v, i) \in V \times \mathbb{Z}_+$ with $i \leq b(v)$
5: $\quad C \longleftarrow \sum_{j=x(v)+1}^{i} c(v, j)$
6: $\quad$ If $c(x) + C > k$, output x and terminate in probability $1 - (k - c(x))/C$
7: $\quad x(v) \longleftarrow \max\{x(v), i\}$
8: $\quad \psi(\rho) \longleftarrow$ the state of ρ for each observed random variable $\rho \in R$
9: **end while**
10: output x

We consider a policy *feasible* if it outputs a vector x with $E[c(x)] \leq k$ for any full realization, where the expectation is over the inner randomness of the policy. The following theorem presents the $(1 - 1/e)$-approximation guarantee of Policy 1.

Theorem 3. *Let π_k be the policy presented in Policy 1. If f is adaptive monotone submodular, then $f_{\mathrm{avg}}(\pi_k) \geq (1 - 1/e) f_{\mathrm{avg}}(\pi^\star)$ holds for any feasible policy $\pi^\star$. Moreover, π_k is feasible, and it always outputs a vector x such that $c(x) \leq 2k$.*

Proof. Let $j \in \{1, \ldots, k\}$. We give an lower bound on $f_{\mathrm{avg}}(\pi_j) - f_{\mathrm{avg}}(\pi_{j-1})$. Suppose that π_j has a solution $x_j \in \mathbb{Z}_+$ and a realization $\psi_j \in \Phi^*_{x_j}$ when its last iteration is beginning, and (v_j, i_j) is chosen in Step 4 of the last iteration. Let $j' := c(x_j)$ and $C := c(x_j \vee i_j \chi_{v_j}) - j'$. The expected increase of the objective function in the last iteration of π_j is $\Delta(v_j, i_j \mid x_j, \psi_j)(j - j')/C$. If $j - 1 > j'$, then π_{j-1} behaves in the same way as π_j until it enters the last iteration, which updates a solution x_j to $x_j \vee i_j \chi_{v_j}$ with probability $(j - 1 - j')/C$. Hence the expected increase of the objective function in the last iteration is $\Delta(v_j, i_j \mid x_j, \psi_j)(j - 1 - j')/C$ in this case. If $j - 1 = j'$, π_{j-1} is the policy that does not execute the last iteration of π_j. In either case, the difference of the expected objective values achieved by π_j and π_{j-1} is

$$\frac{\Delta(v_j, i_j \mid x_j, \psi_j)}{C} \geq \frac{\Delta(v_j, i_j \mid x_j, \psi_j)}{\sum_{i \in [i_j]} c(v_j, i)}. \tag{5}$$

Notice that we are fixing x_j and ϕ_j in this discussion, whereas x_j and ϕ_j depends on the randomness of the variables in R. Taking the expectation of (5) over all full realizations, we have

$$f_{\mathrm{avg}}(\pi_j) - f_{\mathrm{avg}}(\pi_{j-1}) \geq E\left[\frac{\Delta(v_j, i_j \mid x_j, \psi_j)}{\sum_{i \in [i_j]} c(v_j, i)}\right]. \tag{6}$$

Next, we give an upper bound on $f_{\mathrm{avg}}(\pi_{j-1} @ \pi^\star) - f_{\mathrm{avg}}(\pi_{j-1})$. We again discuss with fixing x_j and ϕ_j. Suppose that π_{j-1} terminates with the realization ψ' and outputs a solution y. The $\pi^\star$-portion of $\pi_{j-1} @ \pi^\star$ increases the objective value by $\Delta(\pi^\star \mid y, \psi')$ in expectation. Notice that ψ' and y satisfy $\psi' \sim \psi_j$ and $y \geq x_j$. Hence, the adaptive submodularity of f indicates $\Delta(\pi^\star \mid y, \psi') \leq \Delta(\pi^\star \mid x_j, \psi_j)$. By taking the expectation over all full realizations, we have

$$f_{\mathrm{avg}}(\pi_{j-1} @ \pi^\star) - f_{\mathrm{avg}}(\pi_{j-1}) \leq E[\Delta(\pi^\star \mid x_j, \psi_j)]. \tag{7}$$

By Lemma 2, we have

$$\Delta(\pi^\star \mid x_j, \psi_j) \leq E[c(x_{\phi, \pi^\star}) \mid \phi \sim \psi_j] \max_{(v,i) \in V \times \mathbb{Z}_+} \frac{\Delta(v, i \mid x_j, \psi_j)}{\sum_{i' \in [i]} c(v, i')}. \tag{8}$$

The maximum in the right-hand side of (8) is attained by (v_j, i_j) by the definition. $E[c(x_{\phi, \pi^\star})] \leq k$ for any full realization ϕ because $\pi^\star$ is feasible. Hence, from (7) and (8), we have

$$f_{\mathrm{avg}}(\pi_{j-1} @ \pi^\star) - f_{\mathrm{avg}}(\pi_{j-1}) \leq k \cdot E\left[\frac{\Delta(v_j, i_j \mid x_j, \psi_j)}{\sum_{i' \in [i_j]} c(v_j, i')}\right]. \tag{9}$$

Define $\nu_j := f_{\mathrm{avg}}(\pi^\star) - f_{\mathrm{avg}}(\pi_j)$. Note that $\nu_{j-1} - \nu_j = f_{\mathrm{avg}}(\pi_j) - f_{\mathrm{avg}}(\pi_{j-1})$. Since f is adaptive monotone, Lemma 1 implies that $f_{\mathrm{avg}}(\pi^\star) \leq f_{\mathrm{avg}}(\pi_{j-1} @ \pi^\star)$. Hence, $f_{\mathrm{avg}}(\pi_{j-1} @ \pi^\star) - f_{\mathrm{avg}}(\pi_{j-1}) \geq \nu_{j-1}$. Therefore, (6) and (9) indicate that $\nu_{j-1} - \nu_j \geq \nu_{j-1}/k$ holds for all

$j = 1, \ldots, k$. This implies $\nu_k \leq \nu_0/e$, which is equivalent to $(1 - 1/e) f_{\mathrm{avg}}(\pi^\star) \leq f_{\mathrm{avg}}(\pi_k)$.

By its construction, π_k satisfies $E[c(x_{\phi,\pi_k})] \leq k$ for any full realization ϕ, and hence π_k is feasible. Let x be the vector when the last iteration is beginning, and suppose that the last iteration increases $x(v)$. Then $c(x_{\phi,\pi_k}) \leq c(x) + c(b(v)\chi_v)$ holds. $c(x) < k$ holds since otherwise π_k terminates before the last iteration, and we are assuming $c(b(v)\chi_v) \leq k$. Therefore, $c(x_{\phi,\pi_k}) \leq 2k$ holds for any full realization ϕ. ☐

Remark 2. *Each iteration of Policy 1 chooses a pair (v, i) that maximizes $\Delta(v, i \mid x, \psi)/\sum_{j\in[i]} c(v, j)$. Since increasing $x(v)$ to i costs $\sum_{j=x(v)+1}^{i} c(v, j)$, one may feel that (v, i) should maximize $\Delta(v, i \mid x, \psi)/\sum_{j=x(v)+1}^{i} c(v, j)$. In fact, we can prove the same guarantee even if (v, i) maximizes $\Delta(v, i \mid x, \psi)/\sum_{j=x(v)+1}^{i} c(v, j)$ because Lemma 2 still holds even after $\sum_{j\in[i]} c(v, j)$ is replaced by $\sum_{j=x(v)+1}^{i} c(v, j)$.*

Policy 1 can be modified so that it always outputs a vector x with $x \leq b$ and $c(x) \leq k$ whereas the approximation guarantee is reduced to $(e - 1)/(2e)$. See Algorithm 2 for its detail.

Policy 2 $(e - 1)/(2e)$-Approximation Policy

Input: a finite set V, an adaptive monotone submodular function $\{f_\phi \colon \mathbb{Z}_+^V \to \mathbb{R}_+\}_{\phi\in\Phi}$, $k \in \mathbb{R}_+$, $b \in \mathbb{Z}_+^V$, and $c \colon V \times \mathbb{Z}_+ \to \mathbb{R}_+$
Output: $x \in \mathbb{Z}_+^V$ such that $c(x) \leq k$, and $x \leq b$

```
x(v) ⟵ 0 for each v ∈ V
ψ(ρ) ⟵ ∗ for each ρ ∈ R
v′ ⟵ argmax_{v∈V} f_avg(b(v)χ_v)
In probability 1/2, set x(v′) ⟵ b(v′), and ψ(ρ) ⟵ the
  state of ρ for each observed random variable ρ ∈ R
while ∃(v, i) ∈ V × Z_+ such that Δ(v, i | x, ψ) > 0 and
  x ∨ iχ_v is feasible do
    (v, i) ⟵ argmax Δ(v, i | x, ψ)/Σ_{j∈[i]} c(v, j)
      where the maximization is over (v, i) ∈ V × Z_+ such
      that x ∨ iχ_v is feasible
    x(v) ⟵ i
    ψ(ρ) ⟵ the state of ρ for each observed random vari-
      able ρ ∈ R
end while
output x
```

Theorem 4. *Let π' denote Policy 2. If $f := \{f_\phi \mid \phi \in \Phi\}$ is adaptive monotone submodular, π' always outputs a vector x with $x \leq b$ and $c(x) \leq k$, and it achieves $f_{\mathrm{avg}}(\pi') \geq (e - 1)/(2e) \cdot f_{\mathrm{avg}}(\pi^\star)$ for any feasible policy $\pi^\star$.*

Proof. Define π as described in the paragraph before Policy 1. Since $f_{\mathrm{avg}}(\pi) \geq f_{\mathrm{avg}}(\pi_k)$, Theorem 3 indicates that $f_{\mathrm{avg}}(\pi) \geq (1 - 1/e) f_{\mathrm{avg}}(\pi^\star)$ holds. Let x' be the vector kept by the policy π when the last iteration begins, and suppose that π increases $x'(v)$ to i in the last iteration. Note that $f_{\mathrm{avg}}(\pi) = E[f_\phi(x' \vee i\chi_v)]$.

Let π_1 denote the policy that behaves as π with the exception that π_1 does not execute the last iteration of π. In addition, let π_2 be the policy that always outputs $b(v')\chi_{v'}$, where v' maximizes $f_{\mathrm{avg}}(b(v')\chi_{v'})$. We first prove $f_{\mathrm{avg}}(\pi) \leq f_{\mathrm{avg}}(\pi_1) + f_{\mathrm{avg}}(\pi_2)$. Clearly, $f_{\mathrm{avg}}(\pi_1) = E[f_\phi(x')]$. Let $y \in \mathbb{Z}_+^V$ denote the zero-vector, and ψ denote the realization such that $\psi(\rho) = *$ for all $\rho \in R$. The adaptive submodularity of f indicates $E[f_\phi(x' \vee i\chi_v)] - E[f_\phi(x')] \leq \Delta(v, i \mid y, \psi)$. Since $f_{\mathrm{avg}}(\pi_2) = f_{\mathrm{avg}}(b(v')\chi_{v'}) \geq \Delta(v, i \mid y, \psi)$, we obtain the inequality.

For each full realization, the objective value of a vector output by π' is at least $f(\pi_1)$ in probability $1/2$; otherwise, this is at least $f(\pi_2)$. Therefore, $f_{\mathrm{avg}}(\pi') \geq (f_{\mathrm{avg}}(\pi_1) + f_{\mathrm{avg}}(\pi_2))/2 \geq f_{\mathrm{avg}}(\pi)/2 \geq (e - 1)/(2e) \cdot f_{\mathrm{avg}}(\pi^\star)$. ☐

Let us briefly discuss the running time of Policies 1 and 2. Suppose that $\Delta(v, i \mid x, \phi)$ can be computed in $O(T)$ time for any $v \in V$, $i \in [b(v)]$, $x \in \mathbb{Z}_+^V$, and $\phi \in \Phi_x^*$. Then, both policies decide the next behavior in $O(T \sum_{v\in V} b(v))$ time in each iteration. In the budget allocation problem with the bipartite influence model, we have $T = O(|U| \max_{v\in V} b(v))$. Note that this does not consider the time for observing the states of random variables, which corresponds to checking whether each customer is influenced or not in the budget allocation problem.

Remark 3. *Asadpour and Nazerzadeh [2009] showed that an adaptive algorithm is better than any non-adaptive algorithm by a factor $e/(e-1) > 1.58$ for the stochastic maximum k-cover problem, which is a special case of the budget allocation problem in the bipartite influence model. Their proof gave an instance of the stochastic maximum k-cover problem for which an adaptive policy achieves an objective value L and any non-adaptive solution does not achieve an objective value better than $(1 - 1/e)L$. In fact, the adaptive policy in their proof coincides with Policy 1. This indicates that our adaptive policies improve the objective value by at least 58% than arbitrary non-adaptive algorithms for those instances.*

5 Experiments

We implemented three adaptive policies: Policies 1 and 2, and a sensitive greedy policy defined as follows. Suppose that the policy maintains a vector $x \in \mathbb{Z}_+^V$ and a realization $\phi \in \Phi_x^*$ when a certain iteration begins. In this iteration, the policy computes $(v, i) := \arg\max \Delta(v, i \mid x, \phi)/(\sum_{j\in[i]} c(v, j))$ and increases $x(v)$ by one, where the maximization is taken over all $(v, i) \in V \times \mathbb{Z}_+$ such that $x \vee i\chi_v$ is feasible. In addition to the adaptive policies, we implemented a non-adaptive greedy $(1 - 1/e)$-approximation algorithm [Soma *et al.*, 2014].

5.1 Bipartite influence model

We run the algorithms for instances of the bipartite influence model. As a bipartite graph, we prepared a synthetic graph $(V, U; E)$ over a media set V and a customer set U. The degree distribution on V follows the power law, and $|V| = 100$ and $|U| = 10{,}000$. We randomly chose $b(v)$ from $\{20, 21, \cdots, 30\}$ for each $v \in V$. We prepared two

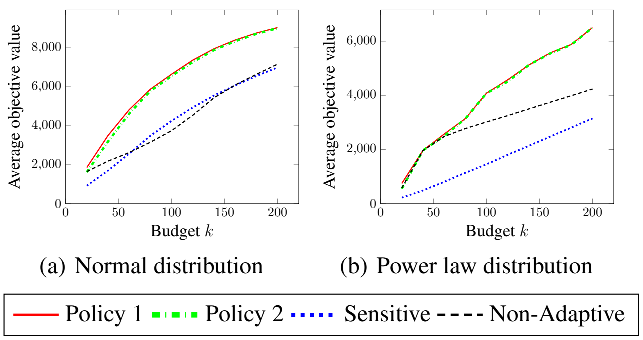

(a) Normal distribution (b) Power law distribution

Figure 1: Experimental results on the bipartite influence model

types of the probabilities q_{vu}, $vu \in E$: In the normal distribution, $q_{vu}(i)$ is given by $\exp(-(i-15)^2/50)/\sqrt{50\pi}$ for each $i \in \{1, \ldots, 30\}$ and $vu \in E$; In the power law distribution, $q_{vu}(i)$ is given by $\exp(0.2(i-30))/10$ for each $i \in \{1, \ldots, 30\}$ and $vu \in E$.

We compute budget allocations over 500 instances by the policies, and compare their objective values by $f_{\mathrm{avg}}(\pi)$ for a policy π. By preliminary experiments, we verified that 500 instances are enough to compare the average objective values.

Figure 1 indicates the average objective values when the budget k is set to a value in $\{20, 40, \ldots, 200\}$. For setting the probabilities q_{vu} of each edge vu, the normal distribution is used in (a), and the power law distribution is used in (b).

Although the theoretical performance guarantee of Policy 2 (in Theorem 4) is inferior to Policy 1 (in Theorem 3), we observe from the experimental results that performances of Policies 1 and 2 are almost same in all instances. Recall that Policy 2 always outputs a feasible allocation whereas Policies 1 does not. Moreover, they are clearly superior to the other two algorithms. In the normal distribution instances, the non-adaptive algorithm is sometimes worse even than the sensitive algorithm. In the power law distribution instances, the non-adaptive algorithm outperforms the sensitive algorithm, but it is clearly worse than Policies 1 and 2.

5.2 General influence model

Adaptive submodularity over an integer lattice is a useful notion, and it has numerous applications other than the budget allocation in the bipartite influence model. Taking advantage of this feature, we extend the budget allocation problem from the bipartite influence model to another influence model defined over general directed graphs. This general influence model can be captured by an adaptive monotone submodular function over an integer lattice, and hence all algorithmic results proposed in this paper can be applied to it. Here, we report empirical performance of the adaptive strategies and the non-adaptive $(1-1/e)$-approximation algorithm in this general influence model.

We do not describe the detail of the model due to the space limitation. The non-adaptive setting of this model is obtained by introducing multiple influence levels into the independence cascade model studied by Kempe et al. [2003] in a context of influence maximization. We note that a similar attempt can be found in Demaine et al. [2014], but our model is different from theirs. Alon et al. [2012] also mentioned that their bipartite influence model can be naturally extended to general graphs, but they do not seem to consider the multiple influence levels on all nodes.

For the experiments, we prepared a graph that represents user-user following information in Twitter [KONECT, 2014]. Each node represents a user, and an arc from a node i to another node j represents that the user corresponding to i is followed by the user corresponding to j. The graph consists of 23370 nodes and 33101 arcs. In this graph, we choose 500 nodes that have largest out-degrees, and consider allocating budgets to these nodes. The parameters in the instances are set as follows: $b(v) = 15$ for all chosen nodes v, and the objective of the problem is defined as the maximization of the number of nodes influenced at least once. Budget k is set to a value in $\{20, 40, \ldots, 200\}$, and the objective values are averaged over 500 instances for each k.

Figure 2 shows the results. Performance of Policies 1 and 2 are nearly equal, the non-adaptive policy is behind them, and the sensitive policy is clearly worse than the others.

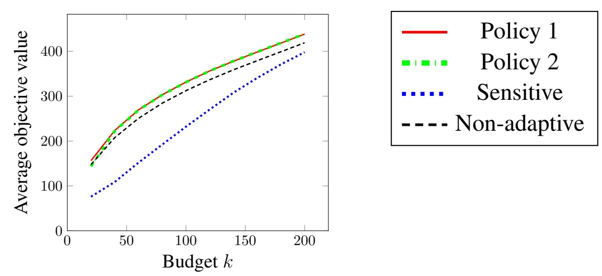

Figure 2: Experimental results on the general influence model with a Twitter graph

6 Conclusion

In this paper, we analyzed adaptive greedy policies for the budget allocation problem. Our contributions are based on the new concept of the adaptive submodularity; we extended the adaptive submodularity defined for the set functions to the functions over integer lattices. We believe that this new concept has other applications than the budget allocation problem. Indeed, we have already studied its applications to data summarization and sensor management. We will report them in the full version of the present paper.

References

[Alon *et al.*, 2012] Noga Alon, Iftah. Gamzu, and Moshe Tennenholtz. Optimizing budget allocation among channels and influencers. In *21st International Conference on World Wide Web (WWW)*, pages 381–388, 2012.

[Asadpour and Nazerzadeh, 2009] Arash Asadpour and Hamid Nazerzadeh. Maximizing Stochastic Monotone Submodular Functions. *ArXiv e-prints*, August 2009.

[Chen and Krause, 2013] Yuxin Chen and Andreas Krause. Near-optimal batch mode active learning and adaptive submodular optimization. In *30th International Conference on Machine Learning (ICML)*, pages 160–168, 2013.

[Chen *et al.*, 2014] Yuxin Chen, Hiroaki Shioi, Cesar Fuentes Montesinos, Lian Pin Koh, Serge Wich, and Andreas Krause. Active detection via adaptive submodularity. In *31th International Conference on Machine Learning (ICML)*, pages 55–63, 2014.

[Chen *et al.*, 2015] Yuxin Chen, Shervin Javdani, Amin Karbasi, J. Andrew (Drew) Bagnell, Siddhartha Srinivasa, and Andreas Krause. Submodular surrogates for value of information. In *The Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-15)*, pages 3511–3518, 2015.

[Demaine *et al.*, 2014] Erik D. Demaine, MohammadTaghi Hajiaghayi, Hamid Mahini, David L. Malec, S. Raghavan, Anshul Sawant, and Morteza Zadimoghaddam. How to influence people with partial incentives. In *23rd International World Wide Web Conference (WWW)*, pages 937–948, 2014.

[Deshpande *et al.*, 2014] Amol Deshpande, Lisa Hellerstein, and Devorah Kletenik. Approximation algorithms for stochastic boolean function evaluation and stochastic submodular set cover. In *Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1453–1466, 2014.

[Domingos and Richardson, 2001] Pedro M. Domingos and Matthew Richardson. Mining the network value of customers. In *Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 57–66, 2001.

[Gabillon *et al.*, 2013] Victor Gabillon, Branislav Kveton, Zheng Wen, Brian Eriksson, and S. Muthukrishnan. Adaptive submodular maximization in bandit setting. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 2697–2705, 2013.

[Gabillon *et al.*, 2014] Victor Gabillon, Branislav Kveton, Zheng Wen, Brian Eriksson, and S. Muthukrishnan. Large-scale optimistic adaptive submodularity. In *Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI-14)*, pages 1816–1823, 2014.

[Golovin and Krause, 2011a] Daniel Golovin and Andreas Krause. Adaptive submodular optimization under matroid constraints. *CoRR*, abs/1101.4450, 2011.

[Golovin and Krause, 2011b] Daniel Golovin and Andreas Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *J. Artif. Intell. Res. (JAIR)*, 42:427–486, 2011.

[Golovin *et al.*, 2010] Daniel Golovin, Andreas Krause, and Debajyoti Ray. Near-optimal bayesian active learning with noisy observations. In *Advances in Neural Information Processing Systems 23: 24th Annual Conference on Neural Information Processing Systems 2010 (NIPS)*, pages 766–774, 2010.

[Gotovos *et al.*, 2015] Alkis Gotovos, Amin Karbasi, and Andreas Krause. Non-monotone adaptive submodular maximization. In *Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI-15)*, pages 1996–2003, 2015.

[Kempe *et al.*, 2003] David Kempe, Jon M. Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 137–146, 2003.

[KONECT, 2014] KONECT. Twitter lists network dataset, 2014. http://konect.uni-koblenz.de/networks/ego-twitter.

[Krause *et al.*, 2014] Andreas Krause, Daniel Golovin, and Sarah J. Converse. Sequential decision making in computational sustainability via adaptive submodularity. *AI Magazine*, 35(2):8–18, 2014.

[Richardson and Domingos, 2002] Matthew Richardson and Pedro M. Domingos. Mining knowledge-sharing sites for viral marketing. In *8th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 61–70, 2002.

[Soma *et al.*, 2014] Tasuku Soma, Naonori Kakimura, Kazuhiro Inaba, and Ken-ichi Kawarabayashi. Optimal budget allocation: Theoretical guarantee and efficient algorithm. In *31th International Conference on Machine Learning (ICML)*, pages 351–359, 2014.

[The Washington Post, 2012] The Washington Post, 2012. http://www.washingtonpost.com/wp-srv/special/politics/track-presidential-campaign-ads-2012/.