

Improved Approximation Factor for Adaptive Influence Maximization via Simple Greedy Strategies

Gianlorenzo D'Angelo ✉

Gran Sasso Science Institute, L'Aquila, Italy

Debashmita Poddar ✉

Gran Sasso Science Institute, L'Aquila, Italy

Cosimo Vinci ✉

Gran Sasso Science Institute, L'Aquila, Italy

Abstract

In the *adaptive influence maximization problem*, we are given a social network and a budget k , and we iteratively select k nodes, called seeds, in order to maximize the expected number of nodes that are reached by an influence cascade that they generate according to a stochastic model for influence diffusion. The decision on the next seed to select is based on the observed cascade of previously selected seeds. We focus on the *myopic feedback model*, in which we can only observe which neighbors of previously selected seeds have been influenced and on the *independent cascade* model, where each edge is associated with an independent probability of diffusing influence. While adaptive policies are strictly stronger than non-adaptive ones, in which all the seeds are selected beforehand, the latter are much easier to design and implement and they provide good approximation factors if the adaptivity gap, the ratio between the adaptive and the non-adaptive optima, is small. Previous works showed that the adaptivity gap is at most 4, and that simple adaptive or non-adaptive greedy algorithms guarantee an approximation of $\frac{1}{4} \left(1 - \frac{1}{e}\right) \approx 0.158$ for the adaptive optimum. This is the best approximation factor known so far for the adaptive influence maximization problem with myopic feedback.

In this paper, we directly analyze the approximation factor of the non-adaptive greedy algorithm, without passing through the adaptivity gap, and show an improved bound of $\frac{1}{2} \left(1 - \frac{1}{e}\right) \approx 0.316$. Therefore, the adaptivity gap is at most $\frac{2e}{e-1} \approx 3.164$. To prove these bounds, we introduce a new approach to relate the greedy non-adaptive algorithm to the adaptive optimum. The new approach does not rely on multi-linear extensions or random walks on optimal decision trees, which are commonly used techniques in the field. We believe that it is of independent interest and may be used to analyze other adaptive optimization problems. Finally, we also analyze the adaptive greedy algorithm, and show that guarantees an improved approximation factor of $1 - \frac{1}{\sqrt{e}} \approx 0.393$.

2012 ACM Subject Classification Theory of computation → Stochastic approximation; Theory of computation → Online algorithms; Theory of computation → Probabilistic computation; Theory of computation → Graph algorithms analysis

Keywords and phrases Adaptive Optimization, Influence Maximization, Submodular Optimization, Stochastic Optimization

Digital Object Identifier 10.4230/LIPIcs.ICALP.2021.59

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: <https://arxiv.org/abs/2007.09065>

Funding This work has been partially supported by the Italian MIUR PRIN 2017 Project AL-GADIMAR “Algorithms, Games, and Digital Markets”.



© Gianlorenzo D'Angelo, Debashmita Poddar, and Cosimo Vinci;
licensed under Creative Commons License CC-BY 4.0

48th International Colloquium on Automata, Languages, and Programming (ICALP 2021).

Editors: Nikhil Bansal, Emanuela Merelli, and James Worrell; Article No. 59; pp. 59:1–59:19

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



1 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 models 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, which iteratively selects a seed with maximum marginal gain, guarantees a $1 - \frac{1}{e}$ approximation factor for the IM problem [34]. Since its definition by Domingos and Richardson [19, 40] and formalization as an optimization problem by Kempe et al. [33, 34], the IM problem and its variants have been extensively investigated, motivated by applications in viral marketing [13], adoption of technological innovations [24], and outbreak or failure detection [35]. See [10, 36] for surveys on the IM problem.

Recently, Golovin and Krause [25] 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 [25] 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 an adaptive greedy algorithm achieves a $1 - \frac{1}{e}$ approximation factor for the adaptive IM problem. In their arXiv version, they retracted a claim (appeared in their previous conference version) in which they mistakenly showed that the adaptive submodularity holds even under the myopic feedback model, and this property would have guaranteed that the adaptive greedy is a constant factor approximation algorithm (under the myopic feedback model). Anyway, they conjectured that there exists a constant factor approximation algorithm for the myopic feedback model, which indeed has been found by Peng and Chen [39] who showed that both the adaptive and non-adaptive greedy algorithms guarantee a $\frac{1}{4}(1 - \frac{1}{e})$ -approximation. In particular, they showed that 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, is upper-bounded by 4. By combining this bound with the approximation factor of both the adaptive and non-adaptive greedy algorithms for the non-adaptive problem they obtain a $\frac{1}{4}(1 - \frac{1}{e})$ -approximation factor. To prove their upper-bound on the adaptivity gap, Peng and Chen use an approach that is inspired by Bradac et al. [6], which in turn is based on the approach introduced by Gupta et al. [29, 30] in the context of stochastic probing. To relate a non-adaptive solution to an optimal adaptive one, they consider the decision tree of an optimal adaptive solution and construct a non-adaptive policy by performing a random root-leaf walk in the tree, according to a probability distribution induced by the tree. Note that computing the non-adaptive policy that guarantees the upper-bound on the adaptivity gap would require to know an optimal decision tree. However, this approach only requires to show the existence of such a non-adaptive policy since it allows us to bound

the adaptivity gap and non-adaptive approximation factor is obtained by combining the non-adaptive approximation factor with the adaptivity gap. In the same paper, Peng and Chen showed that both the adaptive greedy and the non-adaptive greedy algorithms cannot achieve a factor better than $\frac{e^2+1}{(e+1)^2} < 0.607 < 1 - \frac{1}{e}$ of the adaptive optimum, and that the adaptivity gap is at least $\frac{e}{e-1} \approx 1.582$.

In [11], 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. In order to show these bounds, they followed an approach introduced by Asadpour and Nazerzadeh [3] 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, a non-constant upper bound for the adaptivity gap under the full-adoption feedback has been recently shown in [16].

Our Contribution

In this paper, we focus on the myopic model and analyze the approximation factor of the non-adaptive greedy algorithm without passing through the adaptivity gap. We show that the algorithm achieves at least a fraction of $\frac{1}{2} \left(1 - \frac{1}{e}\right) \approx 0.316$ of the adaptive optimum (Theorem 5). By definition, this implies that the adaptivity gap is at most $\frac{2e}{e-1} \approx 3.164$ (Remark 12). For both approximation ratio and adaptivity gap we obtain a substantial improvement with respect to the upper bounds obtained in [39], which are $\frac{1}{4} \left(1 - \frac{1}{e}\right) \approx 0.158$ and 4, respectively.

Non-adaptive policies are strictly weaker than adaptive ones, since the latter can implement the former by simply ignoring any kind of feedback. On the other hand, adaptive policies are difficult to implement as they require to probe suitable seeds and to observe the corresponding feedback, which can be expensive and error-prone. Moreover, they may consist of exponentially-large decision trees that are hard to compute and store. In contrast, non-adaptive policies are easy to design and implement and are independent from the feedback. In particular, the non-adaptive greedy algorithm has been extensively studied and successfully applied in the field of influence maximization. For the non-adaptive setting, several efficient implementation of the greedy algorithm have been devised that allows us to use it in large real-world networks [14, 27, 35, 37, 47, 48]. Our results show that the simple non-adaptive greedy algorithm performs well, even in the adaptive setting where we compare it with the adaptive optimum.

To show our bounds, we introduce a new approach that relate the non-adaptive greedy policy to an optimal adaptive solution. The new approach is not based on multilinear extensions and poisson processes (like, e.g. [3, 7, 9, 11]) neither on random walks on the optimal decision trees (like, e.g. [6, 29, 30, 39]), which are the main tools used so far to relate adaptive and non-adaptive policies, and to bound adaptivity gaps. Previous techniques derive adaptive approximation factors by combining non-adaptive approximation ratios with a bound on the adaptivity gap which is obtained by showing the existence of a “good” non-adaptive policy. However such a policy is hard to compute as it is usually constructed by using an optimal adaptive policy. Our approach, instead, directly analyzes a non-adaptive policy and therefore provides the exact policy that gives the desired adaptivity gap and adaptive approximation factor. We believe that our approach is of independent interest and may be used to bound approximation factors and adaptivity gaps of different adaptive optimization problems.

Our new approach consists in defining a simple randomized non-adaptive policy whose performance is not higher than that guaranteed by the greedy algorithm, and to relate such randomized non-adaptive policy with the optimal adaptive policy. In order to recover good properties of the objective function (like, e.g. submodularity) that usually guarantee good approximations when adopting greedy strategies, we introduce an artificial diffusion process in which each seed has two chances to influence its neighbours. A similar process was introduced by Peng and Chen [39], who consider a diffusion model in which the seeds appear in multiple copies of the influence graphs, so that, roughly speaking, each node has several chances to influence the neighbours, and the main machinery they consider to relate optimal adaptive strategies with optimal non-adaptive ones is that in [6]. Our direct and more refined analysis of the non-adaptive greedy algorithm improves at the same time both the approximation ratio and the adaptivity gap.

To illustrate our approach, in Section 2 we first apply our machinery to the simpler setting of adaptive monotone submodular maximization under cardinality constraint. As observed by Asadpour and Nazerzadeh [3], a constant factor approximation of the non-adaptive greedy policy applied to such setting can be obtained by combining the approximation ratio over the non-adaptive optimum and the adaptivity gap (which is equal to $(1 - \frac{1}{e})$ [3]); this leads to an approximation guarantee of $(1 - \frac{1}{e})^2 \approx 0.399$. We give a more refined analysis of the non-adaptive greedy policy and we show that its approximation ratio is at least $\frac{1}{2} (1 - \frac{1}{e^2}) \approx 0.432$. Asadpour and Nazerzadeh [3] also showed that a so-called continuous greedy policy [7, 9] achieves an approximation ratio of $1 - \frac{1}{e} - \epsilon$ (in polynomial time w.r.t. $\frac{1}{\epsilon}$). Since the continuous greedy policy is a non-adaptive policy, this bound is strictly better than ours. However, the non-adaptive greedy policy is simpler and deterministic, while the continuous greedy policy is randomized and more complex.

Finally, in Section 5, we analyze the adaptive version of the greedy algorithm applied to the adaptive influence maximization problem; by resorting again to an artificial diffusion process, we show that such adaptive algorithm guarantees an approximation ratio of $1 - \frac{1}{\sqrt{e}} \approx 0.393$; thus, we further improve the upper bound shown for the non-adaptive greedy algorithm, and we also give a more refined analysis of the adaptive greedy algorithm than that of Peng and Chen [39], who showed a $\frac{1}{4} (1 - \frac{1}{e}) \approx 0.158$ approximation factor.

Due to the lack of space, some proofs are deferred to the full version of the paper [15].

Related Work

Adaptive Influence Maximization. The adaptive influence maximization problem under the independent cascade model has been studied by [11, 12, 31, 39, 45, 46, 49, 51, 52, 53]. These include studies on several classes of graphs and different feedback models.

The most studied feedback model is the full-adoption feedback, in which the entire influence spread generated by each selected node is observed. Golovin and Krause [25] show that the full-adoption feedback model satisfies the adaptive submodularity property and, by exploiting such property, provide a $(1 - 1/e)$ approximation adaptive algorithm for the problem of finding the best adaptive policy. Chen and Peng [11] study the adaptivity gap in the full-adoption feedback model under certain restrictions on the graph topology. In particular, they show that the adaptivity gap of in-arborescence (resp. out-arborescence) graphs belongs to $[e/(e-1), 2e/(e-1)]$ (resp. $[e/(e-1), 2]$). For general graphs with n nodes (resp. in-arborescence graphs), an upper bound of $O(n^{1/3})$ (resp. $2e^2/(e^2-1)$) for the adaptivity gap under the full-adoption feedback has been recently shown in [16].

Golovin and Krause [25] conjectured that the influence maximization problem under the myopic feedback model admits a constant approximation algorithm and a constant adaptivity gap, despite the adaptive submodularity does not hold under such feedback model. Since

then, several studies have been conducted on the myopic feedback model. Some recent works include that of Salha et al. [42], in which they consider a modified version of the independent cascade model which gives multiple chances to the seeds to activate their neighbours, and consider a different utility function which needs to be maximized. They demonstrate that the myopic feedback model is adaptive submodular under such modified diffusion model, and provide an adaptive greedy policy that achieves a $1 - 1/e$ approximation ratio for the problem of finding the best adaptive policy. The work of Peng and Chen [39] is the first one that provides a constant upper bound on the adaptivity gap under the myopic feedback model. They introduce a policy in which each seed appears in multiple copies of the original graph; furthermore, this hybrid policy connects the adaptive and the non-adaptive policies via a machinery used by [29, 30, 6] in the context of stochastic probing. They show that the adaptivity gap is in $[e/(e-1), 4]$, and the upper bound is turned into $(1 - 1/e)/4$ approximation algorithm.

Other diffusion and feedback models have been also studied, e.g., the multi-round diffusion model [45], the general feedback model [49], and the partial feedback model [53]. Han et al. [31] conduct a study on the batch selection of seeds at each step of the diffusion process. Tong et al. [50] introduce the dynamic independent cascade model, which captures the dynamic nature of real-world social networks. Finally, Singer et al. [5, 41, 43, 44], in their line of research on adaptivity gaps, studied a two-stage process called adaptive seeding, that exhibits some similarities with the influence maximization problem under the myopic feedback model.

Other Topics in Adaptive Optimization. Beyond influence maximization problems, the adaptive optimization and the adaptivity gap have been generally studied for many other stochastic settings [1, 2, 3, 4, 6, 8, 17, 18, 20, 21, 22, 23, 25, 26, 28, 29, 30, 32, 38].

A general adaptive optimization framework deals with the fact that an item will reveal its actual state only when it has been irrevocably included in the final solution, and the main goal is to optimize an objective function under such uncertainty. Stochastic variants of packing integer programs, 0/1 knapsack problems, and covering problems, have been studied under the perspective of adaptive optimization in [17, 18, 23], respectively.

Asadpour et al. [3, 4] study the adaptivity gap of the stochastic submodular maximization problem under a matroid constraint, in which the goal is to select a subset of items satisfying a matroid constraint, that maximizes the value of a monotone submodular value function defined on the random states of the selected items. In [3], the authors consider an adaptive greedy policy (often denoted as myopic policy) to approximate the optimal value of the best adaptive policy, and they show that it achieves a $1/2$ approximation ratio under general matroid constraint, and a $1 - \frac{1}{e}$ approximation ratio under cardinality constraint. Interesting variants or extensions of the above optimization problem have been considered in [6, 29, 30].

Chan and Farias [8] study the efficiency of adaptive greedy policies applied to a general class of stochastic optimization problems, called stochastic depletion problems, in which the adaptive policy, at each step, chooses an action that generates a reward and depletes some of the items; they show that, under certain structural properties, a simple adaptive greedy policy guarantees a constant factor approximation of the best adaptive policy. Hellerstein et al. [32] use an optimal decision tree to build a connection between the adaptive and the non-adaptive policies, and show that the adaptivity gap of stochastic submodular maximization under cardinality constraint is $1 - \frac{1}{e^\tau}$, where τ is the minimum value of the probability that an item is in some state.

Organization of the Paper

In the next section we introduce our new approach by applying it to a simpler setting. In Section 3 we introduce the adaptive influence maximization problem along with the necessary notation and definitions. In Section 4 we give the main results of the paper, that is the improved approximation factor for the non-adaptive greedy algorithm and the upper bound on the adaptivity gap for adaptive influence maximization. In Section 5 we show an improved approximation ratio for the adaptive greedy algorithm. In Section 6 we outline possible research directions.

2 Overview of the Approach: Stochastic Submodular Maximization

In this section, we illustrate our machinery by applying part of it to the problem of maximizing a stochastic submodular set function under cardinality constraints [3].

For two integers h and k , $h \leq k$, let $[k]_h := \{h, h + 1, \dots, k\}$ and $[k] := [k]_1$. A function $f : \mathbb{R}_{\geq 0}^n \rightarrow \mathbb{R}_{\geq 0}$ is a *monotone submodular value function* if, for any vectors $x, y \in \mathbb{R}_{\geq 0}^n$, we get $f(x \vee y) + f(x \wedge y) \leq f(x) + f(y)$, where $x \wedge y$ denotes the componentwise minimum and $x \vee y$ denotes the componentwise maximum.

Let $[n]$ be a finite set of n items, and let $\theta = (\theta_1, \dots, \theta_n)$ be a vector of n real, non-negative, and independent *state* random variables, where each θ_i returns the state $\theta_i \in \mathbb{R}_{\geq 0}$ associated to each item i following a certain probability distribution. Let $f : \mathbb{R}_{\geq 0}^n \rightarrow \mathbb{R}_{\geq 0}$ be the *objective function*, that is a monotone submodular value function. For any $S \subseteq [n]$, let $\theta(S) := (\theta_1(S), \dots, \theta_n(S))$ be the *partial state* random variable, that is a random vector defined as $\theta_i(S) = \theta_i$ if $i \in S$, $\theta_i(S) = 0$ otherwise. With a little abuse of notation, we assume that vector $\theta(S)$ gives also information on the set S which $\theta(S)$ is based on.

For a given integer $k \geq 0$, we aim at selecting a subset $S \subseteq [n]$ subject to cardinality constraint $|S| = k$, that maximizes the expected value $\mathbb{E}_\theta[f(\theta(S))]$. To guarantee a (possibly) better solution we can resort to an *adaptive policy*, that, at each step, observes the partial state $\xi \sim \theta(U)$, where U denotes the set of items previously selected, and selects another further item $\pi(\xi) \in [n] \setminus U$; after k iterations, the policy returns a set $U_{\theta,k}(\pi) \subseteq [n]$ with $|U_{\theta,k}(\pi)| = k$, which is a random set depending on the state of θ . The *stochastic monotone submodular maximization problem* (SMSM) takes as input a set of items $[n]$, a random vector θ , a monotone submodular value function f , and an integer $k \in [n]$, and asks to find an adaptive policy π that maximizes the expected value $\mathbb{E}_\theta[f(\theta(U_{\theta,k}(\pi)))]$.

In general, computing an optimal adaptive strategy is a computationally hard problem. Furthermore, in many contexts it is difficult to implement adaptive strategies, and non-adaptive strategies (in which the solution is chosen without observing the states of the random variables) is a more feasible choice. Asadpour and Nazerzadeh [3] show that a non-adaptive randomized continuous greedy algorithm guarantees a $(1 - \frac{1}{e} - \epsilon)$ approximation for the SMSM problem (in polynomial time w.r.t. $1/\epsilon$); however, the proposed approach resorts to a quite sophisticated randomized algorithm. They also consider a simpler and deterministic *non-adaptive greedy algorithm* as approximation algorithm, that starts from an empty set $S := \emptyset$ and, at each iteration $t \in [k]$, adds in S the item $i \in [n] \setminus S$ maximizing $\mathbb{E}_\theta[f(\theta(S \cup \{i\}))]$. They show that the greedy algorithm guarantees an approximation factor of $\frac{1}{2}(1 - \frac{1}{e}) \approx 0.316$ if the chosen subsets are subject to a matroid constraint, that can be reduced to $(1 - \frac{1}{e})^2 \approx 0.399$ when considering a cardinality constraint $|S| \leq k$ only (i.e., uniform matroid constraint).

In the following theorem, we give a better analysis of the greedy algorithm under cardinality constraints, and we show that the approximation factor increases to $\frac{1}{2}(1 - \frac{1}{e^2}) \approx 0.432$.

► **Theorem 1.** *The non-adaptive greedy algorithm is a $\frac{1}{2} (1 - \frac{1}{e^2})$ approximation algorithm for the SMSM problem.*

For the proof of Theorem 1 we relate the non-adaptive greedy solution with the optimal adaptive solution. We assume w.l.o.g that $k \geq 2$, otherwise the approximation ratio is equal to 1. For any $t \in [k]_0$, let S_t be the set of t items computed by the greedy algorithm at iteration t (where $S_0 := \emptyset$). Let π be an optimal adaptive policy, and let $x = (x_1, \dots, x_n) \in [0, 1]^n$ be the vector such that x_i is the probability that node $i \in [n]$ is selected by policy π . Let $OPT_A(k)$ denote the value of policy π (i.e., the optimal value of the problem). Given $i \in [n]$, let $e^i := (e_1^i, \dots, e_n^i)$ be a random vector where e_j^i has the same distribution of θ_j if $i = j$, and $e_j^i = 0$ otherwise; given a partial state ξ , let $\Delta(i|\xi) := \mathbb{E}_{e^i}[f(\xi \vee e^i) - f(\xi)]$, i.e., the expected increment of the objective function when adding an item i under partial state ξ .

For any $t \in [k-1]_0$, as intermediate step of our analysis, we consider a *randomized non-adaptive policy* Rand_t that, starting from the greedy solution S_t , computes a random set $S_{\rho,t} := S_t \cup \{\rho\}$, where $\rho \in [n]$ is a random item such that $\mathbb{P}[\rho = i] = x_i/k$ for any $i \in [n]$ and selected independently from any other event. Observe that the above random variable is well-defined, as $\sum_{v \in V} (x_v/k) = k/k = 1$; furthermore, we observe that the expected value of f under Rand_t is $\mathbb{E}_{\theta, \rho}[f(\theta(S_{\rho,t}))]$. Furthermore, for any $t \in [k-1]_0$, we consider a *hybrid adaptive policy* Hyb_t , that first runs the adaptive policy π , and then merges the items of S_t with the items of $U_{\hat{\theta},k}(\pi)$ selected by π , where $\hat{\theta}$ is a random state that follows the same distribution of θ but is independent from θ ; finally, the expected value of f under Hyb_t is defined as $\mathbb{E}_{\theta, \hat{\theta}}[f(\theta(S_t) \vee \hat{\theta}(U_{\hat{\theta},k}(\pi) \cup S_t))]$. A similar hybrid adaptive policy has been also considered by Asadpour and Nazerzadeh [3], but in place of the randomized non-adaptive policy considered in our work, they resort to a non-adaptive strategy defined by a Poisson process based on the multilinear extension of the expected value function $\mathbb{E}_{\theta}[f(*)]$. Before showing the theorem, we give some preliminary lemmas, which relate the randomized non-adaptive policy with the hybrid adaptive policy.

► **Lemma 2.** *We have that $\mathbb{E}_{\theta, \rho}[f(\theta(S \cup \{\rho\}))] - \mathbb{E}_{\theta}[f(\theta(S))] = \sum_{i \in [n] \setminus S} x_i \cdot \mathbb{E}_{\theta}[\Delta(i|\theta(S))]$ for any $S \subseteq [n]$.*

► **Lemma 3.** *For any $S \subseteq [n]$, we have $\mathbb{E}_{\theta, \hat{\theta}}[f(\theta(S) \vee \hat{\theta}(U_{\hat{\theta},k}(\pi) \cup S))] \leq \mathbb{E}_{\theta, \hat{\theta}}[f(\theta(S) \vee \hat{\theta}(S))] + \sum_{i \in [n] \setminus S} x_i \cdot \mathbb{E}_{\theta}[\Delta(i|\theta(S))]$.*

► **Lemma 4.** *We have that $OPT_A(k) \leq \mathbb{E}_{\theta, \hat{\theta}}[f(\theta(S) \vee \hat{\theta}(U_{\hat{\theta},k}(\pi) \cup S))]$ for any $S \subseteq [n]$.*

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

Proof of Theorem 1. For any $t \in [k]_0$, let $GR_N(t) := \mathbb{E}_{\theta}[f(\theta(S_t))]$ denote the expected value of f obtained at the t -th iteration of the non-adaptive greedy algorithm. We have that

$$GR_N(t+1) - GR_N(t) = \max_{v \in V} [\mathbb{E}_{\theta}[f(\theta(\{v\} \cup S_t))] - \mathbb{E}_{\theta}[f(\theta(S_t))] \quad (1)$$

$$\begin{aligned} &\geq \overbrace{\mathbb{E}_{\theta, \rho}[f(\theta(S_t \cup \{\rho\}))]}^{\text{Exp. value of Rand}_t} - \mathbb{E}_{\theta}[f(\theta(S_t))] \\ &\geq \frac{1}{k} \cdot \sum_{i \in [n] \setminus S_t} x_i \cdot \mathbb{E}_{\theta}[\Delta(i|\theta(S_t))] \end{aligned} \quad (2)$$

$$\geq \frac{1}{k} \cdot \left(\overbrace{\mathbb{E}_{\theta, \hat{\theta}}[f(\theta(S_t) \vee \hat{\theta}(U_{\hat{\theta},k}(\pi) \cup S_t))]}^{\text{Exp. value of Hyb}_t} - \mathbb{E}_{\theta, \hat{\theta}}[f(\theta(S_t) \vee \hat{\theta}(S_t))] \right) \quad (3)$$

$$\geq \frac{1}{k} \cdot \left(OPT_A(k) - \mathbb{E}_{\theta, \hat{\theta}}[f(\theta(S_t) \vee \hat{\theta}(S_t))] \right) \quad (4)$$

$$\geq \frac{1}{k} \cdot OPT_A(k) - \frac{1}{k} \cdot \left(\mathbb{E}_{\theta, \hat{\theta}} [f(\theta(S_t)) + f(\hat{\theta}(S_t))] \right) \quad (5)$$

$$\begin{aligned} &\geq \frac{1}{k} \cdot OPT_A(k) - \frac{1}{k} \cdot \left(\mathbb{E}_{\theta} [f(\theta(S_t))] + \mathbb{E}_{\hat{\theta}} [f(\hat{\theta}(S_t))] \right) = \frac{1}{k} \cdot OPT_A(k) - \frac{2}{k} \cdot \mathbb{E}_{\theta} [f(\theta(S_t))] \\ &= \frac{1}{k} \cdot OPT_A(k) - \frac{2}{k} \cdot GR_N(t), \end{aligned} \quad (6)$$

where (1) comes from the fact that the greedy strategy, at iteration $t+1$, adds to S_t the item i guaranteeing the best expected value of f , (2) comes from Lemma 2, (3) comes from Lemma 3, (4) comes from Lemma 4, and (5) holds since f is a monotone submodular value function (as $f(\theta(S_t) \vee \hat{\theta}(S_t)) \leq f(\theta(S_t) \vee \hat{\theta}(S_t)) + f(\theta(S_t) \wedge \hat{\theta}(S_t)) \leq f(\theta(S_t)) + f(\hat{\theta}(S_t))$). Thus, by (6) and some manipulations, we get $GR_N(t+1) \geq \frac{1}{k} \cdot OPT_A(k) + (1 - \frac{2}{k}) \cdot GR_N(t)$ for any $t \in [k-1]_0$. By applying iteratively the above inequality, we get $GR_N(k) \geq \frac{1}{k} \cdot \sum_{t=0}^{k-1} (1 - \frac{2}{k})^t \cdot OPT_A(k) = \frac{1}{2} \left(1 - (1 - \frac{2}{k})^k \right) \cdot OPT_A(k)$, that leads to $\frac{GR_N(k)}{OPT_A(k)} \geq \frac{1}{2} \left(1 - (1 - \frac{2}{k})^k \right) \geq \frac{1}{2} \left(1 - \frac{1}{e^2} \right)$, and this shows the claim. \blacktriangleleft

3 Adaptive Influence Maximization under the Myopic Feedback Model: Preliminaries

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 edges are directed and $p_{uv} \in [0, 1]$ is an *activation probability* associated to each edge $(u, v) \in E$. Given a set of *seeds* $S \subseteq V$ 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. The size of $\bigcup_{t \leq r} A_t$, i.e. the number of nodes activated/reached by the diffusion process, is the *influence spread*.

The above diffusion process can be equivalently defined as follows. The *live-edge graph* $\mathbf{L} = (V, \mathbf{L}(E))$ is a random graph made from G , such that each edge $(u, v) \in E$ is included in $\mathbf{L}(E)$ with probability p_{uv} , independently from the other edges, i.e., $\mathbb{P}[\mathbf{L} = L] = \prod_{(u,v) \in L} p_{uv} \prod_{(u,v) \in E \setminus L} (1 - p_{uv})$. With a little abuse of notation, we may denote $\mathbf{L}(E)$ with \mathbf{L} . Given $L \subseteq E$, let $R_L(S) := \{v \in V : \text{there exists a path from } u \text{ to } v \text{ in } L \text{ for some } u \in S\}$, i.e., the set of nodes reached by nodes in S in the graph L . Informally, if S is the set of seeds, and L is a realisation of the live-edge graph, $R_L(S)$ equivalently denotes the set of nodes which are reached/activated by the above diffusion process. Let $\sigma_L(S) := |R_L(S)|$ denote the *influence spread* generated by the set of seeds S if the realised live-edge graph is L , and let $\sigma(S) := \mathbb{E}_{\mathbf{L}}[\sigma_{\mathbf{L}}(S)]$ be the *expected influence spread* generated by S .

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 seeds $S \subseteq V$ with $|S| \leq k$ such that $\sigma(S)$ is maximized. Without loss of generality, we assume that $k \in [n]$ and, since the objective function is monotone, $|S| = k$ for any solution S .

Kempe et al. [33, 34] showed that function σ is monotone and submodular, therefore the following greedy algorithm achieves a $1 - \frac{1}{e}$ approximation factor: (i) start with an empty set of seeds $S := \emptyset$; (ii) at each iteration $t \in [k]$, add to S the node v that maximizes the expected influence spread $\sigma(S \cup \{v\})$. Note that the greedy algorithm requires at each

iteration to compute the value of function σ , for some set of seeds and this has been shown to be computationally intractable as it is $\#P$ -hard [13]. However, standard Chernoff bounds allows us estimate the value of σ through a polynomial number of Monte-Carlo simulations by introducing an arbitrarily small additive error $\epsilon > 0$, which depends on the number of simulations [34]. In the remainder of the paper we will omit the additional term ϵ to avoid unnecessary complicated formulas. We will refer to this algorithm as the *non-adaptive greedy algorithm*.

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

An adaptive policy under the myopic feedback model is formally defined as follows. Given $L \subseteq E$, the *realisation* $\phi_L : V \rightarrow 2^V$ associated to L assigns to each node $v \in V$ the value $\{z \in V : (v, z) \in L\} \cup \{v\}$, i.e., the set containing v and the neighbours activated by seed v when $\mathbf{L} = L$. Let Φ denote the *random realisation*, i.e., the random variable such that $\mathbb{P}[\Phi = \phi_L] = \mathbb{P}[\mathbf{L} = L]$ for any $L \subseteq E$. Given a set $S \subseteq V$, a *partial realisation* $\psi : S \rightarrow 2^V$ is the restriction to S of the domain of some realisation, i.e., there exists $L \subseteq E$ 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 , and let $Im(\psi) := \bigcup_{v \in dom(\psi)} \psi(v)$. A partial realisation ψ' is a *sub-realisation* of a partial realisation 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 any partial realisation ψ can be equivalently represented as $\{(v, \phi_L(v)) : v \in dom(\psi)\}$ for some $L \subseteq E$.

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 seeds, e.g., by returning a string $\pi(\psi) := STOP$. In particular, an adaptive policy π can be run as follows: (i) start from an empty realisation $\psi := \emptyset$; (ii) if $\pi(\psi) \neq STOP$ set $\psi \leftarrow \psi \cup \{(v, \phi_L(v))\}$ and repeat (ii) until $\pi(\psi) = STOP$; (iii) at the end, return $\psi_\pi := \psi$. Let Ψ_π be the random partial realisation returned by the execution of policy π . The *expected influence spread* of an adaptive policy π is defined as $\sigma(\pi) := \mathbb{E}_{\mathbf{L}}[\sigma_{\mathbf{L}}(dom(\Psi_\pi))]$, i.e., it is the expected value of the number of nodes reached by the diffusion process at the end of the execution of policy π . We say that $|\pi| = k$ if policy π always return a partial realisation ψ_π with $|dom(\psi_\pi)| = k$. The *adaptive influence maximization problem (under IC and myopic feedback)* is the computational problem that, given an influence graph G and $k \in [n]$, asks to find an adaptive policy π subject to $|\pi| = k$ that maximizes the expected influence spread $\sigma(\pi)$.

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 an integer $k \in [n]$, the *k-adaptivity gap* of G is defined as $AG(G, k) := \frac{OPT_A(G, k)}{OPT_N(G, k)}$, and measures how much an adaptive policy outperforms a non-adaptive solution for the influence maximization problem applied to influence graph G , when the maximum number of seeds is k . The *adaptivity gap* of G is defined as $AG(G) := \sup_{k \in [n]} AG(G, k)$. We observe that for $k = 1$ the k -adaptivity gap and the approximation factors are trivially equal to 1, thus we omit such case in the following.

4 The Efficiency of the Non-adaptive Greedy Algorithm

In this section, we show that a simple non-adaptive algorithm guarantees an approximation ratio of $\frac{1}{2} \left(1 - \frac{1}{e}\right) \approx 0.316$ for the adaptive influence maximization problem, thus improving the approximation ratio of $\frac{1}{4} \left(1 - \frac{1}{e}\right) \approx 0.158$ given in [39]. The algorithm provided by [39] is the usual *non-adaptive greedy algorithm* given in [34] and reported in the previous section. We observe that such algorithm is non-adaptive, i.e., despite it is used for adaptive optimization, does not resort to the use of any adaptive policy and all the seeds are selected without observing any partial realisation.

► **Theorem 5.** *Given an influence graph G with n nodes and $k \in [n]_2$, the non-adaptive greedy algorithm is a $\frac{1}{2} \left(1 - \left(1 - \frac{1}{k}\right)^k\right) \geq \frac{1}{2} \left(1 - \frac{1}{e}\right)$ approximation algorithm for the adaptive influence maximization problem (under IC and myopic feedback) applied to (G, k) .*

In the proof of Theorem 5 (see Subsection 4.4) we relate the expected influence spread coming from the non-adaptive greedy algorithm with that of the optimal adaptive policy. We first need some notation and preliminary results. Let $G = (V = [n], E, (p_{uv})_{(u,v) \in E})$ be an influence graph, and let $k \in [n]_2$. For any $t \in [k]_0$, let S_t denote the set of the first t seeds selected by the greedy algorithm, so that $\sigma(S_k)$ is the expected influence spread of the solution returned by the algorithm. Let π be an optimal adaptive policy, and let $x = (x_1, \dots, x_n)$ be the vector such that x_i is the probability that node i is selected by π . As in Section 2, we resort again to a randomized non-adaptive policy to relate the expected influence spread of the greedy algorithm with that of the adaptive policy.

Randomized Non-adaptive Policy. For any $t \in [k-1]_0$, as intermediate step of our analysis, we consider again the randomized non-adaptive policy Rand_t defined in Section 2: starting from the greedy solution S_t , we compute a random set $S_{\rho,t} := S_t \cup \{\rho\}$, where $\rho \in [n]$ is a random item such that $\mathbb{P}[\rho = i] = x_i/k$ for any $i \in [n]$ and selected independently from any other event. We observe that the expected value of f under Rand_t is $\mathbb{E}_{\rho}[\sigma(S_{\rho,t})]$.

As a support of our analysis, we also define a new diffusion model and a hybrid adaptive policy. In particular, the new diffusion model will be used to recover certain properties connected with the submodularity that, as in Section 2, will allow us to relate the expected influence spread of the randomized non-adaptive policy with that of the hybrid adaptive policy. Furthermore, following again the approach of Section 2, the hybrid adaptive policy is obtained by combining the greedy (non-adaptive) solution and the optimal adaptive one, and it will be used in our analysis to get an upper bound on the optimal adaptive spread.

2-level Diffusion Model. In the 2-level diffusion model each selected seed $u \in V$ has two chances to influence its neighbours, and all the non-seeds have one chance only, i.e., the activation probability for all the edges (u, v) is $1 - (1 - p_{u,v})^2$ if u is a seed, and $p_{u,v}$ otherwise. More formally, let $\hat{\mathbf{L}}$ be a live-edge graph distributed as \mathbf{L} and independent from \mathbf{L} . Given a set of seeds S , the *2-level live-edge graph* can be defined as $\mathbf{L}^2(S) := \mathbf{L} \cup \{\hat{\mathbf{L}} \cap \{(u, v) \in E : u \in S\}\}$. Given a set of nodes S , let $\sigma_{\mathbf{L}, \hat{\mathbf{L}}}^2(S) := \sigma_{\mathbf{L}^2(S)}(S)$ denote the influence spread induced by S in live-edge graph $\mathbf{L}^2(S)$, and let $\sigma^2(S) := \mathbb{E}_{\mathbf{L}, \hat{\mathbf{L}}}[\sigma_{\mathbf{L}, \hat{\mathbf{L}}}^2(S)]$ denote the expected influence spread induced by S under the 2-level diffusion model. Let Φ and $\hat{\Phi}$ denote the random realisations associated to live-edge graphs \mathbf{L} and $\hat{\mathbf{L}}$, respectively.

2-Level Hybrid Adaptive Policy. For any $t \in [k-1]_0$, let Hyb_t^2 be a hybrid adaptive policy defined as follows: (i) Hyb_t^2 selects all the nodes in S_t as seeds; (ii) then, Hyb_t^2 adds to S_t all the nodes that the optimal adaptive policy π would have select when starting from the

empty realisation, and observing, at each step, partial realisations coming from the live-edge graph $\hat{\mathbf{L}}$ only; in other words, for any seed $v \in V$ selected by the policy, the new set of nodes that the policy observes (to choose the next seed) is $\hat{\Phi}(v)$; (iii) finally, denoting with $\hat{\Psi}_\pi$ the random realisation returned by π , the expected influence spread of Hyb_t^2 is defined as $\mathbb{E}_{\mathbf{L}, \hat{\mathbf{L}}}[\sigma_{\mathbf{L}, \hat{\mathbf{L}}}^2(\text{dom}(\hat{\Psi}_\pi) \cup S_t)]$, i.e., it is the expected influence spread determined $\text{dom}(\hat{\Psi}_\pi) \cup S_t$, according to the 2-level live-edge graph $\mathbf{L}^2(\text{dom}(\hat{\Psi}_\pi) \cup S_t)$.

In what follows we give some technical results which are based on the above definitions and will be used in Subsection 4.4 to show the main theorem. In particular, we show that the 2-level diffusion model satisfies certain properties connected with submodular set functions (see Subsection 4.1); then, we use such properties to relate the expected influence spread in the ordinary diffusion model to that of the 2-level diffusion model (see Subsection 4.2); finally, by using a similar approach as in Section 2, we use the above relations to relate the efficiency of the randomized non-adaptive policy with that of the optimal adaptive policy, via the 2-level hybrid adaptive policy (see Subsection 4.3).

4.1 Adaptive Submodularity in the 2-Level Diffusion Model

The adaptive submodularity [25] is a property that extends the well-known concept of submodularity to the adaptive framework, and allows us to design efficient adaptive approximation algorithms. In particular, the adaptive submodularity states that, given two subrealisations $\psi \subseteq \psi'$ and a node $v \in V$, adding node v under partial realisation ψ' causes an expected increment of the influence spread that is not higher than that caused under partial realisation ψ . Unfortunately, as shown in the arXiv version of [25], the myopic feedback model, in general, does not satisfy the adaptive submodularity.

Anyway, by resorting to the 2-level diffusion model, we can recover a similar property as the adaptive submodularity. Given $S \subseteq V$, a partial realisation $\hat{\psi}$, and $v \in V$, let

$$\begin{aligned} \Delta_S^2(v|\hat{\psi}) \\ := \mathbb{E}_{\mathbf{L}, \hat{\mathbf{L}}} \left[\sigma_{\mathbf{L}^2(\{v\} \cup \text{dom}(\hat{\psi}) \setminus S)}(\{v\} \cup S \cup \text{dom}(\hat{\psi})) - \sigma_{\mathbf{L}^2(\text{dom}(\hat{\psi}) \setminus S)}(S \cup \text{dom}(\hat{\psi})) \mid \hat{\psi} \subseteq \hat{\Phi} \right] \end{aligned}$$

denote the expected increment of the influence spread w.r.t. the 2-level diffusion model when adding seed v to the set of nodes $S \cup \text{dom}(\hat{\psi})$, but assuming that the nodes in S have a unique chance to influence their neighbors, and that partial realisation $\hat{\psi}$ has been observed. We say that the 2-level diffusion model is *adaptive submodular* if, for any $S \subseteq V$, any partial realisations $\hat{\psi}, \hat{\psi}'$ with $\hat{\psi} \subseteq \hat{\psi}'$, and any $v \in V$, we have that $\Delta_S^2(v|\hat{\psi}) \geq \Delta_S^2(v|\hat{\psi}')$.

► **Lemma 6.** *The 2-level diffusion model is adaptive submodular.*

4.2 From the Ordinary to the 2-level Diffusion Model

In this subsection, we see how to relate the 2-level diffusion model to the ordinary one.

► **Lemma 7.** *We have that $\sigma^2(S) \leq 2 \cdot \sigma(S)$ for any $S \subseteq V$.*

Now, given a set $S \subseteq V$ and $v \in V$, let $\Delta_S(v) := \sigma(S \cup \{v\}) - \sigma(S)$, that is the expected increment under the ordinary diffusion model when adding a new node to a set S , and let $\Delta_S^2(v) := \mathbb{E}_{\mathbf{L}, \hat{\mathbf{L}}}[\sigma_{\mathbf{L}^2(\{v\})}(S \cup \{v\}) - \sigma_{\mathbf{L}}(S)]$, that is the above expected increment, with the further assumption that v has two chances to influence its neighbors.

► **Lemma 8.** *We have that $\Delta_S^2(v) \leq 2 \cdot \Delta_S(v)$ for any $S \subseteq V$ and $v \in V$.*

4.3 From the Adaptive to the Randomized Non-adaptive Policy

The following three lemmas (Lemma 9, Lemma 10, and 11) relate the randomized non-adaptive policy with the 2-level hybrid adaptive policy, and then with the optimal adaptive policy of the ordinary diffusion model. In particular, the proofs of Lemma 9, Lemma 10, and 11, resort to a similar approach of the proofs of Lemma 2, Lemma 3, and Lemma 4 of Section 2.

► **Lemma 9.** *For any $S \subseteq V$, we have that $k \cdot (\mathbb{E}_\rho[\sigma(\{\rho\} \cup S)] - \sigma(S)) \geq \sum_{v \in V \setminus S} x_v \cdot \Delta_S(v)$.*

The following lemma relates the adaptive setting with the non-adaptive one, and its proof resorts to the adaptive submodularity defined in Subsection 4.1.

► **Lemma 10.** *We have $\mathbb{E}_{\mathbf{L}, \hat{\mathbf{L}}}[\sigma_{\mathbf{L}, \hat{\mathbf{L}}}^2(\text{dom}(\hat{\Psi}_\pi) \cup S)] \leq \sigma^2(S) + \sum_{v \in V \setminus S} x_v \cdot \Delta_S^2(v)$ for any $S \subseteq V$.*

The following lemma shows that the optimal adaptive influence spread is upper bounded by that expected influence spread of the 2-level hybrid adaptive policy.

► **Lemma 11.** *We have $OPT_A(G, k) \leq \mathbb{E}_{\mathbf{L}, \hat{\mathbf{L}}}[\sigma_{\mathbf{L}, \hat{\mathbf{L}}}^2(\text{dom}(\hat{\Psi}_\pi) \cup S)]$ for any $S \subseteq V$.*

4.4 Proof of Theorem 5

Armed with the above results, we can now prove Theorem 5. For any $t \in [k]_0$, let $GR_N(G, t) := \sigma(S_t)$ denote the expected influence spread $\sigma(S_t)$ obtained when the first t seeds have been selected by the greedy algorithm. We have that

$$\begin{aligned} & GR_N(G, t+1) - GR_N(G, t) \\ &= \max_{v \in V} [\sigma(\{v\} \cup S_t)] - \sigma(S_t) \\ &\geq \overbrace{\mathbb{E}_\rho[\sigma(\{\rho\} \cup S_t)]}^{\text{Exp. value of Rand}_t} - \sigma(S_t) \end{aligned} \tag{7}$$

$$\geq \frac{1}{k} \sum_{v \in V \setminus S_t} x_v \cdot \Delta_{S_t}(v) \tag{8}$$

$$\geq \frac{1}{2k} \sum_{v \in V \setminus S_t} x_v \cdot \Delta_{S_t}^2(v) \tag{9}$$

$$\geq \frac{1}{2k} \overbrace{(\mathbb{E}_{\mathbf{L}, \hat{\mathbf{L}}}[\sigma_{\mathbf{L}, \hat{\mathbf{L}}}^2(\text{dom}(\hat{\Psi}_\pi) \cup S_t)] - \sigma^2(S_t))}^{\text{Exp. value of Hyb}_t^2} \tag{10}$$

$$\geq \frac{1}{2k} (OPT_A(G, k) - \sigma^2(S_t)) \tag{11}$$

$$\geq \frac{1}{2k} (OPT_A(G, k) - 2 \cdot \sigma(S_t)) \tag{12}$$

$$= \frac{1}{2k} OPT_A(G, k) - \frac{1}{k} GR_N(G, t), \tag{13}$$

where (7) holds since the greedy strategy adds to S_t the node v maximizing $\sigma(S_t \cup \{v\})$, (8) comes from Lemma 9, (9) comes from Lemma 8, (10) comes from Lemma 10, (11) comes from Lemma 11, and (12) comes from Lemma 7. Thus, by (13) and some manipulations we get the following recursive relation: $GR_N(G, t+1) \geq \frac{1}{2k} \cdot OPT_A(G, k) + (1 - \frac{1}{k}) \cdot GR_N(G, t)$

for any $t \in [k-1]_0$. By applying iteratively the above inequality, we get $GR_N(G, k) \geq \frac{1}{2k} \cdot \sum_{t=0}^{k-1} \left(1 - \frac{1}{k}\right)^t \cdot OPT_A(G, k) = \frac{1}{2} \left(1 - \left(1 - \frac{1}{k}\right)^k\right) \cdot OPT_A(G, k)$, that leads to $\frac{GR_N(G, k)}{OPT_A(G, k)} \geq \frac{1}{2} \left(1 - \left(1 - \frac{1}{k}\right)^k\right) \geq \frac{1}{2} \left(1 - \frac{1}{e}\right)$, and this shows the claim. \blacktriangleleft

► **Remark 12.** By Theorem 5, we can easily show that, for any influence graph G with n nodes, the k -adaptivity gap of G is at most $2 \left(1 - \left(1 - \frac{1}{k}\right)^k\right)^{-1} \leq \frac{2e}{e-1} \approx 3.164$.

5 The Efficiency of the Adaptive Greedy Algorithm

We show that the adaptive version of the greedy algorithm guarantees an even better approximation ratio of $1 - \frac{1}{\sqrt{e}} \approx 0.393$ for the adaptive influence maximization problem. The *adaptive greedy algorithm* is an adaptive policy π_k^{GR} that selects k seeds in k steps, and at each step t selects the t -th seed that maximizes the expected influence spread conditioned by the observed realisation.

► **Theorem 13.** *Given an influence graph G with n nodes and $k \in [n]_2$, the adaptive greedy algorithm is a $1 - \left(1 - \frac{1}{2k}\right)^k \geq 1 - \frac{1}{\sqrt{e}}$ approximation algorithm for the adaptive influence maximization problem (under IC and myopic feedback) applied to (G, k) .*

In the proof of Theorem 13 we relate the expected influence spread coming from the adaptive greedy algorithm with that of the optimal adaptive policy, passing through a new hybrid adaptive policy. In the same spirit of Theorem 5, we also consider a new diffusion model and a new notion of adaptive submodularity (slightly different from that of Subsection 4.1). To show the main theorem (see Subsection 5.1) we need some notation and preliminary results. Let $G = (V = [n], E, (p_{uv})_{(u,v) \in E})$ be an influence graph, and let $k \in [n]_2$. Let π be an optimal adaptive policy, and let $x = (x_1, \dots, x_n)$ be the vector such that x_i is the probability that node i is selected by π .

Strong 2-level Diffusion Model. We define $L, \hat{L}, L^2(S), \Phi, \hat{\Phi}$ as in the 2-level diffusion model considered in Section 4. Given a partial realisation ψ and a set $S \subseteq V$, let $\sigma_{L, \hat{L}, \psi}^2(S) := \sigma_{L^2(S \setminus \text{dom}(\psi))}^2(S)$ denote the influence spread induced by S in live-edge graph $L^2(S \setminus \text{dom}(\psi))$, and let $\sigma_{\psi}^2(S) := \mathbb{E}_{L, \hat{L}}[\sigma_{L, \hat{L}, \psi}^2(S) | \psi \subseteq \Phi]$ denote the above influence spread in expectation, conditioned by partial realisation ψ .

Strong 2-level Hybrid Adaptive Policy. Given a partial realisation ψ , let Hyb_{ψ}^2 be a hybrid adaptive policy defined as follows: (i) Hyb_{ψ}^2 selects all the nodes in $\text{dom}(\psi)$ as seeds; (ii) then, Hyb_{ψ}^2 adds to $\text{dom}(\psi)$ all the nodes that policy π would have select when starting from the empty realisation, and observing, at each step, partial realisations coming from the live-edge graph \hat{L} only; (iii) finally, denoting with $\hat{\Psi}_{\pi}$ the random realisation returned by policy π , the expected influence spread of Hyb_{ψ}^2 is defined as $\mathbb{E}_{L, \hat{L}}[\sigma_{L, \hat{L}, \psi}^2(\text{dom}(\psi) \cup \text{dom}(\hat{\Psi}_{\pi})) | \psi \subseteq \Phi]$, i.e., it is the expected influence spread determined $\text{dom}(\psi) \cup \text{dom}(\hat{\Psi}_{\pi})$ in the strong 2-level diffusion model, conditioned by partial realisation ψ . Differently from the 2-level hybrid policy defined in Section 4, the expected influence spread of the hybrid policy defined here is conditioned by partial realisation ψ and the unique seeds that have two chances to influence their neighbors are those in $\text{dom}(\hat{\Psi}_{\pi}) \setminus \text{dom}(\psi)$.

Strong Adaptive Submodularity of the Strong 2-level Diffusion Model. Given two partial realisations $\psi, \hat{\psi}$ and $v \in V$, let

$$\begin{aligned} & \Delta_{\psi}^2(v|\hat{\psi}) \\ & := \mathbb{E}_{\mathbf{L}, \hat{\mathbf{L}}} \left[\sigma_{\mathbf{L}, \hat{\mathbf{L}}, \psi}^2(\{v\} \cup \text{dom}(\psi) \cup \text{dom}(\hat{\psi})) - \sigma_{\mathbf{L}, \hat{\mathbf{L}}, \psi}^2(\text{dom}(\psi) \cup \text{dom}(\hat{\psi})) \mid \psi \subseteq \Phi, \hat{\psi} \subseteq \hat{\Phi} \right], \end{aligned}$$

i.e., $\Delta_{\psi}^2(v|\hat{\psi})$ is the expected increment of the influence spread in the strong 2-level diffusion model when adding seed v to the set of nodes $\text{dom}(\hat{\psi}) \cup \text{dom}(\psi)$, conditioned by the observation of partial realisations $\hat{\psi}$ and ψ . We say that the strong 2-level diffusion model is *strongly adaptive submodular* if, for any partial realisations $\hat{\psi}, \hat{\psi}'$ with $\hat{\psi} \subseteq \hat{\psi}'$, any partial realisation ψ , and any $v \in V$, we have that $\Delta_{\psi}^2(v|\hat{\psi}) \geq \Delta_{\psi}^2(v|\hat{\psi}')$.

► **Lemma 14.** *The strong 2-level diffusion model is strongly adaptive submodular.*

From the Ordinary to the Strong 2-level Diffusion Model. Given a partial realisation ψ , and $v \in V$, let $\Delta_{\psi}(v) := \mathbb{E}_{\mathbf{L}}[\sigma_{\mathbf{L}}(\{v\} \cup \text{dom}(\psi)) - \sigma_{\mathbf{L}}(\text{dom}(\psi)) \mid \psi \subseteq \Phi]$, that is the expected increment under the ordinary diffusion model when adding a new node to the nodes in $\text{dom}(\psi)$ conditioned by partial realisation ψ , and let $\Delta_{\psi}^2(v) := \Delta_{\psi}^2(v|\emptyset) = \mathbb{E}_{\mathbf{L}, \hat{\mathbf{L}}}[\sigma_{\mathbf{L}, \hat{\mathbf{L}}, \psi}(\{v\} \cup \text{dom}(\psi)) - \sigma_{\mathbf{L}, \hat{\mathbf{L}}, \psi}(\text{dom}(\psi)) \mid \psi \subseteq \Phi]$, that is the above conditional expectation, but w.r.t. the strong 2-level diffusion model. The following lemma can be shown analogously to Lemma 8.

► **Lemma 15.** *We have that $\Delta_{\psi}^2(v) \leq 2 \cdot \Delta_{\psi}(v)$ for any partial realisation ψ and $v \in V$.*

From the Optimal to the Greedy Adaptive Policy. The following lemma will be used to relate the strong 2-level hybrid adaptive policy with the adaptive greedy policy; its proof is similar to that of Lemma 10, but uses the concept of strong adaptive submodularity.

► **Lemma 16.** *For any partial realisation ψ , we have*

$$\mathbb{E}_{\mathbf{L}, \hat{\mathbf{L}}}[\sigma_{\mathbf{L}, \hat{\mathbf{L}}, \psi}^2(\text{dom}(\psi) \cup \text{dom}(\hat{\Psi}_{\pi})) \mid \psi \subseteq \Phi] \leq \mathbb{E}_{\mathbf{L}}[\sigma_{\mathbf{L}}(\text{dom}(\psi)) \mid \psi \subseteq \Phi] + \sum_{v \in V \setminus \text{dom}(\psi)} x_v \cdot \Delta_{\psi}^2(v).$$

The following lemma shows that the optimal adaptive influence spread is upper bounded by that expected influence spread of the strong 2-level hybrid adaptive policy, and its proof is completely analogue to that of Lemma 11.

► **Lemma 17.** *We have $\text{OPT}_A(G, k) \leq \mathbb{E}_{\mathbf{L}, \hat{\mathbf{L}}}[\sigma_{\mathbf{L}, \hat{\mathbf{L}}, \psi}^2(\text{dom}(\psi) \cup \text{dom}(\hat{\Psi}_{\pi})) \mid \psi \subseteq \Phi]$ for any partial realisation ψ .*

5.1 Proof of Theorem 13

Armed with the above lemmas, and by using a similar approach as in the proof of Theorem 5, we can prove Theorem 13. Given $t \in [k]_0$, let \mathbf{S}_t denote the (random) set of the first t seeds selected by the adaptive greedy policy π_k^{GR} , and let Ψ_t be the (random) partial realisation such that $\text{dom}(\Psi_t) = \mathbf{S}_t$ (i.e., the partial realisation observed by policy π_k^{GR} at the end of step t); let $\text{GR}_A(G, t) := \mathbb{E}_{\mathbf{L}}[\sigma_{\mathbf{L}}(\mathbf{S}_t)]$ denote the expected influence spread of the adaptive greedy policy after selecting the first t seeds. For any $t \in [k-1]_0$, we have that

$$\begin{aligned}
& GR_A(G, t+1) - GR_A(G, t) \\
&= \mathbb{E}_{\mathbf{L}}[\sigma(\mathbf{S}_{t+1}) - \sigma(\mathbf{S}_t)] \\
&= \mathbb{E}_{\Psi_t} \left[\max_{v \in V \setminus \text{dom}(\Psi_t)} \mathbb{E}_{\mathbf{L}}[\sigma_{\mathbf{L}}(\{v\} \cup \text{dom}(\Psi_t)) - \sigma_{\mathbf{L}}(\text{dom}(\Psi_t)) | \Psi_t] \right] \tag{14}
\end{aligned}$$

$$\begin{aligned}
&\geq \mathbb{E}_{\Psi_t} \left[\sum_{v \in V \setminus \text{dom}(\Psi_t)} \frac{x_v}{k} \cdot \mathbb{E}_{\mathbf{L}}[\sigma_{\mathbf{L}}(\{v\} \cup \text{dom}(\Psi_t)) - \sigma_{\mathbf{L}}(\text{dom}(\Psi_t)) | \Psi_t] \right] \\
&= \frac{1}{k} \cdot \mathbb{E}_{\Psi_t} \left[\sum_{v \in V \setminus \text{dom}(\Psi_t)} x_v \cdot \mathbb{E}_{\mathbf{L}}[\sigma_{\mathbf{L}}(\{v\} \cup \text{dom}(\Psi_t)) - \sigma_{\mathbf{L}}(\text{dom}(\Psi_t)) | \Psi_t] \right] \\
&= \frac{1}{k} \cdot \mathbb{E}_{\Psi_t} \left[\sum_{v \in V \setminus \text{dom}(\Psi_t)} x_v \cdot \Delta_{\Psi_t}(v) \right] \\
&\geq \frac{1}{2k} \cdot \mathbb{E}_{\Psi_t} \left[\sum_{v \in V \setminus \text{dom}(\Psi_t)} x_v \cdot \Delta_{\Psi_t}^2(v) \right] \tag{15}
\end{aligned}$$

$$\geq \frac{1}{2k} \cdot \mathbb{E}_{\Psi_t} \left[\mathbb{E}_{\mathbf{L}, \hat{\mathbf{L}}}[\sigma_{\mathbf{L}, \hat{\mathbf{L}}}^2(\text{dom}(\Psi_t) \cup \text{dom}(\hat{\Psi}_\pi)) | \Psi_t] - \mathbb{E}_{\mathbf{L}}[\sigma_{\mathbf{L}}(\text{dom}(\Psi_t)) | \Psi_t] \right] \tag{16}$$

$$\geq \frac{1}{2k} \cdot \mathbb{E}_{\Psi_t} [OPT_A(G, k) - \mathbb{E}_{\mathbf{L}}[\sigma_{\mathbf{L}}(\text{dom}(\Psi_t)) | \Psi_t]] \tag{17}$$

$$\begin{aligned}
&\geq \frac{1}{2k} \cdot OPT_A(G, k) - \frac{1}{2k} \cdot \mathbb{E}_{\Psi_t} [\mathbb{E}_{\mathbf{L}}[\sigma_{\mathbf{L}}(\text{dom}(\Psi_t)) | \Psi_t]] \\
&= \frac{1}{2k} \cdot OPT_A(G, k) - \frac{1}{2k} \cdot \mathbb{E}_{\mathbf{L}}(\sigma_{\mathbf{L}}(\mathbf{S}_t)) \\
&= \frac{1}{2k} \cdot OPT_A(G, k) - \frac{1}{2k} \cdot GR_A(G, t), \tag{18}
\end{aligned}$$

where (14) holds since the adaptive greedy strategy at each step adds the node maximizing the expected influence spread (conditioned by the partial realisation coming from the previous selected seeds), (15) comes from Lemma 15, (16) comes from Lemma 16, (17) comes from Lemma 17. Thus, by (18) and some manipulations we get the following recursive relation:

$$GR_A(G, t+1) \geq \frac{1}{2k} \cdot OPT_A(G, k) + \left(1 - \frac{1}{2k}\right) \cdot GR_A(G, t), \quad \forall t \in [k-1]_0. \tag{19}$$

By applying iteratively (19), we get

$$GR_A(G, k) \geq \frac{1}{2k} \cdot \sum_{t=0}^{k-1} \left(1 - \frac{1}{2k}\right)^t \cdot OPT_A(G, k) = 1 - \left(1 - \frac{1}{2k}\right)^k \cdot OPT_A(G, k),$$

that leads to

$$\frac{GR_A(G, k)}{OPT_A(G, k)} \geq 1 - \left(1 - \frac{1}{2k}\right)^k \geq 1 - \frac{1}{\sqrt{e}},$$

and this shows the claim.

6 Conclusions and Future Work

In the context of adaptive optimization, we have introduced a new approach to relate the solution provided by a simple non-adaptive greedy policy with the adaptive optimum. The new approach allowed us to establish better bounds for the adaptive influence maximization problem under myopic feedback, specifically we improve both the approximation ratio of the non-adaptive greedy policy and the adaptivity gap.

Our results open several research directions in the context of influence maximization and in more general adaptive optimization settings. The approximation factor of the non-adaptive (resp. adaptive) greedy algorithm is between our lower bound of $\frac{1}{2} \left(1 - \frac{1}{e}\right) \approx 0.316$ (resp. $1 - \frac{1}{\sqrt{e}} \approx 0.393$) and the upper bound of $\frac{e^2+1}{(e+1)^2} \approx 0.606$ [39], and the adaptivity gap is between the lower bound of $\frac{e}{e-1} \approx 1.582$ [39] and our upper bound of $\frac{2e}{e-1} \approx 3.164$. The first problem left open by our result is to close these gaps. Furthermore, the techniques introduced in this paper to relate non-adaptive policies with adaptive ones might be useful to find better bounds in several variants of the adaptive influence maximization problem, like a combination of the following settings: different feedback models (e.g., the full-adoption feedback), different diffusion models (e.g., the general triggering model [33]), and different graph classes. Finally, as shown in Section 2, we believe that our new approach could be efficiently used to analyze non-adaptive greedy algorithms in other adaptive optimization problems, like e.g. the stochastic probing problem [6, 29, 30].

References

- 1 Marek Adamczyk, Maxim Sviridenko, and Justin Ward. Submodular stochastic probing on matroids. In *31st International Symposium on Theoretical Aspects of Computer Science, (STACS)*, pages 29–40, 2014. doi:10.4230/LIPIcs.STACS.2014.29.
- 2 Georgios Amanatidis, Federico Fusco, Philip Lazos, Stefano Leonardi, and Rebecca Reifenhäuser. Fast adaptive non-monotone submodular maximization subject to a knapsack constraint. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems, (NeurIPS)*, 2020.
- 3 Arash Asadpour and Hamid Nazerzadeh. Maximizing stochastic monotone submodular functions. *Management Science*, 62(8):2374–2391, 2016.
- 4 Arash Asadpour, Hamid Nazerzadeh, and Amin Saberi. Stochastic submodular maximization. In *Internet and Network Economics, 4th International Workshop, (WINE)*, pages 477–489, 2008.
- 5 Ashwinkumar Badanidiyuru, Christos Papadimitriou, Aviad Rubinfeld, Lior Seeman, and Yaron Singer. Locally adaptive optimization: Adaptive seeding for monotone submodular functions. *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA)*, 1:414–429, 2016.
- 6 Domagoj Bradac, Sahil Singla, and Goran Zuzic. (near) optimal adaptivity gaps for stochastic multi-value probing. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*, pages 49:1–49:21, 2019.
- 7 Gruia Călinescu, Chandra Chekuri, Martin Pál, and Jan Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM J. Comput.*, 40(6):1740–1766, 2011.
- 8 Carri W. Chan and Vivek F. Farias. Stochastic depletion problems: Effective myopic policies for a class of dynamic optimization problems. *Math. Oper. Res.*, 34(2):333–350, 2009. doi:10.1287/moor.1080.0364.
- 9 Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. *SIAM J. Comput.*, 43(6):1831–1879, 2014.

- 10 Wei Chen, Laks V. S. Lakshmanan, and Carlos Castillo. *Information and Influence Propagation in Social Networks*. Synthesis Lectures on Data Management. Morgan & Claypool, 2013.
- 11 Wei Chen and Binghui Peng. On adaptivity gaps of influence maximization under the independent cascade model with full-adoption feedback. *30th International Symposium on Algorithms and Computation, (ISAAC)*, pages 24:1–24:19, 2019.
- 12 Wei Chen, Binghui Peng, Grant Schoenebeck, and Biaoshuai Tao. Adaptive Greedy versus Non-adaptive Greedy for Influence Maximization. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, (AAAI)*, 2020.
- 13 Wei Chen, Chi Wang, and Yajun Wang. 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*, pages 1029–1038, 2010.
- 14 Edith Cohen, Daniel Delling, Thomas Pajor, and Renato F. Werneck. Sketch-based influence maximization and computation: Scaling up with guarantees. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management, (CIKM)*, pages 629–638, 2014.
- 15 Gianlorenzo D'Angelo, Debashmita Poddar, and Cosimo Vinci. Improved approximation factor for adaptive influence maximization via simple greedy strategies. *CoRR*, abs/2007.09065, 2020. [arXiv:2007.09065](https://arxiv.org/abs/2007.09065).
- 16 Gianlorenzo D'Angelo, Debashmita Poddar, and Cosimo Vinci. Better bounds on the adaptivity gap of influence maximization under full-adoption feedback. In *AAAI Conference on Artificial Intelligence, (AAAI)*, 2021. To appear.
- 17 Brian C. Dean, Michel X. Goemans, and Jan Vondrák. Adaptivity and approximation for stochastic packing problems. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA)*, pages 395–404, 2005. URL: <http://dl.acm.org/citation.cfm?id=1070432.1070487>.
- 18 Brian C. Dean, Michel X. Goemans, and Jan Vondrák. Approximating the stochastic knapsack problem: The benefit of adaptivity. *Math. Oper. Res.*, 33(4):945–964, 2008. doi:10.1287/moor.1080.0330.
- 19 Pedro Domingos and Matt Richardson. Mining the network value of customers. *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 57–66, 2001.
- 20 Alina Ene and Huy L. Nguyen. Submodular maximization with nearly-optimal approximation and adaptivity in nearly-linear time. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA)*, pages 274–282, 2019.
- 21 Matthew Fahrbach, Vahab S. Mirrokni, and Morteza Zadimoghaddam. Non-monotone submodular maximization with nearly optimal adaptivity and query complexity. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning, (ICML)*, pages 1833–1842, 2019.
- 22 Matthew Fahrbach, Vahab S. Mirrokni, and Morteza Zadimoghaddam. Submodular maximization with nearly optimal approximation, adaptivity and query complexity. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA)*, pages 255–273, 2019.
- 23 Michel X. Goemans and Jan Vondrák. Stochastic covering and adaptivity. In *Theoretical Informatics, 7th Latin American Symposium, (LATIN)*, pages 532–543, 2006. doi:10.1007/11682462_50.
- 24 Sharon Goldberg and Zhenming Liu. The diffusion of networking technologies. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA)*, pages 1577–1594, 2013.
- 25 Daniel Golovin and Andreas Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *Journal of Artificial Intelligence Research*, 42:427–486, 2011.

- 26 Alkis Gotovos, Amin Karbasi, and Andreas Krause. Non-monotone adaptive submodular maximization. In *Proceedings of the 24th International Conference on Artificial Intelligence (IJCAI)*, page 1996–2003. AAAI Press, 2015.
- 27 Amit Goyal, Wei Lu, and Laks V. S. Lakshmanan. CELF++: optimizing the greedy algorithm for influence maximization in social networks. In *Proceedings of the 20th International Conference on World Wide Web, (WWW)*, pages 47–48, 2011.
- 28 Andrew Guillory and Jeff A. Bilmes. Interactive submodular set cover. In *Proceedings of the 27th International Conference on Machine Learning, (ICML)*, pages 415–422, 2010. URL: <https://icml.cc/Conferences/2010/papers/436.pdf>.
- 29 Anupam Gupta, Viswanath Nagarajan, and Sahil Singla. Algorithms and Adaptivity Gaps for Stochastic Probing. *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms, (SODA)*, 151:1731–1747, 2016.
- 30 Anupam Gupta, Viswanath Nagarajan, and Sahil Singla. Adaptivity gaps for stochastic probing: Submodular and XOS functions. *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA)*, pages 1688–1702, 2017.
- 31 Kai Han, Keke Huang, Xiaokui Xiao, Jing Tang, Aixin Sun, and Xueyan Tang. Efficient algorithms for adaptive influence maximization. *Proceedings of the VLDB Endowment*, 11(9):1029–1040, 2018.
- 32 Lisa Hellerstein, Devorah Kletenik, and Patrick Lin. Discrete stochastic submodular maximization: Adaptive vs. non-adaptive vs. offline. In *Algorithms and Complexity - 9th International Conference, (CIAC)*, pages 235–248, 2015. doi:10.1007/978-3-319-18173-8_17.
- 33 David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 137–146. ACM, 2003.
- 34 David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. *Theory of Computing*, 11:105–147, 2015.
- 35 Jure Leskovec, Andreas Krause, Carlos Guestrin, Christos Faloutsos, Jeanne M. VanBriesen, and Natalie S. Glance. Cost-effective outbreak detection in networks. In *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 420–429, 2007.
- 36 Y. Li, J. Fan, Y. Wang, and K. Tan. Influence maximization on social graphs: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 30(10):1852–1872, 2018.
- 37 Hung T. Nguyen, My T. Thai, and Thang N. Dinh. Stop-and-stare: Optimal sampling algorithms for viral marketing in billion-scale networks. In *Proceedings of the 2016 International Conference on Management of Data, (SIGMOD)*, pages 695–710, 2016.
- 38 Srinivasan Parthasarathy. An analysis of the greedy algorithm for stochastic set cover. *CoRR*, abs/1803.07639, 2018. arXiv:1803.07639.
- 39 Binghui Peng and Wei Chen. Adaptive influence maximization with myopic feedback. *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing System, (NeurIPS)*, pages 5575–5584, 2019.
- 40 Matthew Richardson and Pedro M. Domingos. Mining knowledge-sharing sites for viral marketing. In *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 61–70, 2002.
- 41 Aviad Rubinfeld, Lior Seeman, and Yaron Singer. Approximability of adaptive seeding under knapsack constraints. *Proceedings of the 2015 ACM Conference on Economics and Computation, (EC)*, pages 797–814, 2015.
- 42 Guillaume Salha, Nikolaos Tziortziotis, and Michalis Vazirgiannis. Adaptive submodular influence maximization with myopic feedback. *Proceedings of the 2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 455–462, 2018.
- 43 Lior Seeman and Yaron Singer. Adaptive seeding in social networks. *Proceedings - Annual IEEE Symposium on Foundations of Computer Science, (FOCS)*, pages 459–468, 2013.

- 44 Yaron Singer. Influence maximization through adaptive seeding. *ACM SIGecom Exchanges*, 15(1):32–59, 2016.
- 45 Lichao Sun, Weiran Huang, Philip S. Yu, and Wei Chen. Multi-round influence maximization. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 2249–2258, 2018.
- 46 Jing Tang, Laks V.S. Lakshmanan, Keke Huang, Xueyan Tang, Aixin Sun, Xiaokui Xiao, and Andrew Lim. Efficient approximation algorithms for adaptive seed minimization. *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pages 1096–1113, 2019.
- 47 Youze Tang, Yanchen Shi, and Xiaokui Xiao. Influence maximization in near-linear time: A martingale approach. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*, pages 1539–1554, 2015.
- 48 Youze Tang, Xiaokui Xiao, and Yanchen Shi. Influence Maximization : Near-Optimal Time Complexity Meets Practical Efficiency. *SIGMOD '14: Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*, pages 75–86, 2014.
- 49 Guangmo Tong and Ruiqi Wang. Adaptive Influence Maximization under General Feedback Models, 2019. [arXiv:1902.00192v3](https://arxiv.org/abs/1902.00192v3).
- 50 Guangmo Tong, Ruiqi Wang, Zheng Dong, and Xiang Li. Time-constrained Adaptive Influence Maximization, 2020. [arXiv:2001.01742v2](https://arxiv.org/abs/2001.01742v2).
- 51 Guangmo Tong, Weili Wu, Shaojie Tang, and Ding Zhu Du. Adaptive Influence Maximization in Dynamic Social Networks. *IEEE/ACM Transactions on Networking*, 25(1):112–125, 2017.
- 52 Sharan Vaswani and Laks V. S. Lakshmanan. Adaptive influence maximization in social networks: Why commit when you can adapt? *CoRR*, abs/1604.08171, 2016. [arXiv:1604.08171](https://arxiv.org/abs/1604.08171).
- 53 Jing Yuan and Shaojie Tang. No time to observe: Adaptive influence maximization with partial feedback. *IJCAI International Joint Conference on Artificial Intelligence*, pages 3908–3914, 2017.