Better Bounds on the Adaptivity Gap of Influence Maximization under Full-adoption Feedback

Gianlorenzo D'Angelo, Debashmita Poddar, and Cosimo Vinci
Gran Sasso Science Institute, L'Aquila, Italy
{gianlorenzo.dangelo, debashmita.poddar, cosimo.vinci}@gssi.it

Abstract

In the *influence maximization (IM)* problem, we are given a social network and a budget k, and we look for a set of k nodes in the network, called seeds, that maximize the expected number of nodes that are reached by an influence cascade generated by the seeds, according to some stochastic model for influence diffusion. Extensive studies have been done on the IM problem, since his definition by Kempe, Kleinberg, and Tardos (2003). However, most of the work focuses on the *non-adaptive* version of the problem where all the k seed nodes must be selected before that the cascade starts. In this paper we study the *adaptive* IM, where the nodes are selected sequentially one by one, and the decision on the i-th seed can be based on the *observed* cascade produced by the first $i-1$ seeds. We focus on the *full-adoption feedback* in which we can observe the entire cascade of each previously selected seed and on the independent cascade model where each edge is associated with an independent probability of diffusing influence.

Previous works showed that there are constant upper bounds on the adaptivity gap, which compares the performance of an adaptive algorithm against a non-adaptive one, but the analyses used to prove these bounds only works for specific graph classes such as in-arborescences, out-arborescences, and one-directional bipartite graphs. Our main result is the first sub-linear upper bound that holds for any graph. Specifically, we show that the adaptivity gap is upper-bounded by $\sqrt[3]{n}+1$, where n is the number of nodes in the graph. Moreover we improve over the known upper bound for in-arborescences from $2e/(e-1) \approx 3.16$ to $2e^2/(e^2-1) \approx 2.31$. Finally, we study α-bounded graphs, a class of undirected graphs in which the sum of node degrees higher than two is at most α, and show that the adaptivity gap is upper-bounded by $\sqrt{\alpha}+O(1)$. Moreover, we show that in 0-bounded graphs, i.e. undirected graphs in which each connected component is a path or a cycle, the adaptivity gap is at most $3e^3/(e^3-1) \approx 3.16$.

To prove our bounds, we introduce new techniques to relate adaptive policies with non-adaptive ones that might be of their own interest.

Introduction

In the Influence Maximization (IM) problem, we are given a social network, a stochastic model for diffusion of influence over the network, and a budget k, and we ask to find a set of k nodes, called *seeds*, that maximize their *spread of influence*, which is the expected number of nodes reached by a cascade of influence diffusion generated by the seeds according to the given diffusion model. One of the most studied model for influence diffusion is the Independent Cascade (IC), where each edge is associated with an independent probability of transmitting influence from the source node to the tail node. In the IC model the spread of influence is a monotone submodular function of the seed set, therefore a greedy algorithm guarantees a $1-\frac{1}{e}$ approximation factor for the IM problem (Kempe, Kleinberg, and Tardos 2015). Since his definition (Domingos and Richardson 2001; Richardson and Domingos 2002) and formalization as an optimization problem (Kempe, Kleinberg, and Tardos 2003, 2015), the IM problem and its variants have been extensively investigated, motivated by applications in viral marketing (Chen, Wang, and Wang 2010), adoption of technological innovations (Goldberg and Liu 2013), and outbreak or failure detection (Leskovec et al. 2007). See Chen, Lakshmanan, and Castillo (2013); Li et al. (2018) for surveys on the IM problem.

Recently, Golovin and Krause (2011) initiated the study of the IM problem under the framework of adaptive optimization, where, instead of selecting all the seeds at once at the beginning of the process, we can select one seed at a time and observe, to some extent, the portion of the network reached by a new selected seed. The advantage is that the decision on the next seed to choose can be based on the observed spread of previously selected seeds, usually called *feedback*. Two main feedback models have been introduced: in the *full-adoption* feedback the whole spread from each seed can be observed, while in the *myopic* feedback one can only observe the direct neighbors of each seed.

Golovin and Krause considered the Independent Cascade model and showed that, under full-adoption feedback, the objective function satisfies the property of *adaptive submodularity* (introduced in the same paper) and therefore a greedy algorithm achieves a $1-\frac{1}{e}$ approximation for the adaptive IM problem. They also conjectured that there exists a constant factor approximation algorithm for the myopic feedback model, which has been indeed found by Peng and Chen (2019) who proposed a $\frac{1}{4}\left(1-\frac{1}{e}\right)$-approximation algorithm.

However, the approximation ratio for the adaptive IM problem, which compares a given adaptive algorithm with an optimal adaptive one, does not measure the benefits of implementing adaptive policies over non-adaptive ones. To this aim, Chen and Peng (Chen and Peng 2019; Peng and Chen 2019) introduced the adaptivity gap, which is the supremum, over all possible inputs, of the ratio between the spread of an optimal adaptive policy and that of an optimal non-adaptive one. In (Peng and Chen 2019), Peng and Chen considered independent cascade model with myopic feedback and showed that the adaptivity gap is between $\frac{e}{e-1}$ and 4 for any graph. In (Chen and Peng 2019), the same authors showed some upper and lower bounds on the adaptivity gap in the case of full-adoption feedback, still under independent cascade, for some particular graph classes. Specifically, they showed that the adaptivity gap is in the interval $\left[\frac{e}{e-1}, \frac{2e}{e-1}\right]$ for in-arborescences and it is in the interval $\left[\frac{e}{e-1}, 2\right]$ for out-arborescences. Moreover, it is equal to $\frac{e}{e-1}$ in one-directional bipartite graphs. In order to show these bounds, they followed an approach introduced by Asadpour and Nazerzadeh (2016) which consists in transforming an adaptive policy into a non-adaptive one by means of multilinear extensions, and constructing a Poisson process to relate the influence spread of the non-adaptive policy to that of the adaptive one. For general graphs and full-adoption feedback, the only known upper bounds on the adaptivity gap are linear in the size of the graph and can be trivially derived.

In this paper, we consider the independent cascade model with full-adoption feedback, and show the first sub-linear upper bound on the adaptivity gap that holds for general graphs. In detail we show that that the adaptivity gap is at most $\lceil n^{1/3} \rceil$, where n is the number of nodes in the graph. Moreover, we tighten the upper bound on the adaptivity gap for in-arborescences by showing that it is at most $\frac{2e^2}{e^2-1} < \frac{2e}{e-1}$. Using similar techniques we study the adaptivity gap of α-*bounded graphs*, which is the class of undirected graphs where the sum of node degrees higher than two is at most α. We show that the adaptivity gap is upper-bounded by $\sqrt{\alpha} + O(1)$, which is smaller that $O(n^{1/3})$ for several graph classes. In 0-bounded graphs, i.e. undirected graphs in which each connected component is a path or a cycle, the adaptivity gap is at most $\frac{3e^3}{e^3-1}$.

To prove our bounds, we introduce new techniques to connect adaptive policies with non-adaptive ones that might be of their own interest (further details are given in paragraph "General outline of the proof technique" in Section). In particular, we resort to a simple and randomized *hybrid non-adaptive policy*, that is not based on the Poisson process and the multi-linear extensions, which instead represent the main probabilistic tools adopted by (Asadpour and Nazerzadeh 2016; Chen and Peng 2019)).

Related Work

Influence Maximization. Several studies based on general graphs (Lowalekar, Varakantham, and Kumar 2016; Mihara, Tsugawa, and Ohsaki 2015; Schoenebeck and Tao 2019; Tang, Xiao, and Shi 2014) have been conducted since the seminal paper by Kempe, Kleinberg, and Tardos (2015). Schoenebeck and Tao (2019) studied the influence maximization problem on undirected graphs and proves that it is APX-hard for both the independent cascade and the linear threshold problem. Borgs et al. (2014) propose an efficient algorithm that runs in quasilinear time and still guarantees an approximation factor of $1 - \frac{1}{e} - \epsilon$, for any $\epsilon > 0$. Tang, Xiao, and Shi (2014) propose an algorithm which is experimentally close to the optimal one under the independent cascade model. Mihara, Tsugawa, and Ohsaki (2015) consider unknown graphs for the influence maximization problem and devises an algorithm which achieves a fraction between 0.6 and 0.9 of the influence spread with minimal knowledge of the graph topology. Extensive literature reviews on influence maximization and its machinery is provided by Chen, Lakshmanan, and Castillo (2013) and Li et al. (2018).

Several works on the adaptive influence maximization problem (Cautis, Maniu, and Tziortziotis 2019; Han et al. 2018; Sun et al. 2018; Tang et al. 2019; Tong and Wang 2019; Tong et al. 2017; Vaswani and Lakshmanan 2016; Yuan and Tang 2017) evolved after the concept introduced by Golovin and Krause (2011), and explore the adaptive optimization under different feedback models. The myopic model (in which, one can only observe the nodes influenced by the seed nodes) has been studied in (Peng and Chen 2019; Salha, Tziortziotis, and Vazirgiannis 2018). Sun et al. (2018) capture the scenario in which, instead of considering one round, the diffusion process takes over T rounds, and a seed set of at most k nodes is selected at each round. The authors designed a greedy approximation algorithm that guarantees a constant approximation ratio. Tong et al. (2020) introduce a new version of the adaptive influence maximization problem by adding a time constraint. Other than the classic full-adoption and myopic feedback model, Yuan and Tang (2017), and Tong and Wang (2019), have also introduced different feedback models that use different parameters to overcome the need of submodularity to guarantee a good approximation. Han et al. (2018) propose a framework which uses existing non-adaptive techniques to construct a strong approximation for a generalization of the adaptive influence maximization problem in which in each step a batch of node is selected.

Adaptivity Gaps. Adaptivity gaps for the problem of maximizing stochastic monotone submodular functions have been studied by Asadpour and Nazerzadeh (2016). A series of work studied adaptivity gaps for a two-step adaptive influence maximization problem (Badanidiyuru et al. 2016; Rubinstein, Seeman, and Singer 2015; Seeman and Singer 2013; Singer 2016). Gupta et al. (Gupta, Nagarajan, and Singla 2016, 2017) worked on the adaptivity gaps for stochastic probing. A recent line of studies has been conducted (Chen and Peng 2019; Chen et al. 2020; Peng and Chen 2019) which focuses on finding the adaptivity gaps on different graph classes using the classical feedback models. Peng and Chen (2019) confirmed a conjecture of Golovin and Krause (2011), which states that the adaptive greedy algorithm with myopic feedback is a constant approximation of the adaptive optimal solution. They show that the adaptivity gap of the independent cascade model with myopic feedback belongs to $[\frac{e}{e-1}, 4]$. Chen et al. (2020) introduced the greedy adaptivity gap, which compares the performance of the adaptive and the non-adaptive greedy algorithms. They show that the infimum of the greedy adaptivity gap is $1 - \frac{1}{e}$ for every combination of diffusion and feedback models. The most related work to our results is that of Chen and Peng (2019). Chen and Peng (2019) derive upper and lower bounds on the adaptivity gap under the independent cascade model with full-adoption feedback, when the considered graphs are in-arborescences, out-arborescences, or one-directional bipartite graphs. In particular, they show that the adaptivity gaps of in-arborescences and out-arborescences are in the intervals $\left[\frac{e}{e-1}, \frac{2e}{e-1}\right]$ and $\left[\frac{e}{e-1}, 2\right]$, respectively, and they provide a tight bound of $\frac{e}{e-1}$ on the adaptivity gap of one-directional bipartite graphs.

Organization of the Paper

In Section we give the preliminary definitions and notations which this work is based on. Sections – are devoted to the main technical contribution of the paper (i.e., adaptivity gaps of in-arborescences, general graphs, and α-bounded graphs). In Section , we highlight some future research directions. Due to the lack of space, some missing proofs are deferred to the full version of this work.

Preliminaries

For two integers h and k, $h \leq k$, let $[k]_h := \{h, h+1, \ldots, k\}$ and $[k] := [k]_1$.

Independent Cascade Model. In the *independent cascade model* (IC), we have an *influence graph* $G = (V = [n], E, (p_{uv})_{(u,v)\in E})$, where $p_{uv} \in [0, 1]$ is an *activation probability* associated to each edge $(u, v) \in E$. Given a set of *seed nodes* $S \subseteq V$ which are initially *active*, the diffusion process in the IC model is defined in $t \geq 0$ discrete steps as follows: (i) let A_t be the set of active nodes which are activated at each step $t \geq 0$; (ii) $A_0 := S$; (iii) given a step $t \geq 0$, for any edge (u, v) such that $u \in A_t$, node u can activate node v with probability p_{uv} independently from any other node, and, in case of success, v is included in A_{t+1}; (iv) the diffusion process ends at a step $r \geq 0$ such that $A_r = \emptyset$, i.e., no node can be activated at all, and $\bigcup_{t\leq r} A_t$ is the *influence spread*, i.e., the set of nodes activated/reached by the diffusion process.

The above diffusion process can be equivalently defined as follows. The *live-edge graph* $L = (V, L(E))$ of G is a random graph made from G, where $L(E) \subseteq E$ is a subset of edges such that each edge $(u, v) \in E$ is included in $L(E)$ with probability p_{uv}, independently from the other edges. Given a live-edge graph L, let $R(S, L) := \{v \in V :$ there exists a path from u to v in L for some $u \in S\}$, i.e., the set of nodes reached by nodes in S in the live-edge graph L. Informally, if S is the set of seed nodes, and L is a live-edge graph, $R(S, L)$ equivalently denotes the set of nodes which are reached/activated by the above diffusion process. Given a set of seed nodes S, the *expected influence spread* of S is defined as $\sigma(S) := \mathbb{E}_L[|R(S, L)|]$.

Non-adaptive Influence Maximization. The *non-adaptive influence maximization problem under the IC model* is the computational problem that, given an influence graph G and an integer $k \geq 1$, asks to find a set of seed nodes $S \subseteq V$ with $|S| = k$ such that $\sigma(S)$ is maximized.

Adaptive Influence Maximization. Differently from the non-adaptive setting, in which all the seed nodes are activated at the beginning and then the influence spread is observed, an *adaptive policy* activates the seeds sequentially in k steps, one seed node at each step, and the decision on the next seed node to select is based on the feedback resulting from the observed spread of previously selected nodes. The feedback model considered in this work is *full-adoption*: when a node is selected, the adaptive policy observes its entire influence spread.

An adaptive policy under the full-adoption feedback model is formally defined as follows. Given a live-edge graph L, the *realisation* $\phi_L : V \rightarrow 2^V$ associated to L assigns to each node $v \in V$ the value $\phi_L(v) := R(\{v\}, L)$, i.e., the set of nodes activated by v under a live-edge graph L. Given a set $S \subseteq V$, a *partial realisation* $\psi : S \rightarrow 2^V$ is the restriction to S of some realisation, i.e., there exists a live-edge graph L such that $\psi(v) = \phi_L(v)$ for any $v \in S$. Given a partial realisation $\psi : S \rightarrow 2^V$, let $dom(\psi) := S$, i.e., $dom(\psi)$ is the domain of partial realisation ψ, let $R(\psi) := \bigcup_{v\in S} \psi(v)$, i.e., $R(\psi)$ is the set of nodes reached/activated by the diffusion process when the set of seed nodes is S, and let $f(\psi) := |R(\psi)|$. A partial realisation ψ' is a *sub-realisation* of ψ (or, equivalently, $\psi' \subseteq \psi$), if $dom(\psi') \subseteq dom(\psi)$ and $\psi'(v) = \psi(v)$ for any $v \in dom(\psi')$. We observe that a partial realisation ψ can be equivalently represented as $\{(v, R(\{v\}, L)) : v \in dom(\psi)\}$ for some live-edge graph L.

An adaptive policy π takes as input a partial realisation ψ and, either returns a node $\pi(\psi) \in V$ and activates it as seed, or interrupts the activation of new seed nodes, e.g., by returning a string $\pi(\psi) := STOP$. An adaptive policy π can be run as in Algorithm 1. The *expected influence spread* of an adaptive policy π is defined as $\sigma(\pi) := \mathbb{E}_L[f(\psi_{\pi,L})]$, i.e., it is the expected value (taken on all the possible live-edge graphs) of the number of nodes reached by the diffusion process at the end of Algorithm 1. We say that $|\pi| = k$ if policy π always return a partial realisation $\psi_{\pi,L}$ with $|dom(\psi_{\pi,L})| = k$. The *adaptive influence maximization problem under the IC model* is the computational problem that, given an influence graph G and an integer $k \geq 1$, asks to find an adaptive policy π that maximizes the expected influence spread $\sigma(\pi)$ subject to constraint $|\pi| = k$.

Adaptivity gap. Given an influence graph G and an integer $k \geq 1$, let $OPT_N(G, k)$ (resp. $OPT_A(G, k)$) denote the optimal value of the non-adaptive (resp. adaptive) influence maximization problem with input G and k. Given a class of influence graphs $\mathcal{G}$ and an integer $k \geq 1$, the *k-adaptivity gap* of $\mathcal{G}$ is defined as

$$AG(\mathcal{G}, k) := \sup_{G\in\mathcal{G}} \frac{OPT_A(G, k)}{OPT_N(G, k)},$$

Algorithm 1 Adaptive algorithm

Require: an influence graph G and an adaptive policy π;
Ensure: a partial realisation;
let L be the live-edge graph;
let $\psi := \emptyset$ (i.e., ψ is the empty partial realisation);
while $\pi(\psi) \neq STOP$ **do**
  $v := \pi(\psi)$;
  $\psi := \psi \cup \{(v, R(\{v\}, L))\}$;
end while
return $\psi_{\pi,L} := \psi$;

and measures how much an adaptive policy outperforms a non-adaptive solution for the influence maximization problem applied to influence graphs in $\mathcal{G}$, when the maximum number of seed nodes is k. The *adaptivity gap* of $\mathcal{G}$ is defined as $AG(\mathcal{G}) := \sup_{k\geq 1} AG(\mathcal{G}, k)$. We observe that for $k = 1$ or $n \leq k$ the k-adaptivity gap is trivially equal to 1, thus we omit such cases in the following.

Adaptivity Gap for In-arborescences

An *in-arborescence* is a graph $G = (V, E)$ that can be constructed from a rooted tree $T = (V, F)$, by adding in E an edge (v, u) if u is a father of v in tree T. An upper bound of $\frac{2e}{e-1} \approx 3.16$ on the adaptivity gap of in-arborescences has been provided in (Chen and Peng 2019). In this section we provide an improved upper bound for such graphs.

Theorem 1. *If $\mathcal{G}$ is the class of all the in-arborescences, then*

$$AG(\mathcal{G}, k) \leq \frac{2}{1-(1-2/k)^k} \leq \frac{2e^2}{e^2-1} \approx 2.31, \ \forall k \geq 2.$$

Let $G = (V = [n], E, (p_{uv})_{(u,v)\in E})$ be an in-arborescence, where $n > k$ is the number of nodes. To show the claim of Theorem 1, we need some preliminary notations and lemmas. Given a partial realisation ψ, and a node $i \in [n]$, let

$$\Delta(i|\psi) := \mathbb{E}_L[f(\psi \cup \{(i, R(\{i\}, L))\}) - f(\psi)|\psi \subseteq \phi_L],$$

i.e., $\Delta(i|\psi)$ is the expected increment of the influence spread due to node i when the observed partial realisation is ψ. We have the following claim (from (Golovin and Krause 2011)), holding even for general graphs, whose proof is trivial.

Claim 1. *Adaptive Submodularity, (Golovin and Krause 2011) Let G be an arbitrary influence graph. For any partial realisations ψ, ψ' of G such that $\psi \subseteq \psi'$, and any node $i \notin R(\psi')$, we have that $\Delta(i|\psi') \leq \Delta(i|\psi)$.*

An adaptive policy π is called *randomized* if, for any partial realisation ψ, node $\pi(\psi)$ is not selected deterministically (in general), but randomly (according to a probability distribution p_ψ depending on ψ). Given a vector $\boldsymbol{y} = (y_1, \ldots, y_n)$ such that $y_i \in [0, 1]$ for any $i \in [n]$, we say that $\mathbb{P}(\pi) = \boldsymbol{y}$ if the probability that each node i belongs to $dom(\psi_{\pi,L})$ is y_i, where $\psi_{\pi,L}$ is the partial realisation returned by Algorithm 1 with policy π. Let $OPT_A(G, \boldsymbol{y})$ be the optimal expected influence spread $\sigma(\pi)$ over all the randomized adaptive policies π such that $\mathbb{P}(\pi) = \boldsymbol{y}$.[1]

Let π^* be an optimal adaptive policy for the adaptive influence maximization problem (with $|\pi^*| = k$), and let $\boldsymbol{x} = (x_1, \ldots, x_n)$ be the vector such that $\mathbb{P}(\pi^*) = \boldsymbol{x}$. As $|\pi^*| = k$, we have that $\sum_{i\in[n]} x_i = k$.

For any $t \in [k]_0$, let S_t be the optimal set of t seed nodes in the non-adaptive influence maximization problem, i.e., such that $OPT_N(G, t) = \mathbb{E}_L(|R(S_t, L)|)$. Let $\psi_{t,L}$ be the random variable denoting the sub-realisation of ϕ_L such that $dom(\psi_{t,L}) = S_t$. Let ρ be the random variable equal to node $i \in [n]$ with probability x_i/k. Observe that the above random variable is well-defined, as $\sum_{i\in[n]}(x_i/k) = k/k = 1$. For any $t \in [k]$, let $\psi_{\rho,t,L}$ be the random variable denoting the sub-realisation of ϕ_L such that $dom(\psi_{\rho,t,L}) = S_{t-1} \cup \{\rho\}$.

General outline of the proof technique. We observe that $\psi_{\rho,t,L}$ is the partial realisation coming from the following *hybrid non-adaptive policy*: initially, we activate the first $t-1$ seed nodes as in the optimal non-adaptive solution guaranteeing an expected influence spread of $OPT_N(G, t-1)$; then, we randomly choose a node v according to random variable ρ and we select v as t-th seed node (if not already selected as seed). We use this hybrid non-adaptive policy as a main tool to obtain an improved upper bound on the adaptivity gap for in-arborescences. In Lemma 1, holding even for general graphs, we relate the hybrid non-adaptive policy and the optimal non-adaptive solution, with the optimal adaptive policy. Lemma 1, together with Lemma 2 (that is similar to Lemma 3.8 in (Chen and Peng 2019)), is used in the main proof of the theorem to relate $OPT_N(G, t)$ with $OPT_A(G, k)$ for any $t \in [k]$, and this leads to our upper bound.

The proof structure of Lemma 1 exhibits some similarities with Lemma 6 of (Asadpour and Nazerzadeh 2016) and Lemma 3.3 of (Chen and Peng 2019), but in their approach, they relate non-adaptive policies based on the Poisson process and multi-linear extensions, with the optimal adaptive policy. One disadvantage of the Poisson process adopted in (Chen and Peng 2019) is that the number X of seed nodes selected by the corresponding non-adaptive policy is equal to k under expectation (i.e., $\mathbb{E}(X) = k$), and determining the expected influence spread w.r.t. random variable X has implied a further loss in the final upper bound (see Lemma 3.9 and inequality (21) of Theorem 3.1 in (Chen and Peng 2019)). Instead, by using the hybrid-non-adaptive policy, we guarantee that the number of selected seed nodes at each step $t \in [k]$ is exactly equal to t, independently from the considered random execution. This property allow us to avoid the expectations w.r.t. the number of selected seed nodes, and this leads to a further improvement of the resulting upper bound on the adaptivity gap.

Lemma 1. *Let G be an arbitrary influence graph. For any $t \in [k]$, and any fixed partial realisation ψ of G such that $\mathbb{P}[\psi_{t-1,L} = \psi] > 0$, we have $OPT_A(G, k) \leq \sigma(R(\psi)) + k \cdot \mathbb{E}_{L,\rho}\left[f(\psi_{\rho,t,L}) - f(\psi_{t-1,L})|\psi_{t-1,L} = \psi\right]$.*

[1] We observe that, if $\boldsymbol{y}$ is arbitrary, a deterministic policy π verifying $\mathbb{P}(\pi) = \boldsymbol{y}$ might not exists, and the introduction of randomization solves this issue.

Proof. We have

$$
\begin{aligned}
& k \cdot \mathbb{E}_{L,\rho}\left[f(\psi_{\rho,t,L}) - f(\psi_{t-1,L}) | \psi_{t-1,L} = \psi\right] \\
=& k \cdot \sum_{i\in[n]} \mathbb{P}[\rho = i] \cdot \Delta(i|\psi) \\
=& k \cdot \sum_{i\in[n]\setminus R(\psi)} \frac{x_i}{k} \cdot \Delta(i|\psi) \quad (1)\\
=& \sum_{i\in[n]\setminus R(\psi)} x_i \cdot \Delta(i|\psi), \quad (2)
\end{aligned}
$$

where (1) holds since $\Delta(i|\psi) = 0$ for any $i \in R(\psi)$.

Let $\boldsymbol{x}' = (x_1', \dots x_n')$ be the vector such that $x_i' = 1$ if $i \in R(\psi)$, and $x_i' = x_i$ otherwise. As $x_i' \geq x_i$ for any $i \in [n]$ we have

$$OPT_A(G, k) \leq OPT_A(G, \boldsymbol{x}) \leq OPT_A(G, \boldsymbol{x}'). \quad (3)$$

Let π' be the optimal randomized adaptive policy such that $\mathbb{P}(\pi') = \boldsymbol{x}'$. Policy π' selects each node in $R(\psi)$ with probability 1, thus we can assume that such seed nodes are selected at the beginning and that the adaptive policy starts by observing the resulting partial realisation. Furthermore, we can assume that, for any partial realisation ψ', π' does not select any node $i \in R(\psi')$, otherwise there is no increase of the influence spread. Given $j \in [n]$, let $\Delta'(j)$ denote the expected increment of the influence spread when π' selects the j-th seed node (in order of selection, and without considering in the count the initial seeds of $R(\psi)$); analogously, let $\Delta'(j|i)$ denote the expected increment of the influence spread when π' selects the j-th seed node, conditioned by the fact that the j-th seed is node i.[2] We get

$$
\begin{aligned}
& OPT_A(G, \boldsymbol{x}') \\
=& \sigma(R(\psi)) + \sum_j \Delta'(j) \\
=& \sigma(R(\psi)) + \sum_j \sum_{i\in[n]\setminus R(\psi)} \mathbb{P}[\text{the } j\text{-th seed node is } i] \cdot \Delta'(j|i) \\
=& \sigma(R(\psi)) + \sum_{i\in[n]\setminus R(\psi)} \sum_j \mathbb{P}[\text{the } j\text{-th seed node is } i] \cdot \Delta'(j|i) \\
=& \sigma(R(\psi)) + \sum_{i\in[n]\setminus R(\psi)} \sum_j \mathbb{P}[\text{the } j\text{-th seed node is } i] \cdot \\
& \cdot \mathbb{E}_{\pi'}[\Delta(i|\psi') | i = \pi'(\psi') \text{ for some } \psi' \supseteq \psi \text{ observed at step } j] \\
\leq& \sigma(R(\psi)) + \sum_{i\in[n]\setminus R(\psi)} \sum_j \mathbb{P}[\text{the } j\text{-th seed node is } i] \cdot \Delta(i|\psi) \quad (4)\\
=& \sigma(R(\psi)) + \sum_{i\in[n]\setminus R(\psi)} \mathbb{P}[i \text{ is selected as seed}] \cdot \Delta(i|\psi) \\
=& \sigma(R(\psi)) + \sum_{i\in[n]\setminus R(\psi)} x_i' \cdot \Delta(i|\psi) \\
=& \sigma(R(\psi)) + \sum_{i\in[n]\setminus R(\psi)} x_i \cdot \Delta(i|\psi), \quad (5)
\end{aligned}
$$

[2] If an execution of π' requires less than j steps, we assume that the increase of the influence spread at step j (that contributes to the expected values $\Delta'(j)$ and $\Delta'(j|i)$) is null.

where (4) holds since $\Delta(i|\psi') \leq \Delta(i|\psi)$ for any partial realisation $\psi' \supseteq \psi$ by adaptive submodularity (Claim 1). By putting together (2), (3), and (5), we get

$$
\begin{aligned}
& \sigma(R(\psi)) + k \cdot \mathbb{E}_{L,\rho}\left[f(\psi_{\rho,t,L}) - f(\psi_{t-1,L}) | \psi_{t-1,L} = \psi\right] \\
=& \sigma(R(\psi)) + \sum_{i\in[n]\setminus R(\psi)} x_i \cdot \Delta(i|\psi) \\
\geq& OPT_A(G, \boldsymbol{x}') \\
\geq& OPT_A(G, k),
\end{aligned}
$$

thus showing the claim. □

Lemma 2. *If the input influence graph G is an in-arborescence, then $\sigma(R(\psi_{t-1,L})) \leq f(\psi_{t-1,L}) + OPT_N(G, t-1)$ for any live-edge graph L and $t \in [k]$.*

Armed with the above lemmas, we can now prove Theorem 1.

Proof of Theorem 1. For any $t \in [k]$, we have

$$
\begin{aligned}
& k \cdot (OPT_N(G, t) - OPT_N(G, t-1)) \\
=& k \cdot (\sigma(S_t) - \sigma(S_{t-1})) \\
=& k \cdot (\mathbb{E}_L[f(\psi_{t,L})] - \mathbb{E}_L[f(\psi_{t-1,L})]) \\
\geq& k \cdot (\mathbb{E}_{L,\rho}[f(\psi_{\rho,t,L})] - \mathbb{E}_L[f(\psi_{t-1,L})]) \quad (6)\\
=& k \cdot (\mathbb{E}_{L,\rho}[f(\psi_{\rho,t,L})] - \mathbb{E}_{L,\rho}[f(\psi_{t-1,L})]) \\
=& k \cdot \mathbb{E}_{L,\rho}[f(\psi_{\rho,t,L}) - f(\psi_{t-1,L})] \\
=& \mathbb{E}_{\psi_{t-1,L}}\left[k \cdot \mathbb{E}_{L,\rho}[f(\psi_{\rho,t,L}) - f(\psi_{t-1,L}) | \psi_{t-1,L}]\right] \\
\geq& \mathbb{E}_{\psi_{t-1,L}}[OPT_A(G, k) - \sigma(R(\psi_{t-1,L}))] \quad (7)\\
\geq& \mathbb{E}_{\psi_{t-1,L}}[OPT_A(G, k) - f(\psi_{t-1,L}) - OPT_N(G, t-1)] \quad (8)\\
=& \mathbb{E}_{\psi_{t-1,L}}[OPT_A(G, k)] - \mathbb{E}_{\psi_{t-1,L}}[f(\psi_{t-1,L})] \quad (9)\\
& - \mathbb{E}_{\psi_{t-1,L}}[OPT_N(G, t-1)] \\
=& OPT_A(G, k) - \sigma(S_{t-1}) - OPT_N(G, t-1) \\
=& OPT_A(G, k) - 2 \cdot OPT_N(G, t-1), \quad (10)
\end{aligned}
$$

where (6) holds since $dom(\psi_{t,L})$ is the optimal set of t seed nodes for the non-adaptive influence maximization problem, (7) comes from Lemma 1, and (8) comes from Lemma 2. Thus, by (10), we get $k \cdot (OPT_N(G, t) - OPT_N(G, t-1)) \geq OPT_A(G, k) - 2 \cdot OPT_N(G, t-1)$, that after some manipulations leads to the following recursive relation:

$$OPT_N(G, t) \geq \frac{1}{k} \cdot OPT_A(G, k) + \left(1 - \frac{2}{k}\right) \cdot OPT_N(G, t-1), \quad \forall t \in [k]. \quad (11)$$

By applying iteratively (11), we get

$$
\begin{aligned}
OPT_N(G, k) &\geq \frac{1}{k} \cdot \sum_{t=0}^{k-1} \left(1 - \frac{2}{k}\right)^t \cdot OPT_A(G, k) \\
&= \frac{1 - (1 - 2/k)^k}{2} \cdot OPT_A(G, k),
\end{aligned}
$$

that leads to

$$\frac{OPT_A(G, k)}{OPT_N(G, k)} \leq \frac{2}{1 - (1 - 2/k)^k} \leq \frac{2}{1 - e^{-2}} = \frac{2e^2}{e^2 - 1},$$

and this shows the claim. □

Adaptivity Gap for General Influence Graphs

In this section, we exhibit upper bounds on the k-adaptivity gap of general graphs. In the following theorem, we first give an upper bound that is linear in the number of seed nodes.

Theorem 2. *Given an arbitrary class of influence graphs $\mathcal{G}$ and $k \geq 2$, we get $AG(\mathcal{G}, k) \leq k$.*

In the next theorem we give an upper bound on the adaptivity gap that is sublinear in the number of nodes of the considered graph.

Theorem 3. *If $\mathcal{G}$ is the class of influence graphs having at most n nodes, we get $AG(\mathcal{G}) \leq \lceil n^{1/3} \rceil$.*

Let $G = (V, E, (p_{uv})_{(u,v)\in E})$ be the input influence graph. To show Theorem 3, we recall the preliminary notations considered for the proof of Theorem 1, and we give a further preliminary lemma.

Lemma 3. *Given a set $U \subseteq V$ of cardinality $h \geq k$, we have $\sigma(U) \leq \frac{h}{k} \cdot OPT_N(G, k)$.*

We use Theorem 2 and Lemma 3 to show Theorem 3.

Proof of Theorem 3. We assume w.l.o.g. that $k > \lceil n^{1/3} \rceil$ and that $OPT_N(G, k) < (\lceil n^{1/3} \rceil)^2$. Indeed, if $k \leq \lceil n^{1/3} \rceil$, by Theorem 2 the claim holds, and if $OPT_N(G, k) \geq (\lceil n^{1/3} \rceil)^2$, then $\frac{OPT_A(G,k)}{OPT_N(G,k)} \leq \frac{|V|}{OPT_N(G,k)} \leq \frac{n}{(\lceil n^{1/3} \rceil)^2} \leq \lceil n^{1/3} \rceil$, and the claim holds as well. For any $t \in [k]$, we have

$$
\begin{aligned}
& k \cdot (OPT_N(G, t) - OPT_N(G, t-1)) \\
\geq & k \cdot (\mathbb{E}_{L,\rho}[f(\psi_{\rho,t,L})] - \mathbb{E}_{L,\rho}[f(\psi_{t-1,L})]) \\
= & \mathbb{E}_{\psi_{t-1,L}}\left[k \cdot \mathbb{E}_{L,\rho}[f(\psi_{\rho,t,L}) - f(\psi_{t-1,L}) | \psi_{t-1,L}]\right] \\
\geq & \mathbb{E}_{\psi_{t-1,L}}[OPT_A(G, k) - \sigma(R(\psi_{t-1,L}))] \qquad (12) \\
= & \mathbb{E}_{\psi_{t-1,L}}[OPT_A(G, k)] - \mathbb{E}_{\psi_{t-1,L}}[\sigma(R(\psi_{t-1,L}))] \\
\geq & \mathbb{E}_{\psi_{t-1,L}}[OPT_A(G, k)] - \mathbb{E}_{\psi_{k,L}}[\sigma(R(\psi_{k,L}))] \\
\geq & \mathbb{E}_{\psi_{t-1,L}}[OPT_A(G, k)] - \mathbb{E}_{\psi_{k,L}}\left[\frac{|R(\psi_{k,L})|}{k} \cdot OPT_N(G, k)\right] \qquad (13) \\
= & OPT_A(G, k) - \frac{\mathbb{E}_{\psi_{k,L}}[|R(\psi_{k,L})|]}{k} \cdot OPT_N(G, k) \\
\geq & OPT_A(G, k) - \frac{\mathbb{E}_{\psi_{k,L}}[|R(\psi_{k,L})|]}{\lceil n^{1/3} \rceil + 1} \cdot ((\lceil n^{1/3} \rceil)^2 - 1) \qquad (14) \\
= & OPT_A(G, k) - (\lceil n^{1/3} \rceil - 1) \cdot \mathbb{E}_{\psi_{k,L}}[|R(\psi_{k,L})|] \\
= & OPT_A(G, k) - (\lceil n^{1/3} \rceil - 1) \cdot OPT_N(G, k), \qquad (15)
\end{aligned}
$$

where (12) comes from Lemma 1, (13) comes from Lemma 3, and (14) comes from the hypothesis $k > \lceil n^{1/3} \rceil$ and $OPT_N(G, k) < (\lceil n^{1/3} \rceil)^2$. By (15), we get $OPT_N(G, t) - OPT_N(G, t-1) \geq (OPT_A(G, k) - (\lceil n^{1/3} \rceil - 1) \cdot OPT_N(G, k))/k$ for any $t \in [k]$, and by summing such inequality over all $t \in [k]$, we get

$$
\begin{aligned}
& OPT_N(G, k) \\
= & \sum_{t=1}^{k} (OPT_N(G, t) - OPT_N(G, t-1)) \\
\geq & \sum_{t=1}^{k} \frac{OPT_A(G, k) - (\lceil n^{1/3} \rceil - 1) \cdot OPT_N(G, k)}{k} \\
= & OPT_A(G, k) - (\lceil n^{1/3} \rceil - 1) \cdot OPT_N(G, k). \qquad (16)
\end{aligned}
$$

Finally, (16) implies that $OPT_A(G, k) \leq \lceil n^{1/3} \rceil \cdot OPT_N(G, k)$, and this shows the claim. □

Adaptivity Gap for Other Influence Graphs

In this section, we extend the results obtained in Theorem 1, and we get upper bounds on the adaptivity gap of other classes of influence graphs. In particular, we consider the class of α-*bounded graphs*: a class of undirected graphs parametrized by an integer $\alpha \geq 0$ that includes several known graph topologies. In the following, when we refer to undirected influence graphs, we assume that, for any undirected edge $\{u, v\}$, there are two directed edges (u, v) and (v, u) having respectively two (possibly) distinct probabilities p_{uv} and p_{vu}.

α-bounded graphs. Given an undirected graph $G = (V, E)$ and a node $v \in V$, let $deg_v(G)$ be the degree of node v in graph G. Given an integer $\alpha \geq 0$, graph G is an α-*bounded graph* if $\sum_{v \in V: deg_v(G) > 2} deg_v(G) \leq \alpha$, i.e., the sum all the node degrees higher than 2 is at most α. In the following, we exhibit some interesting classes of α-bounded graphs: (i) the set of 0-bounded graphs is made of all the graphs G such that each connected component of G is either an undirected path, or an undirected cycle; (ii) if graph G is homeomorphic to a star with h edges, then G is a h-bounded graph; (iii) if graph G is homeomorphic to a parallel-link graph with h edges, then G is a $2h$-bounded graph; (iv) if graph G is homeomorphic to a cycle with h chords, then G is a $6h$-bounded graph; (v) if graph G is homeomorphic to a clique with h nodes, then G is a $h(h-1)$-bounded graph.

In the following, we provide an upper bound on the adaptivity gap of α-bounded influence graphs for any $\alpha \geq 0$.

Theorem 4. *Given $\alpha \geq 0$, let $\mathcal{G}$ be the class of α-bounded influence graphs. Then $AG(\mathcal{G}, k) \leq \min\left\{k, \frac{\alpha}{k} + 2 + \frac{1}{1-(1-1/k)^k}\right\} \leq \frac{\sqrt{4(e-1)^2\alpha + (3e-2)^2} + 3e - 2}{2(e-1)}$ for any $k \geq 2$, i.e., $AG(\mathcal{G}) \leq \sqrt{\alpha} + O(1)$.*

Let $G = (V = [n], E, (p_{uv})_{(u,v)\in E})$ be an α-bounded influence graph, and we recall the preliminary notations from Theorem 1. The proof of Theorem 4 is a non-trivial generalization of Theorem 1. In particular, the proof resorts to Theorem 2 to get the upper bound of k, and, by following the approach of Theorem 1, the following technical lemma is used in place of Lemma 2 to get the final upper bound.

Lemma 4. *When the input influence graph G is an α-bounded graph with $\alpha \geq 0$, we have that $\sigma(R(\psi_{t-1,L})) \leq f(\psi_{t-1,L}) + \left(\frac{\alpha}{k} + 2\right) \cdot OPT_N(G, k)$ for any $t \in [k]$ and live-edge graph L.*

Proof. Given a subset $U \subseteq V$, let $\partial U := \{u \in U : \exists (u,v) \in E, v \notin U\}$. We have that $\sigma(R(\psi)) \leq |R(\psi)| + \sigma(\partial R(\psi)) = f(\psi) + \sigma(\partial R(\psi))$ for any partial realisation ψ. Thus, to show the claim, it suffices to show that

$$\sigma(\partial R(\psi_{t-1,L})) \leq \left(\frac{\alpha}{k} + 2\right) \cdot OPT_N(G,k).$$

Let $U \subseteq V$ such that U has at most k connected components. Let A be the set of connected components containing at least one node of degree higher than 2, and let B be the set of the remaining components, i.e., containing nodes with degree in $[2]_0$ only. By definition of A and B, we necessarily have that $|\partial A| \leq \sum_{v \in V: deg_v(G) > 2} deg_v(G) \leq \alpha$ and $|\partial B| \leq 2k$. Thus $|\partial U| \leq |\partial A| + |\partial B| \leq \alpha + 2k$, and the next claim follows.

Claim 2. *Given a subset $U \subseteq V$ made of at most k connected components, then $|\partial U| \leq \alpha + 2k$.*

Now, we have that

$$\begin{aligned}
\sigma(\partial R(\psi_{t-1,L})) \leq & \sigma(\partial R(\psi_{k,L})) \\
\leq & \frac{|\partial R(\psi_{k,L})|}{k} \cdot OPT_N(G,k) \qquad (17) \\
\leq & \frac{\alpha + 2k}{k} \cdot OPT_N(G,k), \qquad (18)
\end{aligned}$$

where (17) comes from Lemma 3, and (18) holds since $R(\psi_{k,L})$ contains at most k connected components and because of Claim 2. Thus, by (18), the claim of the lemma follows. □

We can now prove Theorem 4.

Proof of Theorem 4. For any $t \in [k]$, we have

$$\begin{aligned}
& k \cdot (OPT_N(G,t) - OPT_N(G,t-1)) \\
\geq & k \cdot (\mathbb{E}_{L,\rho}[f(\psi_{\rho,t,L})] - \mathbb{E}_{L,\rho}[f(\psi_{t-1,L})]) \\
= & \mathbb{E}_{\psi_{t-1,L}}\left[k \cdot \mathbb{E}_{L,\rho}[f(\psi_{\rho,t,L}) - f(\psi_{t-1,L}) | \psi_{t-1,L}]\right] \\
\geq & \mathbb{E}_{\psi_{t-1,L}}[OPT_A(G,k) - \sigma(R(\psi_{t-1,L}))] \qquad (19) \\
\geq & \mathbb{E}_{\psi_{t-1,L}}\left[OPT_A(G,k) - f(\psi_{t-1,L}) - \left(\frac{\alpha}{k}+2\right) OPT_N(G,k)\right] \qquad (20) \\
= & \mathbb{E}_{\psi_{t-1,L}}[OPT_A(G,k)] - \mathbb{E}_{\psi_{t-1,L}}[f(\psi_{t-1,L})] \\
& - \left(\frac{\alpha}{k}+2\right) \cdot \mathbb{E}_{\psi_{t-1,L}}[OPT_N(G,k)] \\
= & OPT_A(G,k) - \sigma(S_{t-1}) - \left(\frac{\alpha}{k}+2\right) \cdot OPT_N(G,k) \\
= & OPT_A(G,k) - \left(\frac{\alpha}{k}+2\right) \cdot OPT_N(G,k) - OPT_N(G,t-1), \qquad (21)
\end{aligned}$$

where (19) comes from Lemma 1 and (20) comes from Lemma 4. Thus, by (21), we get the following recursive relation:

$$\begin{aligned}
OPT_N(G,t) \geq & \frac{1}{k}\left(OPT_A(G,k) - \left(\frac{\alpha}{k}+2\right) OPT_N(G,k)\right) \\
& + \left(1 - \frac{1}{k}\right) OPT_N(G,t-1), \qquad (22)
\end{aligned}$$

for any $t \in [k]$. By applying iteratively (22), we get $OPT_N(G,k) \geq \frac{1}{k} \cdot \left(OPT_A(G,k) - \left(\frac{\alpha}{k}+2\right) \cdot OPT_N(G,k)\right) \cdot \sum_{t=0}^{k-1} \left(1 - \frac{1}{k}\right)^j = \left(OPT_A(G,k) - \left(\frac{\alpha}{k}+2\right) \cdot OPT_N(G,k)\right) \cdot \left(1 - \left(1 - \frac{1}{k}\right)^k\right)$, that, after some manipulations, leads to

$$\frac{OPT_A(G,k)}{OPT_N(G,k)} \leq \frac{\alpha}{k} + 2 + \frac{1}{1-(1-1/k)^k} \leq \frac{\alpha}{k} + 2 + \frac{1}{1-e^{-1}}. \qquad (23)$$

By Theorem 2, we have that $\frac{OPT_A(G,k)}{OPT_N(G,k)} \leq k$, thus, by (23), we get

$$\begin{aligned}
\frac{OPT_A(G,k)}{OPT_N(G,k)} & \leq \min\left\{k, \frac{\alpha}{k} + 2 + \frac{1}{1-(1-1/k)^k}\right\} \\
& \leq \min\left\{k, \frac{\alpha}{k} + 2 + \frac{1}{1-e^{-1}}\right\} \\
& \leq \frac{\sqrt{4(e-1)^2\alpha + (3e-2)^2} + 3e - 2}{2(e-1)}, \qquad (24)
\end{aligned}$$

where (24) is equal to the real value of $k \geq 0$ such that $k = \frac{\alpha}{k} + 2 + \frac{1}{1-e^{-1}}$. By (24) the claim follows. □

For the particular case of 0-bounded influence graphs, the following theorem provides a better upper bound on the adaptivity gap (the proof is analogue to that of Theorem 1).

Theorem 5. *Let $\mathcal{G}$ be the class of 0-bounded influence graphs. Then $AG(\mathcal{G},k) \leq \min\left\{k, \frac{3}{1-(\max\{0,1-3/k\})^k}\right\} \leq \frac{3e^3}{e^3-1} \approx 3.16$, for any $k \geq 2$.*

Future Works

The first problem that is left open by our results is the gap between the constant lower bound provided by Chen and Peng (2019) and our upper bound on the adaptivity gap for general graphs. Besides trying to lower the upper bound, a possible direction could be that of increasing the lower bound by finding instances with a non constant adaptivity gap. Since the lower bound given in (Chen and Peng 2019) holds even when the graph is a directed path, one direction could be to exploit different graph topologies.

Although in this work we have improved the upper bound on the adaptivity gap of in-arborescence, there is still a gap between upper and lower bound, thus another open problem is to close it. It would be also interesting to find better bounds on the adaptivity gap of other graph classes, like e.g. out-arborescences. A further interesting research direction is to study the adaptivity gap of some graph classes modelling real-world networks, both theoretically and experimentally.

The study of the adaptive IM problem in the Linear Threshold model is still open, in terms of both approximation ratio and adaptivity gap. We observe that in this case the objective function is not adaptive submodular in both myopic and full-adoption feedbacks and therefore the greedy approach by Golovin and Krause (2011) cannot be applied.

The techniques introduced in this paper to relate adaptive policies with non-adaptive ones might be useful to find better upper bounds on the adaptivity gaps in different feedback models, like e.g. the myopic one, or in different graph classes.

Acknowledgements

This work was partially supported by the Italian MIUR PRIN 2017 Project "ALGADIMAR" Algorithms, Games, and Digital Markets.

References

Asadpour, A.; and Nazerzadeh, H. 2016. Maximizing Stochastic Monotone Submodular Functions. *Management Science* 62(8): 2374–2391.

Badanidiyuru, A.; Papadimitriou, C.; Rubinstein, A.; Seeman, L.; and Singer, Y. 2016. Locally adaptive optimization: Adaptive seeding for monotone submodular functions. *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms* 1: 414–429.

Borgs, C.; Brautbar, M.; Chayes, J. T.; and Lucier, B. 2014. Maximizing Social Influence in Nearly Optimal Time. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014*, 946–957. SIAM.

Cautis, B.; Maniu, S.; and Tziortziotis, N. 2019. Adaptive Influence Maximization. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD*, 3185–3186. ACM.

Chen, W.; Lakshmanan, L. V. S.; and Castillo, C. 2013. *Information and Influence Propagation in Social Networks*. Synthesis Lectures on Data Management. Morgan & Claypool Publishers.

Chen, W.; and Peng, B. 2019. On Adaptivity Gaps of Influence Maximization under the Independent Cascade Model with Full Adoption Feedback. *ISAAC* .

Chen, W.; Peng, B.; Schoenebeck, G.; and Tao, B. 2020. Adaptive Greedy versus Non-adaptive Greedy for Influence Maximization. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI*.

Chen, W.; Wang, C.; and Wang, Y. 2010. Scalable influence maximization for prevalent viral marketing in large-scale social networks. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1029–1038.

Domingos, P.; and Richardson, M. 2001. Mining the network value of customers. *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* 57–66.

Goldberg, S.; and Liu, Z. 2013. The Diffusion of Networking Technologies. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013*, 1577–1594.

Golovin, D.; and Krause, A. 2011. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *Journal of Artificial Intelligence Research* 42: 427–486.

Gupta, A.; Nagarajan, V.; and Singla, S. 2016. Algorithms and Adaptivity Gaps for Stochastic Probing. *SODA '16: Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms* 151: 1731–1747.

Gupta, A.; Nagarajan, V.; and Singla, S. 2017. Adaptivity gaps for stochastic probing: Submodular and XOS functions. *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms* 1688–1702.

Han, K.; Huang, K.; Xiao, X.; Tang, J.; Sun, A.; and Tang, X. 2018. Efficient algorithms for adaptive influence maximization. *Proceedings of the VLDB Endowment* 11(9): 1029–1040.

Kempe, D.; Kleinberg, J.; and Tardos, É. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 137–146. ACM.

Kempe, D.; Kleinberg, J.; and Tardos, É. 2015. Maximizing the spread of influence through a social network. *Theory of Computing* 11: 105–147.

Leskovec, J.; Krause, A.; Guestrin, C.; Faloutsos, C.; VanBriesen, J. M.; and Glance, N. S. 2007. Cost-effective outbreak detection in networks. In *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 420–429. ACM.

Li, Y.; Fan, J.; Wang, Y.; and Tan, K. 2018. Influence Maximization on Social Graphs: A Survey. *IEEE Transactions on Knowledge and Data Engineering* 30(10): 1852–1872.

Lowalekar, M.; Varakantham, P.; and Kumar, A. 2016. Robust influence Maximization. *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS* 1395–1396.

Mihara, S.; Tsugawa, S.; and Ohsaki, H. 2015. Influence maximization problem for unknown social networks. In *2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 1539–1546. IEEE Computer Society.

Peng, B.; and Chen, W. 2019. Adaptive Influence Maximization with Myopic Feedback. *NeurIPS* .

Richardson, M.; and Domingos, P. M. 2002. Mining knowledge-sharing sites for viral marketing. In *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 61–70. ACM.

Rubinstein, A.; Seeman, L.; and Singer, Y. 2015. Approximability of adaptive seeding under knapsack constraints. *EC 2015 - Proceedings of the 2015 ACM Conference on Economics and Computation* 797–814.

Salha, G.; Tziortziotis, N.; and Vazirgiannis, M. 2018. Adaptive submodular influence maximization with myopic feedback. *Proceedings of the 2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM 2018* 455–462.

Schoenebeck, G.; and Tao, B. 2019. Influence Maximization on Undirected Graphs: Towards Closing the (1-1/e) Gap. In Karlin, A.; Immorlica, N.; and Johari, R., eds., *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019*, 423–453. ACM.

Seeman, L.; and Singer, Y. 2013. Adaptive seeding in social networks. *Proceedings - Annual IEEE Symposium on Foundations of Computer Science, FOCS* 459–468.

Singer, Y. 2016. Influence maximization through adaptive seeding. *ACM SIGecom Exchanges* 15(1): 32–59.

Sun, L.; Huang, W.; Yu, P. S.; and Chen, W. 2018. Multi-round influence maximization. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* 2249–2258.

Tang, J.; Lakshmanan, L. V.; Huang, K.; Tang, X.; Sun, A.; Xiao, X.; and Lim, A. 2019. Efficient approximation algorithms for adaptive seed minimization. *Proceedings of the ACM SIGMOD International Conference on Management of Data* 1096–1113.

Tang, Y.; Xiao, X.; and Shi, Y. 2014. Influence Maximization : Near-Optimal Time Complexity Meets Practical Efficiency. *SIGMOD '14: Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data* 75–86.

Tong, G.; and Wang, R. 2019. Adaptive Influence Maximization under General Feedback Models. In *arXiv preprint arXiv: 1902.00192v3*.

Tong, G.; Wang, R.; Dong, Z.; and Li, X. 2020. Time-constrained Adaptive Influence Maximization. In *arXiv preprint arXiv: 2001.01742v2*.

Tong, G.; Wu, W.; Tang, S.; and Du, D. Z. 2017. Adaptive Influence Maximization in Dynamic Social Networks. *IEEE/ACM Transactions on Networking* 25(1): 112–125.

Vaswani, S.; and Lakshmanan, L. V. S. 2016. Adaptive Influence Maximization in Social Networks: Why Commit when You can Adapt? *CoRR* abs/1604.08171.

Yuan, J.; and Tang, S. 2017. No time to observe: Adaptive influence maximization with partial feedback. *IJCAI International Joint Conference on Artificial Intelligence* 3908–3914.