

Adaptive Greedy versus Non-adaptive Greedy for Influence Maximization

Wei Chen
Microsoft Research, China
weic@microsoft.com

Binghui Peng
Columbia University
bp2601@columbia.edu

Grant Schoenebeck
University of Michigan
schoeneb@umich.edu

Biaoshuai Tao
University of Michigan
bstao@umich.edu

Abstract

We consider the *adaptive influence maximization problem*: given a network and a budget k , iteratively select k seeds in the network to maximize the expected number of adopters. In the *full-adoption feedback model*, after selecting each seed, the seed-picker observes all the resulting adoptions. In the *myopic feedback model*, the seed-picker only observes whether each neighbor of the chosen seed adopts. Motivated by the extreme success of greedy-based algorithms/heuristics for influence maximization, we propose the concept of *greedy adaptivity gap*, which compares the performance of the adaptive greedy algorithm to its non-adaptive counterpart. Our first result shows that, for submodular influence maximization, the adaptive greedy algorithm can perform up to a $(1 - 1/e)$ -fraction worse than the non-adaptive greedy algorithm, and that this ratio is tight. More specifically, on one side we provide examples where the performance of the adaptive greedy algorithm is only a $(1 - 1/e)$ fraction of the performance of the non-adaptive greedy algorithm in four settings: for both feedback models and both the *independent cascade model* and the *linear threshold model*. On the other side, we prove that in any submodular cascade, the adaptive greedy algorithm always outputs a $(1 - 1/e)$ -approximation to the expected number of adoptions in the optimal non-adaptive seed choice. Our second result shows that, for the general submodular cascade model with full-adoption feedback, the adaptive greedy algorithm can outperform the non-adaptive greedy algorithm by an unbounded factor. Finally, we propose a risk-free variant of the adaptive greedy algorithm that always performs no worse than the non-adaptive greedy algorithm.

1 Introduction

The *influence maximization problem* (INFMAX) is an optimization problem that asks which seeds a viral marketing campaign should target (e.g. by giving free products) so that propagation from these seeds influences the most people in a social network. That is, given a graph, a *stochastic diffusion model* defining how each node is infected by its neighbors, and a limited budget k , how to pick k seeds such that the expected number of total infected nodes in this graph at the end of the diffusion is maximized. This problem has significant applications in viral marketing, outbreak detection, rumor controls, etc, and has been extensively studied (cf. Chen et al. [9], Li et al. [22]).

For INFMAX, most of the existing work has considered *submodular* diffusion models, especially the *independent cascade model* and the *linear threshold model* [19]. Likewise, we also focus on submodular diffusion models. In submodular diffusion models, a vertex v 's marginal probability of becoming infected after a new neighbor t is infected given S as the set of v 's already infected

neighbors is at least the marginal probability that v is infected after t is newly infected given $T \supseteq S$ as the set of v 's already infected neighbors (see the paragraph before Theorem 2.4 for more details). Intuitively, this means that the influence of infected nodes are substitutes and never have synergy.

When submodular INFMAX is considered, nearly all the known algorithms/heuristics are based on a greedy algorithm that iteratively picks the seed that has the largest marginal influence. Some of them improve the running time of the original greedy algorithm by skipping vertices that are known to be suboptimal [21, 15], while the others improve the scalability of the greedy algorithm by using more scalable algorithms to approximate the expected total influence [4, 33, 34, 10, 25] or computing a score of the seeds that is closely related to the expected total influence [6, 8, 7, 16, 18, 12, 32, 29]. Arora et al. [2] benchmark most of the aforementioned variants of the greedy algorithms.

In this paper, we study the *adaptive influence maximization problem*, where seeds are selected iteratively and feedback is given to the seed-picker after selecting each seed. Two different feedback models have been studied in the past: the *full-adoption feedback model* and the *myopic feedback model* [14]. In the full-adoption feedback model, the seed-picker sees the entire diffusion process of each selected seed, and in the myopic feedback model the seed-picker only sees whether each neighbor of the chosen seed is infected.

Past literature focused on the *adaptivity gap*—the ratio between the performance of the *optimal* adaptive algorithm and the performance of the *optimal* non-adaptive algorithm [14, 26, 5]. However, even in the non-adaptive setting, INFMAX is known to be APX-hard [19, 29]. As a result, in practice, it is not clear whether the adaptivity gap can measure how much better an adaptive algorithm can do.

In this paper, we define and consider the *greedy adaptivity gap*, which is the ratio between the performance of the adaptive greedy algorithm and the non-adaptive greedy algorithm. We focus on the gap between the greedy algorithms for three reasons. First, as we mentioned, the APX-hardness of INFMAX renders the practical implications of the adaptivity gap unclear. Second, as we remarked at the beginning, the greedy algorithm is used almost exclusively in the context of influence maximization. Third, the iterative nature of the original greedy algorithm naturally extends to the adaptive setting.

1.1 Our Results

We show that, for the general submodular diffusion models, with both the full-adoption feedback model and the myopic feedback model, the infimum of the greedy adaptivity gap is exactly $(1 - 1/e)$. In addition, this result can be extended to the two well-studied submodular diffusion models: the independent cascade model and the linear threshold model. This is proved in two steps.

In the first step, we show that there are INFMAX instances where the adaptive greedy algorithm can only produce $(1 - 1/e)$ fraction of the influence of the solution output by the non-adaptive greedy algorithm. This result is surprising: one would expect that the adaptivity is always helpful, as the feedback provides more information to the seed-picker, which makes the seed-picker refine the seed choices in future iterations. Our result shows that this is not the case, and the feedback, if overly used, can make the seed-picker act in a more myopic way, which is potentially harmful.

In the second step, we show that the adaptive greedy algorithm always achieves a $(1 - 1/e)$ -approximation of the non-adaptive optimal solution, so its performance is always at least a $(1 - 1/e)$ fraction of the performance of the non-adaptive greedy algorithm. In particular, combining the two steps, we see that when the adaptive greedy algorithm output only obtains a (nearly) $(1 - 1/e)$ -fraction of the performance of the non-adaptive greedy algorithm, the non-adaptive greedy algorithm is (almost) optimal. This worst-case guarantee indicates that the adaptive greedy algorithm will never be too bad.

model	AG	GAG inf	GAG sup
ICM, full-adoption	at least $e/(e-1)$ [5]	$1 - 1/e$ (Thm 3.1)	unknown
ICM, myopic	at least $e/(e-1)$, at most 4 [26]	$1 - 1/e$ (Thm 3.1)	unknown
LTM, full-adoption	unknown	$1 - 1/e$ (Thm 3.1)	unknown
LTM, myopic	unknown	$1 - 1/e$ (Thm 3.1)	unknown
GSDM, full-adoption	∞ (Thm 4.2)	$1 - 1/e$ (Thm 3.1)	∞ (Thm 4.1)
GSDM, myopic	at least $e/(e-1)$ (implied by [26])	$1 - 1/e$ (Thm 3.1)	unknown

Table 1: Results for the adaptivity gap (AG), the infimum of the greedy adaptivity gap (GAG inf) and the supremum of the greedy adaptivity gap (GAG sup), where GSDM stands for general submodular diffusion model.

As the second result, we show that the supremum of the greedy adaptivity gap is infinity, for the general submodular diffusion model with full-adoption feedback. This indicates that the adaptive greedy algorithm can perform significantly better than its non-adaptive counterpart. We also show, with almost the same proof, that the adaptivity gap in this setting (general submodular model with full-adoption feedback) is also unbounded.

Finally, we propose a risk-free but more conservative variant of the adaptive greedy algorithm, which always performs at least as well as the non-adaptive greedy algorithm. We recommend both the adaptive greedy algorithm and this variant.

1.2 Related Work

The influence maximization problem was initially posed by Domingos and Richardson [11], Richardson and Domingos [27]. Kempe et al. [19] proposed the linear threshold model and the independent cascade model, and show that they are submodular. Whenever a diffusion model is submodular, the greedy algorithm was shown to obtain a $(1 - 1/e)$ -approximation to the optimal number of infections [24, 19, 20, 23].

For adaptive INFMAX, Golovin and Krause [14] showed that INFMAX with the independent cascade model and full-adoption feedback is *adaptive submodular*, which implies that the adaptive greedy algorithm obtains a $(1 - 1/e)$ -approximation to the adaptive optimal solution. On the other hand, INFMAX for the independent cascade model with myopic feedback, as well as INFMAX for the linear threshold model with both feedback models, are not adaptive submodular. In particular, the adaptive greedy algorithm fails to obtain a $(1 - 1/e)$ -approximation for the independent cascade model with myopic feedback [26]. Peng and Chen [26] showed that the adaptivity gap for the independent cascade model with myopic feedback is at most 4 and at least $e/(e-1)$, and they also showed that both the adaptive and non-adaptive greedy algorithms perform a $0.25(1 - 1/e)$ -approximation to the adaptive optimal solution. The adaptivity gap for the independent cascade model with full-adoption feedback, as well as the adaptivity gap for the linear threshold model with both feedback models, are still open problems, although there is some partial progress [5].

Our paper is not the first work studying the adaptive greedy algorithm. Previous work focused on improving the running time of the adaptive greedy algorithm [17]. However, to the best of our knowledge, our work is the first one that compares the adaptive greedy algorithm to its non-adaptive counterpart.

Finally, we remark that there do exist INFMAX algorithms that are not based on greedy [3, 13, 1, 28, 30, 31], but they are typically for non-submodular diffusion models.

We summarize the existing results about the adaptivity gap and our new results about the greedy adaptivity gap in Table 1.

2 Preliminary

All graphs in this paper are simple and directed. Given a graph $G = (V, E)$ and a vertex $v \in V$, let $\Gamma(v)$ and $\deg(v)$ be the set of in-neighbors and the in-degree of v respectively.

2.1 Triggering Model

We consider the well-studied *triggering model* [19], which is commonly used to capture “general” submodular diffusion models. A more general way to capture submodular diffusion models is the *general threshold model* [19] with *submodular local influence functions*. All our results hold under this setting as well. We will discuss this in Appendix B.

Definition 2.1 (Kempe et al. [19]). The *triggering model*, $I_{G,F}$, is defined by a graph $G = (V, E)$ and for each vertex v a distribution \mathcal{F}_v over the subset of its in-neighbors $\{0, 1\}^{|\Gamma(v)|}$. Let $F = \{\mathcal{F}_v \mid v \in V\}$.

On an input seed set $S \subseteq V$, $I_{G,F}(S)$ outputs a set of infected vertices as follows:

1. Initially, only vertices in S are infected. Each vertex v samples a subset of its in-neighbors $T_v \subseteq \Gamma(v)$ from \mathcal{F}_v independently. We call T_v the *triggering set* of v .
2. In each subsequent round, a vertex v becomes infected if a vertex in T_v is infected in the previous round.
3. After a round where no additional vertices are infected, the set of infected vertices is the output.

$I_{G,F}$ in Definition 2.1 can be viewed as a random function $I_{G,F} : \{0, 1\}^{|V|} \rightarrow \{0, 1\}^{|V|}$. In addition, if the triggering set T_v is fixed for each vertex v , then $I_{G,F}$ is deterministic. Given v , its triggering set T_v , and an in-neighbor $u \in \Gamma(v)$, we say that the edge (u, v) is *live* if $u \in T_v$, and we say that (u, v) is *blocked* if $u \notin T_v$. It is easy to see that, when the triggering sets for all vertices are sampled, $I_{G,F}(S)$ is the set of all vertices that are reachable from S when removing all blocked edges from the graph.

We define a *realization* of a graph $G = (V, E)$ as a function $\phi : E \rightarrow \{\mathbf{L}, \mathbf{B}\}$ such that $\phi(e) = \mathbf{L}$ if $e \in E$ is live and $\phi(e) = \mathbf{B}$ if $e \in E$ is blocked. Let $I_{G,F}^\phi : \{0, 1\}^{|V|} \rightarrow \{0, 1\}^{|V|}$ be the deterministic function corresponding to the triggering model $I_{G,F}$ with vertices’ triggering sets following realization ϕ . We write $\phi \sim F$ to indicate that a realization ϕ is sampled according to $F = \{\mathcal{F}_v\}$.

The triggering model captures the well-known independent cascade and linear threshold models. In the two definitions below, we define the two models in terms of the triggering model, which is sufficient for this paper. In Appendix A, we present the original definitions and give some intuitions for the two models for those readers who are not familiar with them.

Definition 2.2. The *independent cascade model* ICM is a special case of the triggering model $I_{G,F}$ where $G = (V, E, w)$ is an edge-weighted graph with $w(u, v) \in (0, 1]$ for each $(u, v) \in E$ and \mathcal{F}_v is the distribution such that each $u \in \Gamma(v)$ is included in T_v with probability $w(u, v)$ independently.

Definition 2.3. The *linear threshold model* LTM is a special case of the triggering model $I_{G,F}$ where $G = (V, E, w)$ is an edge-weighted graph with $w(u, v) > 0$ for each $(u, v) \in E$ and $\sum_{u \in \Gamma(v)} w(u, v) \leq 1$ for each $v \in V$, and \mathcal{F}_v is the distribution defined as follows: order v ’s in-neighbors u_1, \dots, u_T arbitrarily, sample a real number r in $[0, 1]$ uniformly, and

$$T_v = \begin{cases} \{u_t\} & \text{if } r \in \left[\sum_{i=1}^{t-1} w(u_i, v), \sum_{i=1}^t w(u_i, v) \right) \\ \emptyset & \text{if } r \geq \sum_{i=1}^T w(u_i, v) \end{cases}.$$

Intuitively, T_v includes at most one of v 's in-neighbors such that each u_t is included with probability $w(u_t, v)$.

Given a triggering model $I_{G,F}$, let $\sigma_{G,F} : \{0,1\}^{|V|} \rightarrow \mathbb{R}_{\geq 0}$ be the *global influence function* defined as $\sigma_{G,F}(S) = \mathbb{E}_{\phi \sim F}[|I_{G,F}^\phi(S)|]$. We drop the subscripts G, F and write the global influence function as $\sigma(\cdot)$ when there is no ambiguity.

A function f mapping from a set of elements to a non-negative value is *submodular* if $f(A \cup \{v\}) - f(A) \geq f(B \cup \{v\}) - f(B)$ for any two sets A, B with $A \subsetneq B$ and any element $v \notin B$.

Theorem 2.4 (Kempe et al. [19]). *For any triggering model $I_{G,F}$, $\sigma_{G,F}(\cdot)$ is submodular. In particular, $\sigma_{G,F}(\cdot)$ is submodular for both ICM and LTM.*

2.2 INFMAX and Adaptive INFMAX

Definition 2.5. The *influence maximization problem* (INFMAX) is an optimization problem which takes inputs $G = (V, E)$, F , and $k \in \mathbb{Z}^+$, and outputs a seed set S that maximizes the expected total number of infections: $S \in \operatorname{argmax}_{S \subseteq V: |S| \leq k} \sigma(S)$.

In the remaining part of this subsection, we define the adaptive version of the influence maximization problem. We will define two different models: the *full-adoption feedback model* and the *myopic feedback model*. Suppose a seed set $S \subseteq V$ is chosen by the seed-picker, and an underlying realization ϕ is given but not known by the seed-picker. Informally, in the full-adoption feedback model, the seed-picker sees all the vertices that are infected by S in all future iterations, i.e., the seed-picker sees $I_{G,F}^\phi(S)$. In the myopic feedback model, the seed-picker only sees the states of S 's neighbors, i.e., whether each vertex in $\{v \mid \exists s \in S : s \in \Gamma(v)\}$ is infected.

Define a *partial realization* as a function $\varphi : E \rightarrow \{\text{L}, \text{B}, \text{U}\}$ such that $\phi(e) = \text{L}$ if e is known to be live, $\phi(e) = \text{B}$ if e is known to be blocked, and $\phi(e) = \text{U}$ if the status of e is not yet known. We say that a partial realization φ is *consistent with* the full realization ϕ , denoted by $\phi \simeq \varphi$, if $\phi(v) = \varphi(v)$ whenever $\varphi(v) \neq \text{U}$. For the ease of notation, for an edge $(u, v) \in E$, we will write $\phi(u, v), \varphi(u, v)$ instead of $\phi((u, v)), \varphi((u, v))$.

Definition 2.6. Given a triggering model $I_{G=(V,E),F}$ with a realization ϕ , the *full-adoption feedback* is a function $\Phi_{G,F,\phi}^f$ mapping a seed set $S \subseteq V$ to a partial realization φ such that

- $\varphi(u, v) = \phi(u, v)$ for each $u \in I_{G,F}^\phi(S)$, and
- $\varphi(u, v) = \text{U}$ for each $u \notin I_{G,F}^\phi(S)$.

Definition 2.7. Given a triggering model $I_{G=(V,E),F}$ with a realization ϕ , the *myopic feedback* is a function $\Phi_{G,F,\phi}^m$ mapping a seed set $S \subseteq V$ to a partial realization φ such that

- $\varphi(u, v) = \phi(u, v)$ for each $u \in S$, and
- $\varphi(u, v) = \text{U}$ for each $u \notin S$.

An *adaptive policy* π is a function that maps a seed set S and a partial realization φ to a vertex $v = \pi(S, \varphi)$, which corresponds to the next seed the policy π would choose given φ and S being the set of seeds that has already been chosen. Naturally, we only care about $\pi(S, \varphi)$ when $\varphi = \Phi_{G,F,\phi}^f(S)$ or $\varphi = \Phi_{G,F,\phi}^m(S)$, although we define π that specifies an output for any possible inputs S and φ . Notice that we have defined π as a deterministic policy for simplicity, and our results hold for randomized policies. Let Π be the set of all possible adaptive policies.

Notice that an adaptive policy π completely specifies a seeding strategy in an iterative way. Given an adaptive policy π and a realization ϕ , let $\mathcal{S}^f(\pi, \phi, k)$ be the first k seeds selected according to π with the underlying realization ϕ under the full-adoption feedback model. By our definition, $\mathcal{S}^f(\pi, \phi, k)$ can be computed as follows:

1. initialize $S = \emptyset$;
2. update $S = S \cup \{\pi(S, \Phi_{G,F,\phi}^f(S))\}$ for k iterations;
3. output $\mathcal{S}^f(\pi, \phi, k) = S$.

Define $\mathcal{S}^m(\pi, \phi, k)$ similarly for the myopic feedback model, where $\Phi_{G,F,\phi}^m(S)$ instead of $\Phi_{G,F,\phi}^f(S)$ is used in Step 2 above.

Let $\sigma^f(\pi, k)$ be the expected number of infected vertices given that k seeds are chosen according to π , i.e., $\sigma^f(\pi, k) = \mathbb{E}_{\phi \sim F}[|I_{G,F}^\phi(\mathcal{S}^f(\pi, \phi, k))|]$. Define $\sigma^m(\pi, k)$ similarly for the myopic feedback model.

Definition 2.8. The *adaptive influence maximization problem* (adaptive INFMAX) is an optimization problem which takes as inputs $G = (V, E)$, F , and $k \in \mathbb{Z}^+$, and outputs an adaptive policy π that maximizes the expected total number of infections: $\pi \in \operatorname{argmax}_{\pi \in \Pi} \sigma^f(\pi, k)$ or $\pi \in \operatorname{argmax}_{\pi \in \Pi} \sigma^m(\pi, k)$ (depending on the feedback model used).

2.3 Adaptivity Gap and Greedy Adaptivity Gap

The adaptivity gap is defined as the ratio between the performance of the optimal adaptive policy and the performance of the optimal non-adaptive seeding strategy. In this paper, we only consider the adaptivity gap for triggering models.

Definition 2.9. The *adaptivity gap with full-adoption feedback* is

$$\sup_{G,F,k} \frac{\max_{\pi \in \Pi} \sigma^f(\pi, k)}{\max_{S \subseteq V, |S| \leq k} \sigma(S)}.$$

The *adaptivity gap with myopic feedback* is defined similarly.

The (non-adaptive) *greedy algorithm* iteratively picks a seed that has the maximum marginal gain to the objective function $\sigma(\cdot)$:

1. initialize $S = \emptyset$;
2. update for k iterations $S = S \cup \{s\}$, where $s \in \operatorname{argmax}_{s \in V} \sigma(S \cup \{s\})$ with tie broken in an arbitrarily consistent order;
3. return S .

Let $S^g(k)$ be the set of k seeds output by the (non-adaptive) greedy algorithm.

The *greedy adaptive policy* π^g is defined as $\pi^g(S, \varphi) = s$ such that

$$s \in \operatorname{argmax}_{s \in V} \mathbb{E}_{\phi \simeq \varphi} \left[|I_{G,F}^\phi(S \cup \{s\})| \right],$$

with tie broken in an arbitrary consistent order.

Definition 2.10. Given a triggering model $I_{G,F}$ and $k \in \mathbb{Z}^+$, the *greedy adaptivity gap with full-adoption feedback* is $\frac{\sigma^f(\pi^g, k)}{\sigma(S^g(k))}$. The *greedy adaptivity gap with myopic feedback* is defined similarly.

Notice that, unlike the adaptivity gap in Definition 2.9, we leave G, F, k unspecified (instead of taking a supremum over them) when defining the greedy adaptivity gap. This is because we are interested in both supremum and infimum of the ratio $\frac{\sigma^f(\pi^g, k)}{\sigma(S^g(k))}$. Notice that the infimum of the ratio $\frac{\max_{\pi \in \Pi} \sigma^f(\pi, k)}{\max_{S \subseteq V, |S| \leq k} \sigma(S)}$ in Definition 2.9 is 1: the optimal adaptive policy is at least as good as the optimal non-adaptive policy, as the non-adaptive policy can be viewed as a special adaptive policy; on the other hand, it is easy to see that there are INFMAX instances such that the optimal adaptive policy is no better than non-adaptive one (for example, a graph containing k vertices but no edges). For this reason, we only care about the supremum of this ratio.

3 Infimum of Greedy Adaptivity Gap

In this section, we show that the infimum of the greedy adaptivity gap for the triggering model is exactly $(1 - 1/e)$, for both the full-adoption feedback model and the myopic feedback model. This implies that the greedy adaptive policy can perform even worse than the conventional non-adaptive greedy algorithm, but it will never be significantly worse. Moreover, we show that this result also holds for both ICM (Definition 2.2) and LTM (Definition 2.3).

Theorem 3.1. *For the full-adoption feedback model,*

$$\inf_{G, F, k: I_{G, F} \text{ is ICM}} \frac{\sigma^f(\pi^g, k)}{\sigma(S^g(k))} = \inf_{G, F, k: I_{G, F} \text{ is LTM}} \frac{\sigma^f(\pi^g, k)}{\sigma(S^g(k))} = \inf_{G, F, k} \frac{\sigma^f(\pi^g, k)}{\sigma(S^g(k))} = 1 - \frac{1}{e}.$$

The same result holds for the myopic feedback model.

In Sect. 3.1, we show by providing examples that the greedy adaptive policy in the worst case will only achieves $(1 - 1/e + \varepsilon)$ -approximation of the expected number of infected vertices given by the non-adaptive greedy algorithm, for both ICM and LTM.

In Sect. 3.2, we shows that the greedy adaptive policy has performance at least $(1 - 1/e)$ of the performance of the non-adaptive optimal seeds (Theorem 3.5). Theorem 3.5 provides a lower bound on the greedy adaptivity gap for the triggering model and is also interesting on its own. At the end of Sect. 3.2, we prove Theorem 3.1 by putting the results from Sect. 3.1 and Sect. 3.2 together.

3.1 Tight Examples

In this subsection, we show that the adaptive greedy algorithm can perform worse than the non-adaptive greedy algorithm by a factor of $(1 - 1/e + \varepsilon)$, for both ICM and LTM and any $\varepsilon > 0$. This may be surprising, as one would expect that the feedback provided to the seed-picker will refine the seed choices in the future iterations. Here, we provide some intuitions why adaptivity can sometimes hurt. Suppose there are two promising sequences of seed selections, $\{s, u_1, \dots, u_k\}$ and $\{s, v_1, \dots, v_k\}$, such that

- s is the best seed which will be chosen first;
- $\{s, u_1, \dots, u_k\}$ has a better performance;

- the influence of u_1, \dots, u_k are non-overlapping, the influence of v_1, \dots, v_k are non-overlapping, but the influence of u_i, v_j overlaps for each i, j ; moreover, if u_1 is picked as the second seed, the greedy algorithm, adaptive or not, will continue to pick u_2, \dots, u_k , and if v_1 is picked as the second seed, v_2, \dots, v_k will be picked next;

Now, suppose there is a vertex t elsewhere which can be infected by both s and v_1 , such that

- if t is infected by s , which slightly reduces the marginal influence of v_1 , v_1 will be less promising than u_1 ;
- if t is not infected by s , v_1 is more promising than u_1 ;
- in average, when there is no feedback, v_1 is still less promising than u_1 , even after adding the increment in t 's infection probability to v_1 's expected marginal influence.

In this case, the non-adaptive greedy algorithm will “go to the right trend” by selecting u_1 as the second seed; the adaptive greedy algorithm, if receiving feedback that t is not infected by s , will “go to the wrong trend” by selecting v_1 next.

As a high-level description of the lesson we learned, both versions of the greedy algorithms are intrinsically myopic, and the feedback received by the adaptive policy may make the seed-picker act in a more myopic way, which could be more hurtful to the final performance.

We will assume in the rest of this section that vertices can have positive integer weights, as justified in the following remark.

Remark 3.2. For both ICM and LTM, we can assume without loss of generality that each vertex has a positive integer weight, so that, in INFMAX, we are maximizing the expected total weight of the infected vertices instead of maximizing the expected number of infected vertices as before. Suppose we want to make a vertex v have weight $W \in \mathbb{Z}^+$. We can construct $W - 1$ vertices w_1, \dots, w_{W-1} , and create $W - 1$ directed edges $(v, w_1), \dots, (v, w_{W-1})$ with weight 1. (Recall from Definition 2.2 and Definition 2.3 that the graphs in both ICM and LTM are edge-weighted, and the weights of edges completely characterize the collection of triggering set distributions F .) It is straightforward from Definition 2.2 and Definition 2.3 that, for both ICM and LTM, each of w_1, \dots, w_{W-1} will be infected with probability 1 if v is infected. In addition, both the greedy algorithm and the greedy adaptive policy will never pick any of w_1, \dots, w_{W-1} as seeds, as seeding v is strictly better. Therefore, we can consider the subgraph consisting of v, w_1, \dots, w_{W-1} as a gadget that representing a vertex v having weight W .

Lemma 3.3. *For any ε , there exists G, F, k such that $I_{G,F}$ is an ICM and*

$$\frac{\sigma^f(\pi^g, k)}{\sigma(S^g(k))} \leq 1 - \frac{1}{e} + \varepsilon, \quad \frac{\sigma^m(\pi^g, k)}{\sigma(S^g(k))} \leq 1 - \frac{1}{e} + \varepsilon.$$

Proof. We will construct an INFMAX instance $(G = (V, E, w), k + 1)$ with $k + 1$ seeds allowed. Let $W \in \mathbb{Z}^+$ be a sufficiently large perfect square divisible by k^{2k} and whose value are to be decided later. The vertex set V contains the following weighted vertices:

- a vertex s that has weight $2W$;
- a vertex t that has weight \sqrt{W}/k ;
- $2k$ vertices $u_1, \dots, u_k, v_1, \dots, v_k$ that have weight 1;
- $k(k + 1)$ vertices $\{w_{ij} \mid i = 1, \dots, k + 1; j = 1, \dots, k\}$

- w_{11}, \dots, w_{1k} have weight $\frac{W}{k}$;
- w_{i1}, \dots, w_{ik} have weight $\frac{1}{k}(1 - \frac{1}{k})^{i-1}W + \sqrt{W}$ for each $i = 2, \dots, k$;
- $w_{(k+1)1}, \dots, w_{(k+1)k}$ have weight $(1 - \frac{1}{k})^k W + \frac{\sqrt{W-k}}{k} - (k-1)\sqrt{W}$.

The edge set E is specified as follow:

- create two edges (v_1, t) and (s, t) ;
- for each $i = 1, \dots, k$, create $k+1$ edges $(u_i, w_{i1}), (u_i, w_{i2}), \dots, (u_i, w_{(k+1)i})$, and create k edges $(v_i, w_{i1}), (v_i, w_{i2}), \dots, (v_i, w_{ik})$.

For the weights of edges, all the edges have weight 1 except for the edge (s, t) which has weight $2k/\sqrt{W}$.

It is straightforward to check that $\sigma(\{s\}) = 2W + \frac{2k}{\sqrt{W}}\bar{w}(t) = 2W + 2$, $\sigma(\{u_i\}) = 1 + \sum_{j=1}^{k+1} \bar{w}(w_{ji}) = 1 + W + \frac{\sqrt{W-k}}{k}$ for each u_i , $\sigma(\{v_1\}) = 1 + \bar{w}(t) + \sum_{j=1}^k \bar{w}(w_{1j}) = 1 + \frac{\sqrt{W}}{k} + W$, $\sigma(\{v_i\}) = 1 + \sum_{j=1}^k \bar{w}(w_{ij}) = 1 + (1 - \frac{1}{k})^{i-1}W + k\sqrt{W}$ for each v_2, \dots, v_k , and the influence of the remaining vertices are significantly less than these.

Since s has the highest influence, both the greedy algorithm and the greedy adaptive policy will choose s as the first seed.

For the non-adaptive greedy algorithm, the next seed will be one of u_1, \dots, u_k , each of which contributes $1 + W + \frac{\sqrt{W-k}}{k}$ infected vertices. To see this, the only seed that is comparable to u_1, \dots, u_k is v_1 . However, since t will be infected by s (which has already been chosen as a seed) with probability $\frac{2k}{\sqrt{W}}$, the marginal contribution of v_1 will be slightly less than $\sigma(\{v_1\})$, and it will be $1 + W + (1 - \frac{2k}{\sqrt{W}})\frac{\sqrt{W}}{k} = 1 + W + \frac{\sqrt{W-2k}}{k}$, which is less than the marginal contribution of each of u_1, \dots, u_k . Since the influence of u_1, \dots, u_k are non-overlapping, it is straightforward to check that the non-adaptive greedy algorithm will choose $\{s, u_1, \dots, u_k\}$, which will infect vertices with a total weight of

$$\bar{w}(s) + \frac{2k}{\sqrt{W}}\bar{w}(t) + k + \sum_{i=1}^{k+1} \sum_{j=1}^k \bar{w}(w_{ij}) = (k+2)W + O(\sqrt{W})$$

in expectation.

The second seed picked by the greedy adaptive policy will depend on whether t is infected by s . Notice that the status of t is available to the policy in both the full-adoption feedback model and the myopic feedback model, so the arguments here, as well as the remaining part of this proof, apply to both feedback models. By straightforward calculations, the greedy adaptive policy will pick v_1 as the next seed if t is not infected by s , and the policy will pick a seed from u_1, \dots, u_k otherwise.

In the latter case, the policy will eventually pick the seed set $\{s, u_1, \dots, u_k\}$, which will infect vertices with a total weight of

$$\bar{w}(s) + \bar{w}(t) + k + \sum_{i=1}^{k+1} \sum_{j=1}^k \bar{w}(w_{ij}) = (k+2)W + O(\sqrt{W})$$

with probability 1 (notice that we are in the scenario that t has been infected by s).

In the former case, we can see that the third seed picked by the policy will be v_2 instead of any of u_1, \dots, u_k . In particular, v_2 contributes $(1 - \frac{1}{k})W + k\sqrt{W}$ infected vertices. On the other

hand, since w_{11}, \dots, w_{ik} have already been infected by v_1 , the marginal contribution for each u_i is $\sigma(\{u_i\}) - \bar{w}(w_{1i}) = (1 - \frac{1}{k})W + 1 + \frac{\sqrt{W}-k}{k}$, which is less than the contribution of v_2 . By similar analysis, we can see that the greedy adaptive policy in this case will pick the seed set $\{s, v_1, \dots, v_k\}$, which will infect vertices with a total weight of

$$\bar{w}(s) + \bar{w}(t) + k + \sum_{i=1}^k \sum_{j=1}^k \bar{w}(w_{ij}) = \left(2 + k \left(1 - \left(1 - \frac{1}{k}\right)^k\right)\right) W + O(\sqrt{W})$$

in expectation (notice that $w_{(k+1)1}, \dots, w_{(k+1)k}$ are not infected).

Since t will be infected with probability $\frac{2k}{\sqrt{W}}$, the expected weight of infected vertices for the greedy adaptive policy is

$$\begin{aligned} \frac{2k}{\sqrt{W}} \left((k+2)W + O(\sqrt{W}) \right) + \left(1 - \frac{2k}{\sqrt{W}}\right) \cdot \left(\left(2 + k \left(1 - \left(1 - \frac{1}{k}\right)^k\right)\right) W + O(\sqrt{W}) \right) \\ = \left(2 + k \left(1 - \left(1 - \frac{1}{k}\right)^k\right)\right) W + O(k^2\sqrt{W}). \end{aligned}$$

Putting together, both $\frac{\sigma^f(\pi^g, k)}{\sigma(S^g(k))}$ and $\frac{\sigma^m(\pi^g, k)}{\sigma(S^g(k))}$ in this case equal to

$$\frac{\left(2 + k \left(1 - \left(1 - \frac{1}{k}\right)^k\right)\right) W + O(k^2\sqrt{W})}{(k+2)W + O(\sqrt{W})},$$

which has limit $1 - 1/e$ when both W and k tend to infinity. □

Lemma 3.4. *For any ε , there exists G, F, k such that $I_{G,F}$ is an LTM and*

$$\frac{\sigma^f(\pi^g, k)}{\sigma(S^g(k))} \leq 1 - \frac{1}{e} + \varepsilon, \quad \frac{\sigma^m(\pi^g, k)}{\sigma(S^g(k))} \leq 1 - \frac{1}{e} + \varepsilon.$$

Proof. We will construct an INFMAX instance $(G = (V, E, w), k+1)$ with $k+1$ seeds allowed. Let $W \in \mathbb{Z}^+$ be a sufficiently large perfect square divisible by k^{2k} and whose value are to be decided later. The vertex set V contains the following weighted vertices:

- a vertex s that has weight $2W$;
- a vertex t that has weight \sqrt{W}/k ;
- k vertices u_1, \dots, u_k that have weight 1;
- k vertices v_1, \dots, v_k such that $\bar{w}(v_1) = W + 1$ and $\bar{w}(v_i) = W(1 - \frac{1}{k})^{i-1} + \sqrt{W}$ for each $i = 2, \dots, k$;
- k vertices v_{k+1}, \dots, v_{2k} such that v_{k+1}, \dots, v_{2k-1} have weight $W(1 - \frac{1}{k})^k$ and $\bar{w}(v_{2k}) = W(1 - \frac{1}{k})^k + \sqrt{W} - k - (k-1)\sqrt{W} - 1$.

The edge set E and the weights of edges are specified as follow:

- create two edges (v_1, t) and (s, t) with weights $1 - \frac{2k}{\sqrt{W}}$ and $\frac{2k}{\sqrt{W}}$ respectively;

- create $2k^2$ edges $\{(u_i, v_j) \mid i = 1, \dots, k; j = 1, \dots, 2k\}$, each of which has weight $\frac{1}{k}$.

It is easy to check that the weights of the incoming edges for each vertex v satisfy $\sum_{u \in \Gamma(v)} w(u, v) \leq 1$, as required in Definition 2.3.

The remaining part of the analysis is similar to the proof of Lemma 3.3. The first seed chosen by both algorithms is s . After this, each u_i has marginal influence $1 + \frac{1}{k} \sum_{i=1}^{2k} \bar{w}(v_i) = 1 + W + \frac{\sqrt{W}-k}{k}$. Since t is infected by s with probability $\frac{2k}{\sqrt{W}}$, the marginal influence of v_1 without any feedback is $(1 - \frac{2k}{\sqrt{W}})\bar{w}(t) + \bar{w}(v_1) = 1 + W + \frac{\sqrt{W}-2k}{k}$. If t is known to be infected, the marginal influence of v_1 is $1 + W$; if t is known to be uninfected, the marginal influence of v_1 is $1 + W + \frac{\sqrt{W}}{k}$. By comparing these values, the non-adaptive greedy algorithm will pick one of u_1, \dots, u_k as the second seed, and the greedy adaptive policy will pick v_1 as the second seed if t is not infected and one of u_1, \dots, u_k as the second seed if t is infected. (Notice that $\bar{w}(v_1) > \bar{w}(v_2) > \dots > \bar{w}(v_k) > \bar{w}(v_{k+1}) = \dots = \bar{w}(v_{2k-1}) > \bar{w}(v_{2k})$.)

Simple analyses show that the non-adaptive greedy algorithm will choose seeds $\{s, u_1, \dots, u_k\}$, which will infect all of v_1, \dots, v_{2k} with probability 1, and the adaptive greedy policy will choose $\{s, v_1, \dots, v_k\}$ with a very high probability $1 - \frac{2k}{\sqrt{W}}$, which will leave v_{k+1}, \dots, v_{2k} uninfected. Since s, v_1, \dots, v_{2k} are the only vertices with weight $\Theta(W)$ and we have both $\sum_{i=1}^k \bar{w}(v_i) = (1 - (1 - \frac{1}{k})^k)W + O(\sqrt{W})$ and $\sum_{i=1}^{2k} \bar{w}(v_i) = W + O(\sqrt{W})$, the lemma follows by taking the limit $W \rightarrow \infty$ and $k \rightarrow \infty$. \square

3.2 Lower Bound

Theorem 3.5. *For a triggering model $I_{G,F}$, we have both*

$$\sigma^f(\pi^g, k) \geq \left(1 - \frac{1}{e}\right) \max_{S \subseteq V, |S| \leq k} \sigma(S) \quad \text{and} \quad \sigma^m(\pi^g, k) \geq \left(1 - \frac{1}{e}\right) \max_{S \subseteq V, |S| \leq k} \sigma(S).$$

For a high-level idea of the proof, let S with $|S| = i$ be the seeds picked by π^g for the first i iterations and S^* be the optimal non-adaptive seed set: $S^* \in \operatorname{argmax}_{|S'| \leq k} \sigma(S')$. Given S as the existing seeds and any feedback (myopic or full-adoption) corresponding to S , we can show that the marginal increment to the expected influence caused by the $(i+1)$ -th seed picked by π^g is at least $1/k$ of the marginal increment to the expected influence caused by S^* . Then, a standard argument showing that the greedy algorithm can achieve a $(1 - 1/e)$ -approximation for any submodular monotone optimization problem can be used to prove this theorem.

Theorem 3.5 is implied by the following three propositions. In the remaining part of this section, we let S^* be an optimal seed set for the non-adaptive INFMAX: $S^* \in \operatorname{max}_{S \subseteq V, |S| \leq k} \sigma(S)$.

We first show that the global influence function after fixing a partial seed set S and any possible feedback of S is still submodular.

Proposition 3.6. *Given a triggering model $I_{G,F}$, any $S \subseteq V$, any feedback model (either full-adoption or myopic) and any partial realization φ that is a valid feedback of S (i.e., $\exists \phi : \varphi = \Phi_{G,F,\phi}^f(S)$ or $\exists \phi : \varphi = \Phi_{G,F,\phi}^m(S)$, depending on the feedback model considered), the function $\mathcal{T} : \{0, 1\}^{|V|} \rightarrow \mathbb{R}_{\geq 0}$ defined as $\mathcal{T}(X) = \mathbb{E}_{\phi \sim \varphi} [I_{G,F}^\phi(S \cup X)]$ is submodular.*

Proof. Fix a feedback model, $S \subseteq V$ and φ that is a valid feedback of S . Let \bar{S} be the set of infected vertices indicated by the feedback of S . Formally, \bar{S} is the set of all vertices that are reachable from S by only using edges e with $\varphi(e) = \text{L}$.

We consider a new triggering model $I_{G',F'}$ defined as follows:

- G' shares the same vertex set with G ;
- The edge set of G' is obtained by removing all edges e in G with $\varphi(e) \neq \mathbf{U}$;
- The distribution \mathcal{F}'_v is normalized from \mathcal{F}_v . Specifically, for each $T_v \subseteq \Gamma(v)$, let $p(T_v)$ be the probability that T_v is chosen as the triggering set under \mathcal{F}_v . Let $\Gamma'(v)$ be the set of v 's in-neighbors in G' , and we have $\Gamma'(v) \subseteq \Gamma(v)$ by our construction. Then, \mathcal{F}'_v is defined such that $T_v \subseteq \Gamma'(v)$ is chosen as the triggering set with probability $p(T_v) / \sum_{T'_v \subseteq \Gamma'(v)} p(T'_v)$.

A simple coupling argument reveals that

$$\mathcal{T}(X) = \mathbb{E}_{\phi \simeq \varphi} \left[\left| I_{G,F}^\phi(S \cup X) \right| \right] = \sigma_{G',F'}(\bar{S} \cup X). \quad (1)$$

We define a coupling of a realization ϕ of G with $\phi \simeq \varphi$ to a realization ϕ' of G' in a natural way: $\phi(e) = \phi'(e)$ for all edges e in G' . From our construction of $F' = \{\mathcal{F}'_v\}$, it is easy to see that, when ϕ is coupled with ϕ' , the probability that ϕ is sampled under $I_{G,F}$ conditioning on $\phi \simeq \varphi$ equals to the probability that ϕ' is sampled under $I_{G',F'}$. Under this coupling, it is easy to see that u is reachable from S by live edges under ϕ if and only if it is reachable from \bar{S} by live edges under ϕ' . This proves Eqn. (1).

Finally, by Theorem 2.4, $\sigma_{G',F'}(\cdot)$ is submodular. Therefore, for any two vertex sets A, B with $A \subsetneq B$ and any $u \notin B$,

$$\mathcal{T}(A \cup \{u\}) - \mathcal{T}(A) = \sigma_{G',F'}(\bar{S} \cup A \cup \{u\}) - \sigma_{G',F'}(\bar{S} \cup A)$$

is weakly larger than

$$\mathcal{T}(B \cup \{u\}) - \mathcal{T}(B) = \sigma_{G',F'}(\bar{S} \cup B \cup \{u\}) - \sigma_{G',F'}(\bar{S} \cup B)$$

if $u \notin \bar{S}$, and

$$\mathcal{T}(A \cup \{u\}) - \mathcal{T}(A) = \mathcal{T}(B \cup \{u\}) - \mathcal{T}(B) = 0$$

if $u \in \bar{S}$. In both case, the submodularity of $\mathcal{T}(\cdot)$ holds. \square

Next, we show that the marginal gain to the global influence function after selecting one more seed according to π^g is at least $1/k$ fraction of the marginal gain of including all the vertices in S^* as seeds.

Proposition 3.7. *Given a triggering model $I_{G,F}$, any $S \subseteq V$, any feedback model and any partial realization φ that is a valid feedback of S , let $s = \pi^g(S, \varphi)$ be the next seed chosen by the greedy policy. We have*

$$\mathbb{E}_{\phi \simeq \varphi} \left[\left| I_{G,F}^\phi(S \cup \{s\}) \right| \right] - \mathbb{E}_{\phi \simeq \varphi} \left[\left| I_{G,F}^\phi(S) \right| \right] \geq \frac{1}{k} \left(\mathbb{E}_{\phi \simeq \varphi} \left[\left| I_{G,F}^\phi(S \cup S^*) \right| \right] - \mathbb{E}_{\phi \simeq \varphi} \left[\left| I_{G,F}^\phi(S) \right| \right] \right).$$

Proof. Let $S^* = \{s_1^*, \dots, s_k^*\}$. By the greedy nature of π^g , we have

$$\forall v : \mathbb{E}_{\phi \simeq \varphi} \left[\left| I_{G,F}^\phi(S \cup \{s\}) \right| \right] \geq \mathbb{E}_{\phi \simeq \varphi} \left[\left| I_{G,F}^\phi(S \cup \{v\}) \right| \right],$$

and this holds for v being any of s_1^*, \dots, s_k^* in particular.

Let $S_i^* = \{s_1^*, \dots, s_i^*\}$ for each $i = 1, \dots, k$ and $S_0^* = \emptyset$, the proposition concludes from the following calculations

$$\begin{aligned}
& \mathbb{E}_{\phi \simeq \varphi} \left[\left| I_{G,F}^\phi(S \cup \{s\}) \right| \right] - \mathbb{E}_{\phi \simeq \varphi} \left[\left| I_{G,F}^\phi(S) \right| \right] \\
& \geq \frac{1}{k} \sum_{i=1}^k \left(\mathbb{E}_{\phi \simeq \varphi} \left[\left| I_{G,F}^\phi(S \cup \{s_i^*\}) \right| \right] - \mathbb{E}_{\phi \simeq \varphi} \left[\left| I_{G,F}^\phi(S) \right| \right] \right) \\
& \geq \frac{1}{k} \sum_{i=1}^k \left(\mathbb{E}_{\phi \simeq \varphi} \left[\left| I_{G,F}^\phi(S \cup S_{i-1}^* \cup \{s_i^*\}) \right| \right] - \mathbb{E}_{\phi \simeq \varphi} \left[\left| I_{G,F}^\phi(S \cup S_{i-1}^*) \right| \right] \right) \quad (\text{Proposition 3.6}) \\
& = \frac{1}{k} \left(\mathbb{E}_{\phi \simeq \varphi} \left[\left| I_{G,F}^\phi(S \cup S^*) \right| \right] - \mathbb{E}_{\phi \simeq \varphi} \left[\left| I_{G,F}^\phi(S) \right| \right] \right),
\end{aligned}$$

where the last equality is by a telescoping sum, by noticing that $S_i^* = S_{i-1}^* \cup \{s_i^*\}$ and $S^* = S_k^*$. \square

Finally, we prove the following proposition which is a more general statement than Theorem 3.5.

Proposition 3.8. *For a triggering model $I_{G,F}$ and any $\ell \in \mathbb{Z}^+$, we have $\sigma^f(\pi^g, \ell) \geq (1 - (1 - 1/k)^\ell) \sigma(S^*)$, and the same holds for the myopic feedback model.*

Proof. We will only consider the full-adoption feedback model, as the proof for the myopic feedback model is identical. We prove this by induction on ℓ . The base step for $\ell = 1$ holds trivially by Proposition 3.7 by considering $S = \emptyset$ in the proposition.

Suppose the inequality holds for $\ell = \ell_0$. We investigate the expected marginal gain to the global influence function by selecting the $(\ell_0 + 1)$ -th seed. For a seed set $S \subseteq V$ with $|S| = \ell_0$ and a partial realization φ , let $P(S, \varphi)$ be the probability that the policy π^g chooses S as the first ℓ_0 seeds and φ is the feedback. That is, $P(S, \varphi) = \Pr_{\phi \sim F} \left(S^f(\pi^g, \phi, \ell_0) = S \wedge \Phi_{G,F,\phi}^f(S) = \varphi \right)$. The mentioned expected marginal gain is

$$\begin{aligned}
& \sigma^f(\pi^g, \ell_0 + 1) - \sigma^f(\pi^g, \ell_0) \\
& = \sum_{S, \varphi: |S| = \ell_0} P(S, \varphi) \left(\mathbb{E}_{\phi \simeq \varphi} \left[\left| I_{G,F}^\phi(S \cup \{\pi^g(S, \varphi)\}) \right| \right] - \mathbb{E}_{\phi \simeq \varphi} \left[\left| I_{G,F}^\phi(S) \right| \right] \right) \\
& \geq \sum_{S, \varphi: |S| = \ell_0} P(S, \varphi) \cdot \frac{1}{k} \left(\mathbb{E}_{\phi \simeq \varphi} \left[\left| I_{G,F}^\phi(S \cup S^*) \right| \right] - \mathbb{E}_{\phi \simeq \varphi} \left[\left| I_{G,F}^\phi(S) \right| \right] \right) \quad (\text{Proposition 3.7}) \\
& \geq \sum_{S, \varphi: |S| = \ell_0} P(S, \varphi) \cdot \frac{1}{k} \left(\mathbb{E}_{\phi \simeq \varphi} \left[\left| I_{G,F}^\phi(S^*) \right| \right] - \mathbb{E}_{\phi \simeq \varphi} \left[\left| I_{G,F}^\phi(S) \right| \right] \right) \\
& = \frac{1}{k} \sigma(S^*) - \frac{1}{k} \sigma^f(\pi^g, \ell_0),
\end{aligned}$$

where the last equality follows from the law of total probability.

By rearranging the above inequality and the induction hypothesis,

$$\begin{aligned}
\sigma^f(\pi^g, \ell_0 + 1) & \geq \frac{1}{k} \sigma(S^*) + \frac{k-1}{k} \sigma^f(\pi^g, \ell_0) \\
& \geq \left(\frac{1}{k} + \frac{k-1}{k} \left(1 - \left(1 - \frac{1}{k} \right)^{\ell_0} \right) \right) \sigma(S^*) \\
& = \left(1 - \left(1 - \frac{1}{k} \right)^{\ell_0 + 1} \right) \sigma(S^*),
\end{aligned}$$

which concludes the inductive step. \square

By taking $\ell = k$ and noticing that $1 - (1 - 1/k)^k > 1 - 1/e$, it is easy to see that Proposition 3.8 implies Theorem 3.5.

Finally, putting Theorem 3.5, Lemma 3.3 and Lemma 3.4 together, Theorem 3.1 can be concluded easily.

Proof of Theorem 3.1. Since ICM and LTM are special cases of triggering models, we have

$$\inf_{G,F,k: I_{G,F} \text{ is ICM}} \frac{\sigma^f(\pi^g, k)}{\sigma(S^g(k))} \geq \inf_{G,F,k} \frac{\sigma^f(\pi^g, k)}{\sigma(S^g(k))}$$

and

$$\inf_{G,F,k: I_{G,F} \text{ is LTM}} \frac{\sigma^f(\pi^g, k)}{\sigma(S^g(k))} \geq \inf_{G,F,k} \frac{\sigma^f(\pi^g, k)}{\sigma(S^g(k))}.$$

Lemma 3.3 and Lemma 3.4 show that both

$$\inf_{G,F,k: I_{G,F} \text{ is ICM}} \frac{\sigma^f(\pi^g, k)}{\sigma(S^g(k))} \quad \text{and} \quad \inf_{G,F,k: I_{G,F} \text{ is LTM}} \frac{\sigma^f(\pi^g, k)}{\sigma(S^g(k))}$$

are at most $1 - 1/e$. On the other hand, Theorem 3.5 implies

$$\frac{\sigma^f(\pi^g, k)}{\sigma(S^g(k))} \geq \frac{\sigma^f(\pi^g, k)}{\sigma(S^*)} \geq 1 - \frac{1}{e}$$

for any triggering model $I_{G,F}$ and any k , where S^* , as usual, denotes the optimal seeds in the non-adaptive setting.

Putting together, Theorem 3.1 concludes for the full-adoption feedback model. Since all those inequalities hold for the myopic feedback model as well, Theorem 3.1 concludes for all feedback models. \square

4 Supremum of Greedy Adaptivity Gap

In this section, we show that, for the full-adoption feedback model, both the adaptivity gap and the supremum of the greedy adaptivity gap are unbounded. As a result, in some cases, the adaptive version of the greedy algorithm can perform significantly better than its non-adaptive counterpart.

Theorem 4.1. *The greedy adaptivity gap with full-adoption feedback is unbounded: there exists a triggering model $I_{G,F}$ and k such that*

$$\frac{\sigma^f(\pi^g, k)}{\sigma(S^g(k))} = 2^{\Omega(\log \log |V| / \log \log \log |V|)}.$$

Theorem 4.2. *The adaptivity gap for the general triggering model with full-adoption feedback is infinity.*

In Sect. 4.1, we consider a variant of INFMAX such that the seeds can only be chosen among a prescribed vertex set $\overline{V} \subseteq V$, where \overline{V} is specified as an input to the INFMAX instance. We show that, under this setting with LTM, both the adaptivity gap and the supremum of the greedy adaptivity gap with the full-adoption feedback model are unbounded (Lemma 4.5). Since it is

common in practice that only a subset of nodes in a network is visible or accessible to the seed-picker, Lemma 4.5 is also interesting on its own. In Sect. 4.2, we show that how Lemma 4.5 can be used to prove Theorem 4.1 and Theorem 4.2. Notice that Theorem 4.1 and Theorem 4.2 hold for the standard INFMAX setting without a prescribed set of seed candidates, but we do not know if they hold for LTM (instead, they are for the more general triggering model).

We first present the following lemma revealing a special additive property for LTM, which will be used later.

Lemma 4.3. *Suppose $I_{G,F}$ is LTM. If $U_1, U_2 \subseteq V$ with $U_1 \cap U_2 = \emptyset$ satisfy that there is no path from any vertices in U_1 to any vertices in U_2 and vice versa, then $\sigma(U_1) + \sigma(U_2) = \sigma(U_1 \cup U_2)$.*

Proof. For any seed set $S \subseteq V$, $\sigma(S)$ can be written as follows:

$$\sigma(S) = \sum_{\phi} \Pr(\phi \text{ is sampled}) \cdot \left| I_{G,F}^{\phi}(S) \right|. \quad (2)$$

For U_1 and U_2 in the lemma statement, since each vertex can only have at most one incoming live edge (in Definition 2.3, each T_v has size at most 1), under any realization ϕ , each vertex $v \in V \setminus (U_1 \cup U_2)$ that is reachable from vertices in $U_1 \cup U_2$ is reachable from either vertices in U_1 or vertices in U_2 , but not both. Therefore, $|I_{G,F}^{\phi}(U_1)| + |I_{G,F}^{\phi}(U_2)| = |I_{G,F}^{\phi}(U_1 \cup U_2)|$ for any ϕ , and the lemma follows from considering the decomposition of $\sigma(U_1)$ and $\sigma(U_2)$ according to (2). \square

4.1 On LTM with Prescribed Seed Candidates

Definition 4.4. The *influence maximization problem with prescribed seed candidates* is an optimization problem which takes as inputs $G = (V, E)$, F , $k \in \mathbb{Z}^+$, and $\overline{V} \subseteq V$, and outputs a seed set $S \subseteq \overline{V}$ that maximizes the expected total number of infections: $S \in \arg\max_{S \subseteq \overline{V}: |S| \leq k} \sigma(S)$. The *adaptive influence maximization problem with prescribed seed candidates* has the same definition as it is in Definition 2.8, with the exception that the range of the function π is now \overline{V} , and Π is the set of all such policies.

Lemma 4.5. *For INFMAX with prescribed seed candidates with LTM and full-adoption feedback, the adaptivity gap is infinity, and the greedy adaptivity gap is $2^{\Omega(\log |V| / \log \log |V|)}$.*

Proof. For $d, W \in \mathbb{Z}^+$ with d being sufficiently large and $W \gg d^{d+1}$, we construct the following (adaptive) INFMAX instance with prescribed seed candidates:

- the edge-weighted graph $G = (V, E, w)$ consists of an $(d + 1)$ -level directed full d -ary tree with the root node being the sink (i.e., an in-arborescence) and W vertices each of which is connected *from* the root node of the tree; the weight of each edge in the tree is $1/d$, and the weight of each edge connecting from the root to those W vertices is 1;
- the number of seeds is given by $k = 2(\frac{d+1}{2})^d$;
- the prescribed set for seed candidates \overline{V} is the set of all the leaves in the tree.

Since the leaves are not reachable from one to another, Lemma 4.3 indicates that choosing any k vertices among \overline{V} , i.e., the leaves, infects the same number of vertices in expectation. It is easy to see that a single seed among the leaves will infect the root node with probability $1/d^d$, and those W vertices will be infected with probability 1 if the root of the tree is infected. Thus, for any seed set $S \subseteq \overline{V}$, by assuming all vertices in the tree are infected (in the sake of finding an upper bound

for $\sigma(S)$), we have $\sigma(S) \leq \frac{1}{d^d} \cdot |S| \cdot W + \sum_{i=0}^d d^i < \frac{|S|W}{d^d} + d^{d+1}$. This gives an upper bound for the performance of both the non-adaptive greedy algorithm and the non-adaptive optimal seed set.

Now, we consider the greedy adaptive policy. If the root is not infected, there always exists a path from a certain leaf to the root such that all vertices on the path are not infected. This is because, if all children of an internal node w are infected, w will be infected with probability 1 (as $f_w = d \times \frac{1}{d} = 1$ which will always be no smaller than θ_w). In other words, if an internal node is uninfected, at least one of its children is uninfected. It is easy to see that, before the root is infected, the greedy adaptive policy will always choose a leaf such that all vertices on the path between the leaf and the root are uninfected.

Next, we evaluate the expected number of seeds required to infect the root, under the greedy adaptive policy. Suppose the tree only has two levels (i.e., a star). The number of seeds among the leaves required to infect the root is a random variable with uniform distribution on $\{1, \dots, d\}$, with expectation $\frac{d+1}{2}$. By induction on the number of levels of the tree, with a d -level tree as it is in our case, the expected number of seeds required to infect the root is $(\frac{d+1}{2})^d$, which equals to $\frac{k}{2}$. By Markov's inequality, the k seeds chosen according to the greedy adaptive policy will infect the root with probability at least $1/2$. Therefore, $\sigma^f(\pi^g, k) \geq \frac{1}{2}W$, and the optimal adaptive policy can only be better: $\max_{\pi \in \Pi} \sigma^f(\pi, k) \geq \sigma^f(\pi^g, k) \geq \frac{1}{2}W$.

Putting together, both the adaptivity gap and the supremum of the greedy adaptivity gap is at least

$$\frac{\frac{1}{2}W}{\frac{kW}{d^d} + d^{d+1}} = \frac{\frac{1}{2}W}{\frac{1}{2^{d-1}}(1 + \frac{1}{d})^d W + d^{d+1}} = \Omega\left(2^d\right),$$

if setting $W = d^{d+10} \gg d^{d+1}$. The lemma concludes by noticing $d = \Omega(\frac{\log |V|}{\log \log |V|})$ (in particular, $|V| = W + o(W) = d^{d+10} + o(d^{d+10})$, so $\log |V| = d \log d + o(d \log d)$, $\log \log |V| = \log d + o(\log d)$, and $d = \Omega(\frac{\log |V|}{\log \log |V|})$ as claimed). \square

4.2 Proof of Theorem 4.1, 4.2

To prove Theorem 4.1 and Theorem 4.2, we construct an INFMAX instance with a special triggering model $I_{G,F}$ which is a combination of ICM and LTM.

Definition 4.6. The *mixture of ICM and LTM* is a triggering model $I_{G,F}$ where $G = (V, E, w)$ is an edge-weighted graph with $w(u, v) \in (0, 1]$ for each $(u, v) \in E$ and each vertex v is labelled either **IC** or **LT** such that T_v is sampled according to \mathcal{F}_v described in Definition 2.2 if v is labelled **IC** and T_v is sampled according to \mathcal{F}_v described in Definition 2.3 if v is labelled **LT**. In addition, each vertex v labelled **L** satisfies $\sum_{u \in \Gamma(v)} w(u, v) \leq 1$.

To conclude Theorem 4.1 and Theorem 4.2, we construct an edge-weighted graph $G = (V, E, w)$ on which the greedy adaptive policy significantly outperforms the non-adaptive greedy algorithm. Let $M \gg d^{d+1}$ be a large integer. We reuse the graph with a tree and W vertices in the proof of Lemma 4.5. We create M such graphs and name them T_1, \dots, T_M . Let $L = d^d$ be the number of leaves in each T_i . Let $\mathbb{Z}_L = \{1, \dots, L\}$ and \mathbb{Z}_L^M be the set of all M -dimensional vectors whose entries are from \mathbb{Z}_L . For each $\mathbf{z} = (z_1, \dots, z_M) \in \mathbb{Z}_L^M$, create a vertex $a_{\mathbf{z}}$ and create a directed edge from $a_{\mathbf{z}}$ to the z_i -th leaf of the tree T_i for each $i = 1, \dots, M$. The weight of each such edge is 1. Let $A = \{a_{\mathbf{z}} \mid \mathbf{z} \in \mathbb{Z}_L^M\}$. Notice that $|A| = L^M$ and each $a_{\mathbf{z}} \in A$ is connected to M vertices from T_1, \dots, T_M respectively. The leaves of T_1, \dots, T_M are labelled as **IC**, and the remaining vertices are labelled as **LT**. Finally, set $k = 2(\frac{d+1}{2})^d$ as before.

Due to that M is large, it is more beneficial to seed a vertex in A than a vertex elsewhere. In particular, seeding a root in certain T_i infects W vertices, while seeding a vertex in A will infects $M \cdot (\frac{1}{d})^d W \gg W$ vertices in expectation.

It is easy to see that, in the non-adaptive setting, the optimal seeding strategy is to choose k seeds from A such that they do not share any out-neighbors, in which case the k chosen seeds will cause the infection of exactly k leaves in each T_i . This is also what the non-adaptive greedy algorithm will do. As before, to find an upper bound for any seed set S with $|S| = k$, we assume that all vertices in each T_i are infected, and we have $\sigma(S) \leq M \left(k \cdot \frac{1}{d^d} W + \sum_{i=0}^d d^i \right)$.

For the similar reason as it is in the proof of Lemma 4.5, the greedy adaptive policy would iteratively pick a seed $a_z \in A$ such that, for each $i = 1, \dots, M$, if the root in T_i is not yet infected, the path from z_i -th leaf to the root in T_i contains no infected vertices. By the same analysis in the proof of Lemma 4.5, by choosing k seeds among A as described above, which is equivalent as choosing k leaves in each of T_1, \dots, T_M simultaneously, the root in each T_i is infected with probability at least $\frac{1}{2}$. Therefore, the expected total number of infected vertices is at least $M \cdot \frac{1}{2} W$.

It may seem problematic that the greedy adaptive policy may start to seed the roots among T_1, \dots, T_M when it sees that there are already a lot of infected roots (so seeding a root is better than seed a vertex in A). However, we can always choose M to be large enough such that, after choosing k seeds adaptively from A , the total number of uninfected roots is still significantly more than d^d with high probability, so that seeding a vertex from A is still more beneficial. This is always possible: supposing p is the probability that a root is infected with k adaptive seeds from A , we only need to choose M such that $M \cdot (1 - p) \gg d^d$.

Putting together as before, both the adaptivity gap and the supremum of the greedy adaptivity gap is at least

$$\frac{M \cdot \frac{1}{2} W}{M(\frac{kW}{d^d} + d^{d+1})} = \frac{\frac{1}{2} W}{\frac{1}{2^{d-1}}(1 + \frac{1}{d})^d W + d^{d+1}} = \Omega(2^d),$$

if fixing $W = d^{d+10} \gg d^{d+1}$. Theorem 4.2 concludes by letting $d \rightarrow \infty$. To conclude Theorem 4.1, we need to show that $d = \Omega(\log \log |V| / \log \log \log |V|)$. To see this, we set $M = d^{d+10}$ which is sufficiently large for our purpose. Since we have $L = d^d$, we have $|V| = L^M + o(L^M) = d^{d^{d+11}} + o(d^{d^{d+11}})$, which implies $d = \Omega(\log \log |V| / \log \log \log |V|)$.

5 A Variant of Greedy Adaptive Policy

Although we have seen that the adaptive version of the greedy algorithm can perform worse than its non-adaptive counterpart, in general, we would still recommend the use of it as long as it is feasible, as it can also perform significantly better than the non-adaptive greedy algorithm (Theorem 4.1) while never being too bad (Theorem 3.5). As we remarked, the adaptivity may be harmful because exploiting the feedback may make the seed-picker too myopic. In this section, we propose a less aggressive risk-free version of the greedy adaptive policy, π^{g-} , in that it balances between the exploitation of the feedback and the focus on the average in the conventional non-adaptive greedy algorithm.

First, we apply the non-adaptive greedy algorithm with $|V|$ seeds to obtain an order \mathcal{L} on all vertices. Then for any $S \subseteq V$ and any partial realization φ , $\pi^{g-}(S, \varphi)$ is defined to be the first vertex v in \mathcal{L} that is not known to be infected. Formally, v is the first vertex in \mathcal{L} that are not reachable from S when removing all edges e with $\varphi(e) \in \{\mathbf{B}, \mathbf{U}\}$. This finishes the description of the policy.

This adaptive policy is always no worse than the non-adaptive greedy algorithm, as it is easy to see that those seeds chosen by π^g are either seeded or infected by previously selected seeds in π^{g-} .

However, π^{g-} can sometimes be conservative. It is possible that π^{g-} has the same performance as the non-adaptive greedy algorithm, but π^g is much better. Especially, when there is no path between any two vertices among the first k vertices in \mathcal{L} , π^{g-} will make the same choice as the non-adaptive greedy algorithm. The INFMAX instance in Sect. 4.2 is an example of this.

We have seen that π^{g-} sometimes performs better than π^g (e.g., in those instances constructed in the proofs of Lemma 3.2 and Lemma 3.3) and sometimes performs worse than the π^g (e.g., in the instance constructed in Sect. 4.2). Therefore, given a *particular* INFMAX instance, for deciding which of π^{g-} and π^g to be used (we should never consider the non-adaptive greedy algorithm if adaptivity is available, as it is always weakly worse than π^{g-}), we recommend a comparison of the two policies by simulations. Notice that the seed-picker can randomly sample a realization ϕ and simulate the feedback the policy will receive. Thus, given $I_{G,F}$, both π^{g-} and π^g can be estimated by taking an average over the numbers of infected vertices in a large number of simulations.

6 Conclusion and Open Problems

We have seen that the infimum of the greedy adaptivity gap is exactly $(1 - 1/e)$ for ICM, LTM, and general triggering models with both the full-adoption feedback model and the myopic feedback model. We have also seen that the supremum of this gap is infinity for the full-adoption feedback model. One natural open problem is to find the supremum of the greedy adaptivity gap for the myopic feedback model. Another natural open problem is to find the supremum of the greedy adaptivity gap for the more specific ICM and LTM.

The greedy adaptivity gap studied in this paper is closely related to the adaptivity gap studied in the past. Since the non-adaptive greedy algorithm is always a $(1 - 1/e)$ -approximation of the non-adaptive optimal solution, a constant adaptivity gap implies a constant greedy adaptivity gap. For example, the adaptivity gap for ICM with myopic feedback is at most 4 [26], so the greedy adaptivity gap in the same setting is at most $\frac{4}{1-1/e}$. In addition, the greedy adaptive policy is known to achieve a $(1 - 1/e)$ -approximation to the adaptive optimal solution for ICM with full-adoption feedback [14], so the adaptivity gap and the greedy adaptivity gap could either be both constant or both unbounded for ICM with full-adoption feedback model, but it remains open which case is true. The adaptivity gap for ICM with full-adoption feedback, as well as the adaptivity gap for LTM with both feedback models, are all important open problems. We believe these problems can be studied together with the greedy adaptivity gap.

References

- [1] Rico Angell and Grant Schoenebeck. Don't be greedy: Leveraging community structure to find high quality seed sets for influence maximization. *WINE*, 2016.
- [2] Akhil Arora, Sainyam Galhotra, and Sayan Ranu. Debunking the myths of influence maximization. In *Proceedings of the 2017 ACM International Conference on Management of Data-SIGMOD'17*, 2017.
- [3] S Bharathi, D Kempe, and M Salek. Competitive influence maximization in social networks. In *WINE*, 2007.

- [4] Christian Borgs, Michael Brautbar, Jennifer Chayes, and Brendan Lucier. Maximizing social influence in nearly optimal time. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 946–957. SIAM, 2014.
- [5] Wei Chen and Binghui Peng. On adaptivity gaps of influence maximization under the independent cascade model with full adoption feedback. Technical report, arXiv:1907.01707, 2019.
- [6] Wei Chen, Yajun Wang, and Siyu Yang. Efficient influence maximization in social networks. In *ACM SIGKDD*, pages 199–208. ACM, 2009.
- [7] Wei Chen, Yifei Yuan, and Li Zhang. Scalable influence maximization in social networks under the linear threshold model. In *2010 IEEE International Conference on Data Mining*, pages 88–97. IEEE, 2010.
- [8] Wei Chen, Yifei Yuan, and Li Zhang. Scalable influence maximization in social networks under the linear threshold model. In *Data Mining (ICDM), 2010 IEEE 10th International Conference on*, pages 88–97. IEEE, 2010.
- [9] Wei Chen, Laks VS Lakshmanan, and Carlos Castillo. Information and influence propagation in social networks. *Synthesis Lectures on Data Management*, 5(4):1–177, 2013.
- [10] Suqi Cheng, Huawei Shen, Junming Huang, Guoqing Zhang, and Xueqi Cheng. Staticgreedy: solving the scalability-accuracy dilemma in influence maximization. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*, pages 509–518. ACM, 2013.
- [11] P. Domingos and M. Richardson. Mining the network value of customers. In *ACM SIGKDD*, 2001.
- [12] Sainyam Galhotra, Akhil Arora, and Shourya Roy. Holistic influence maximization: Combining scalability and efficiency with opinion-aware models. In *Conference on Management of Data*, pages 743–758. ACM, 2016.
- [13] S Goldberg and Z Liu. The diffusion of networking technologies. In *SODA*, 2013.
- [14] Daniel Golovin and Andreas Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *Journal of AI Research*, 42:427–486, 2011.
- [15] Amit Goyal, Wei Lu, and Laks VS Lakshmanan. Celf++: optimizing the greedy algorithm for influence maximization in social networks. In *Proceedings of the 20th international conference WWW*, pages 47–48. ACM, 2011.
- [16] Amit Goyal, Wei Lu, and Laks VS Lakshmanan. Simpath: An efficient algorithm for influence maximization under the linear threshold model. In *Data Mining (ICDM), 2011 IEEE 11th International Conference on*, pages 211–220. IEEE, 2011.
- [17] 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.
- [18] Kyomin Jung, Wooram Heo, and Wei Chen. Irie: Scalable and robust influence maximization in social networks. In *Data Mining (ICDM), 2012 IEEE 12th International Conference on*, pages 918–923. IEEE, 2012.

- [19] David Kempe, Jon M. Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *ACM SIGKDD*, pages 137–146, 2003.
- [20] David Kempe, Jon M. Kleinberg, and Éva Tardos. Influential nodes in a diffusion model for social networks. In *ICALP*, pages 1127–1138, 2005.
- [21] Jure Leskovec, Andreas Krause, Carlos Guestrin, Christos Faloutsos, Jeanne VanBriesen, and Natalie 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. ACM, 2007.
- [22] Yuchen Li, Ju Fan, Yanhao Wang, and Kian-Lee Tan. Influence maximization on social graphs: A survey. *IEEE Trans. on Knowledge and Data Engineering*, 30(10):1852–1872, 2018.
- [23] Elchanan Mossel and Sébastien Roch. Submodularity of influence in social networks: From local to global. *SIAM J. Comput.*, 39(6):2176–2188, 2010.
- [24] George L Nemhauser, Laurence A Wolsey, and Marshall L Fisher. An analysis of approximations for maximizing submodular set functions. *Mathematical Programming*, 14(1):265–294, 1978.
- [25] Naoto Ohsaka, Takuya Akiba, Yuichi Yoshida, and Ken-ichi Kawarabayashi. Fast and accurate influence maximization on large networks with pruned monte-carlo simulations. In *AAAI*, pages 138–144, 2014.
- [26] Binghui Peng and Wei Chen. Adaptive influence maximization with myopic feedback. Technical report, arXiv:1905.11663, 2019.
- [27] M. Richardson and P. Domingos. Mining knowledge-sharing sites for viral marketing. In *ACM SIGKDD*, pages 61–70, 2002.
- [28] Grant Schoenebeck and Biaoshuai Tao. Beyond worst-case (in) approximability of nonsubmodular influence maximization. In *International Conference on Web and Internet Economics*, pages 368–382. Springer, 2017.
- [29] Grant Schoenebeck and Biaoshuai Tao. Influence maximization on undirected graphs: Towards closing the $(1 - 1/e)$ gap. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 423–453. ACM, 2019.
- [30] Grant Schoenebeck and Biaoshuai Tao. Beyond worst-case (in)approximability of nonsubmodular influence maximization. *ACM Trans. Comput. Theory*, 11(3):12:1–12:56, April 2019. ISSN 1942-3454. doi: 10.1145/3313904. URL <http://doi.acm.org/10.1145/3313904>.
- [31] Grant Schoenebeck, Biaoshuai Tao, and Fang-Yi Yu. Think globally, act locally: On the optimal seeding for nonsubmodular influence maximization. In *RANDOM 2019: International Conference on Randomization and Computation*, 2019.
- [32] Jing Tang, Xueyan Tang, Xiaokui Xiao, and Junsong Yuan. Online processing algorithms for influence maximization. In *Proceedings of the 2018 International Conference on Management of Data*, pages 991–1005. ACM, 2018.
- [33] Youze Tang, Xiaokui Xiao, and Yanchen Shi. Influence maximization: Near-optimal time complexity meets practical efficiency. In *SIGMOD international conference on Management of data*, pages 75–86. ACM, 2014.

- [34] 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. ACM, 2015.

A Some Intuitions for ICM and LTM

In the original definition, ICM is defined such that each vertex u attempts only once to infect each of its not-yet-infected out-neighbor v with probability $w(u, v)$.

Definition A.1. The *independent cascade model* IC_G is defined by a directed edge-weighted graph $G = (V, E, w)$ such that $w(u, v) \leq 1$ for each $(u, v) \in E$. On input seed set $S \subseteq V$, $IC_G(S)$ outputs a set of infected vertices as follows:

1. Initially, only vertices in S are infected.
2. In each subsequent round, each vertex u infected in the previous round infects each (not yet infected) out-neighbor v with probability $w(u, v)$ independently.
3. After a round where there is no additional infected vertices, $IC_G(S)$ outputs the set of infected vertices.

It is straightforward to see that this definition is equivalent to Definition 2.2.

The basic idea behind the original LTM is that the influence from the in-neighbors of a vertex is additive.

Definition A.2. The *linear threshold model* LT_G is defined by a directed edge-weighted graph $G = (V, E, w)$ such that $\sum_{u:u \in \Gamma(v)} w(u, v) \leq 1$ for each $v \in V$. On input seed set $S \subseteq V$, $LT_G(S)$ outputs a set of infected vertices as follows:

1. Initially, only vertices in S are infected, and for each vertex v a *threshold* θ_v is sampled uniformly at random from $[0, 1]$ independently.
2. In each subsequent round, a vertex v becomes infected if $\sum_{u:u \in \Gamma(v)} w(u, v) \geq \theta_v$ and u is infected.
3. After an round where there is no additional infected vertices, $LT_G(S)$ outputs the set of infected vertices.

Kempe et al. [19] showed that the definition above is equivalent to Definition 2.3. For an intuition of this, consider a not-yet-infected vertex v and a set of its infected neighbors $IN_v \subseteq \Gamma(v)$. v will be infected by vertices in IN_v with probability $\sum_{u:u \in IN_v} w(u, v)$, as $\Pr(\theta_v \leq \sum_{u:u \in IN_v} w(u, v)) = \sum_{u:u \in IN_v} w(u, v)$. In the case where v becomes infected, we can attribute its infection to exactly one of its infected neighbors. The infection will be attributed to neighboring infected vertex u with probability equal to $w(u, v)$ (in which case $T_v = \{u\}$). Overall, the probability that v includes an incoming edge from $\{(u, v) : u \in IN_v\}$ is exactly $\sum_{u:u \in IN_v} w(u, v)$.

B On General Threshold Model

In this section, we show that all our results in this paper hold for submodular general threshold model, a model that is more general than the triggering model. In Sect. B.1, we define the general threshold model, and we define the two feedback models, the full-adoption and the myopic, based on the general threshold model. In Sect. B.2, we justify that all our results in this paper hold for submodular general threshold model.

B.1 General Threshold Model and Feedback

Definition B.1 (Kempe et al. [19]). The *general threshold model*, $I_{G,F}$, is defined by a graph $G = (V, E)$ and for each vertex v a monotone *local influence function* $f_v : \{0, 1\}^{|\Gamma(v)|} \rightarrow [0, 1]$ with $f_v(\emptyset) = 0$. Let $F = \{f_v \mid v \in V\}$.

On an input seed set $S \subseteq V$, $I_{G,F}(S)$ outputs a set of infected vertices as follows:

1. Initially, only vertices in S are infected, and for each vertex v the threshold θ_v is sampled uniformly at random from the interval $(0, 1]$ independently.
2. In each subsequent round, a vertex v becomes infected if the influence of its infected in-neighbors, $IN_v \subseteq \Gamma(v)$, exceeds its threshold: $f_v(IN_v) \geq \theta_v$.
3. After a round where no additional vertices are infected, the set of infected vertices is the output.

$I_{G,F}$ in Definition B.1 can be viewed as a random function $I_{G,F} : \{0, 1\}^{|V|} \rightarrow \{0, 1\}^{|V|}$. In addition, if the thresholds of all the vertices are fixed, this function becomes deterministic. Correspondingly, we define a *realization* of a graph $G = (V, E)$ as a function $\phi : V \rightarrow (0, 1]$ which encodes the thresholds of all vertices. Let $I_{G,F}^\phi : \{0, 1\}^{|V|} \rightarrow \{0, 1\}^{|V|}$ be the deterministic function corresponding to the general threshold model $I_{G,F}$ with vertices' thresholds following realization ϕ . We will interchangeably consider ϕ as a function defined above or a $|V|$ dimensional vector in $(0, 1]^{|V|}$, and we write $\phi \sim (0, 1]^{|V|}$ to mean a random realization is sampled such that each θ_v is sampled uniformly at random and independently as it is in Definition B.1.

Like the triggering model, the general threshold model also captures the independent cascade and linear threshold models.

- ICM is a special case of the general threshold model $I_{G,F}$ where $G = (V, E, w)$ is an edge-weighted graph with $w(u, v) \in (0, 1]$ for each $(u, v) \in E$ and $f_v(IN_v) = 1 - \prod_{u \in IN_v} (1 - w(u, v))$ for each $f_v \in F$.
- LTM is a special case of the general threshold model $I_{G,F}$ where $G = (V, E, w)$ is an edge-weighted graph with $w(u, v) > 0$ for each $(u, v) \in E$ and $\sum_{u \in \Gamma(v)} w(u, v) \leq 1$ for each $v \in V$ and $f_v(IN_v) = \sum_{u \in IN_v} w(u, v)$ for each $f_v \in F$.

Given a general threshold model $I_{G,F}$, the global influence function is then defined as $\sigma_{G,F}(S) = \mathbb{E}_{\phi \sim (0, 1]^{|V|}} [I_{G,F}^\phi(S)]$. Mossel and Roch [23] showed that $\sigma(\cdot)$ is monotone and submodular if each $f_v(\cdot)$ is monotone and submodular. We normally say that a general threshold model $I_{G,F}$ is submodular if each $f_v \in F$ is submodular. Notice that this implies $\sigma(\cdot)$ is submodular.

In the remaining part of this section, we define the *full-adoption feedback model* and the *myopic feedback model* corresponding to the general threshold model.

When the seed-picker sees that a vertex v is not infected (v may be a vertex adjacent to $I_{G,F}^\phi(S)$ in the full-adoption feedback model, or a vertex adjacent to S in the myopic feedback model), the seed-picker has certain partial information about v 's threshold. Specifically, let IN_v be v 's infected in-neighbors that are observed by the seed-picker. By seeing that v is not infected, the seed-picker knows that the threshold of v is in the range $(f_v(IN_v), 1]$, and the posterior distribution of θ_v is the uniform distribution on this range.

Let the *level* of a vertex v , denoted by o_v , be a value which either equals a character \checkmark indicating that it is infected, or a real value $\vartheta_v \in [0, 1]$ indicating that $\theta_v \in (\vartheta_v, 1]$. Let $O = \{\checkmark\} \cup [0, 1]$ be the space of all possible levels. A *partial realization* φ is a function specifying a level for each vertex: $\varphi : V \rightarrow O$. We say that a partial realization φ is *consistent with* the full realization ϕ , denoted by $\phi \simeq \varphi$, if $\phi(v) > \varphi(v)$ for any $v \in V$ such that $\varphi(v) \neq \checkmark$.

Definition B.2. Given a general threshold model $I_{G=(V,E),F}$ with a realization ϕ , the *full-adoption feedback* is a function $\Phi_{G,F,\phi}^f$ mapping a seed set $S \subseteq V$ to a partial realization φ such that

- $\varphi(v) = \checkmark$ for each $v \in I_{G,F}^\phi(S)$, and
- $\varphi(v) = f_v(I_{G,F}^\phi(S) \cap \Gamma(v))$ for each $v \notin I_{G,F}^\phi(S)$.

Definition B.3. Given a general threshold model $I_{G=(V,E),F}$ with a realization ϕ , the *myopic feedback* is a function $\Phi_{G,F,\phi}^m$ mapping a seed set $S \subseteq V$ to a partial realization φ such that

- $\varphi(v) = \checkmark$ for each $v \in S$, and
- for each $v \notin S$, $\varphi(v) = \checkmark$ if $f_v(S \cap \Gamma(v)) \geq \phi(v)$, and $\varphi(v) = f_v(S \cap \Gamma(v))$ if $f_v(S \cap \Gamma(v)) < \phi(v)$.

Notice that, in both definitions above, a vertex v that does not have any infected neighbor (i.e., $v \notin S$ such that $I_{G,F}^\phi(S) \cap \Gamma(v) = \emptyset$ for the full-adoption feedback model or $S \cap \Gamma(v) = \emptyset$ for the myopic feedback model) always satisfies $\varphi(v) = 0$, as $f_v(\emptyset) = 0$ by Definition 2.1.

After properly defining the two feedback models, the definition of the adaptive policy π , as well as the definitions of the functions $\mathcal{S}^f(\cdot, \cdot, \cdot)$, $\mathcal{S}^m(\cdot, \cdot, \cdot)$, $\sigma^f(\cdot, \cdot)$, $\sigma^m(\cdot, \cdot)$, are exactly the same as they are in Sect. 2.2. The definitions of the adaptivity gap and the greedy adaptivity gap are also the same as they are in Sect. 2.3.

B.2 Extending of Our Results to General Threshold Model

We will show in this section that all our results can be extended to the submodular general threshold model. Recall that a general threshold model is submodular means that all the local influence functions f_v 's are submodular. In this section, whenever we write $I_{G,F}$, we refer to the general threshold model in Definition B.1, not the triggering model in Definition 2.1.

B.2.1 Infimum of Greedy Adaptivity Gap

Theorem 3.1 is extended as follows.

Theorem B.4. *For the full-adoption feedback model,*

$$\inf_{G,F,k: I_{G,F} \text{ is ICM}} \frac{\sigma^f(\pi^g, k)}{\sigma(S^g(k))} = \inf_{G,F,k: I_{G,F} \text{ is LTM}} \frac{\sigma^f(\pi^g, k)}{\sigma(S^g(k))} = \inf_{G,F,k: I_{G,F} \text{ is submodular}} \frac{\sigma^f(\pi^g, k)}{\sigma(S^g(k))} = 1 - \frac{1}{e}.$$

The same result holds for the myopic feedback model.

Recall that Theorem 3.1 can be easily implied by Lemma 3.3, Lemma 3.4 and Theorem 3.5. Since Lemma 3.3 and Lemma 3.4 are for specific models ICM and LTM which are compatible with both the triggering model and the general threshold model, their validity here is clear. Following the same arguments, Theorem B.4 can be implied by Lemma 3.3, Lemma 3.4 and the following theorem which is the counterpart to Theorem 3.5.

Theorem B.5. *If $I_{G,F}$ is a submodular general threshold model, then we have both*

$$\sigma^f(\pi^g, k) \geq \left(1 - \frac{1}{e}\right) \max_{S \subseteq V, |S| \leq k} \sigma(S) \quad \text{and} \quad \sigma^m(\pi^g, k) \geq \left(1 - \frac{1}{e}\right) \max_{S \subseteq V, |S| \leq k} \sigma(S).$$

Similar to the proof of Theorem 3.5, Theorem B.5 can be proved by showing the three propositions: Proposition 3.6, Proposition 3.7 and Proposition 3.8. It is straightforward to check that Proposition 3.7 and Proposition 3.8 hold for the general threshold model with exactly the same proofs. Now, it remains to extend Proposition 3.6 to the general threshold model, which is restated and proved below.

Proposition B.6. *Given a submodular general threshold model $I_{G,F}$, any $S \subseteq V$, any feedback model (either full-adoption or myopic) and any partial realization φ that is a valid feedback of S (i.e., $\exists \phi : \varphi = \Phi_{G,F,\phi}^f(S)$ or $\exists \phi : \varphi = \Phi_{G,F,\phi}^m(S)$, depending on the feedback model considered), the function $\mathcal{T} : \{0,1\}^{|V|} \rightarrow \mathbb{R}_{\geq 0}$ defined as $\mathcal{T}(X) = \mathbb{E}_{\phi \simeq \varphi} [I_{G,F}^\phi(S \cup X)]$ is submodular.*

Proof. Fix a feedback model, $S \subseteq V$ and φ that is a valid feedback of S . Let $T = \{v \mid \varphi(v) = \checkmark\}$ be the set of infected vertices indicated by the feedback of S . We consider a new general threshold model $I_{G',F'}$ defined as follows:

- G' is obtained by removing vertices in T from G (and the edges connecting from/to vertices in T are also removed);
- For any $v \in V' = V \setminus T$, $\Gamma(v) \cap T$ is the set of in-neighbors of v that are removed. Define $f'_v(Y) = \frac{f_v((\Gamma(v) \cap T) \cup Y) - \varphi(v)}{1 - \varphi(v)}$ for each subset Y of v 's in-neighbors in the new graph G' : $Y \subseteq \Gamma(v) \cap V'$.

Notice that f'_v is a valid local influence function. f'_v is clearly monotone. For each $v \in V'$, we have $\varphi(v) = f_v(\Gamma(v) \cap T)$, as this is exactly the feedback received from the fact that v has not yet infected. It is then easy to see that f'_v is always non-negative and $f'_v(\emptyset) = 0$.

A simple coupling argument can show that

$$\mathbb{E}_{\phi \simeq \varphi} [I_{G,F}^\phi(S \cup X)] = \sigma_{G',F'}(X \setminus T) + |T|. \quad (3)$$

To define the coupling, for each $v \in V'$, the threshold of v in G , θ_v , is coupled with the threshold of v in G' as $\theta'_v = \frac{\theta_v - \varphi(v)}{1 - \varphi(v)}$. This is a valid coupling: by $\phi \simeq \varphi$, we know that θ_v is sampled uniformly at random from $(\varphi(v), 1]$, which indicates that the marginal distribution of θ'_v is the uniform distribution on $(0, 1]$, which makes $I_{G',F'}$ a valid general threshold model.

Under this coupling, on the vertices V' , the cascade in G with seeds $S \cup X$ and partial realization φ is identical to the cascade in G' with seeds $X \setminus T$. To see this, consider an arbitrary vertex $v \in V'$ and let IN_v and IN'_v be v 's infected neighbors in G and G' respectively. For induction hypothesis, suppose the two cascade processes before v 's infection are identical. We have $IN_v = IN'_v \cup (\Gamma(v) \cap T)$ and $IN'_v \cap (\Gamma(v) \cap T) = \emptyset$. It is easy to see from our construction that v is infected in G if and only if v is infected in G' :

$$f_v(IN_v) \geq \theta_v \Leftrightarrow f'_v(IN'_v) = \frac{f_v(IN_v) - \varphi(v)}{1 - \varphi(v)} \geq \theta'_v.$$

This proves Eqn. (3).

Finally, since each $f_v(\cdot)$ is assumed to be submodular, it is easy to see that each $f'_v(\cdot)$ is submodular by our definition. Thus, $I_{G',F'}$ is a submodular model. This, combined with Eqn. (3), proves the proposition. \square

B.2.2 Supremum of Greedy Adaptivity Gap

All the results in Sect. 4 about the supremum of the greedy adaptivity gap can be extended easily to the submodular general threshold model. In particular, Lemma 4.3 and Lemma 4.5 are under LTM, which is compatible with the submodular general threshold model. Theorem 4.1 and Theorem 4.2 are proved by providing an example with a diffusion model that is a combination of ICM and LTM, and the diffusion model constructed in Definition 4.6 can be easily described in the formulation of the general threshold model, since both ICM and LTM can be described in the general threshold model.