

Online Influence Maximization under the Independent Cascade Model with Node-Level Feedback*

Zhijie Zhang,^{1,2} Wei Chen,³ Xiaoming Sun,^{1,2} Jialin Zhang^{1,2,†}

¹ Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China

² University of Chinese Academy of Sciences, Beijing, China

³ Microsoft Research Asia, Beijing, China

{zhangzhijie,sunxiaoming,zhangjialin}@ict.ac.cn, weic@microsoft.com

Abstract

We study the online influence maximization (OIM) problem in social networks, where the learner repeatedly chooses seed nodes to generate cascades, observes the cascade feedback, and gradually learns the best seeds that generate the largest cascade in multiple rounds. In the demand of the real world, we work with *node-level feedback* instead of the common *edge-level feedback* in the literature. The edge-level feedback reveals all edges that pass through information in a cascade, whereas the node-level feedback only reveals the activated nodes with timestamps. The node-level feedback is arguably more realistic since in practice it is relatively easy to observe who is influenced but very difficult to observe from which relationship (edge) the influence comes. Previously, there is a nearly optimal $\tilde{O}(\sqrt{T})$ -regret algorithm for OIM problem under the *linear threshold* (LT) diffusion model with node-level feedback. It remains unknown whether the same algorithm exists for the *independent cascade* (IC) diffusion model. In this paper, we resolve this open problem by presenting an $\tilde{O}(\sqrt{T})$ -regret algorithm for OIM problem under the IC model with node-level feedback.

1 Introduction

Social networks have gained great attention in the past decades as a model for describing relationships between humans. Typically, researchers show great interest in how information, ideas, news, influence, etc spread over social networks, starting from a small set of nodes called *seeds*. To this end, a variety of diffusion models are proposed to formulate the propagation in reality, and the most well-known ones are the *independent cascade* (IC) model and the *linear threshold* (LT) model (Kempe, Kleinberg, and Tardos

2003). A corresponding optimization problem, known as *influence maximization* (IM), asks how to maximize the influence spread, under a specific diffusion model, by selecting a limited number of “good” seeds. The problem has found enormous applications, including advertising, viral marketing, news transmission, etc.

In the canonical setting, the IM problem takes as input a social network, which is formulated as an edge-weighted directed graph. The problem is NP-hard but can be well-approximated (Kempe, Kleinberg, and Tardos 2003). For the past decade, more efficient and effective algorithms have been designed (Borgs et al. 2014; Tang, Xiao, and Shi 2014; Tang, Shi, and Xiao 2015), leading to an almost complete resolution of the problem. However, the canonical IM is sometimes difficult to apply in practice, as edge parameters of the network are often *unknown* in many scenarios. A possible way to circumvent such difficulty is to learn the edge parameters from past observed diffusion cascades, and then maximize the influence based on the learned parameters. The learning task is referred to as *network inference*, and has been extensively studied in the literature (Gomez-Rodriguez, Leskovec, and Krause 2010; Myers and Leskovec 2010; Gomez-Rodriguez, Balduzzi, and Schölkopf 2011; Du et al. 2012; Netrapalli and Sanghavi 2012; Abrahao et al. 2013; Daneshmand et al. 2014; Du et al. 2013, 2014; Narasimhan, Parkes, and Singer 2015; Pouget-Abadie and Horel 2015; He et al. 2016; Chen et al. 2021). However, this approach does not take into account the cost of the learning process and fails to balance between exploration and exploitation when future diffusion cascades come. This motivates the study of *online influence maximization* (OIM) problem considered in this paper.

In OIM, the learner faces an unknown social network and runs T rounds in total. At each round, the learner chooses a seed set to generate cascades, observes the cascade feedback, and receives the influence value as a reward. The goal is to maximize the influence values received over T rounds, or equivalently, to minimize the cumulative regret compared with the optimal seed set that generates the largest influence. The most widely studied feedback in the literature

*The published version of this paper claims there are OIM algorithms with node-level feedback that use standard offline oracles. However, we later found there are severe bugs in our initial algorithms. In this latest version, we retrieve most of our results by presenting an OIM algorithm that uses pair oracles.

†Corresponding author.

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Feedback	Diffusion model	Regret	Pair oracle	Reference
Edge-level	IC	$\tilde{O}(n^3\sqrt{T})$	No	Wang and Chen (2017)
Node-level	LT	$\tilde{O}(n^{9/2}\sqrt{T})$	Yes	Li et al. (2020)
Node-level	IC	$\tilde{O}(n^{7/2}\sqrt{T}/\gamma)$	Yes	Theorem 2

Table 1: Comparison of results on OIM problems

is the *edge-level* feedback (Chen, Wang, and Yuan 2013; Chen et al. 2016; Wang and Chen 2017; Wen et al. 2017; Wu et al. 2019), where the learner can observe whether an edge passes through the information received by its start point. The *node-level* feedback was only investigated very recently in (Vaswani, Lakshmanan, and Schmidt 2016; Li et al. 2020), where the learner can only observe which nodes receive the information at each time step during a diffusion process. In practice, the node-level feedback is more realistic than the edge-level feedback, not only because it reveals less information, but also because it is usually easy to observe who is influenced but very difficult to observe from which edge the influence comes from. For example, in the social network platform, it is easy to learn whether and when a user buys some specific product or service but is difficult to learn based on whose recommendations or comments the user makes such a decision.

In light of this, it is interesting to study the OIM problem with node-level feedback. For the LT model, Li et al. (2020) recently presents a nearly optimal $\tilde{O}(\text{poly}(|G|)\sqrt{T})$ -regret algorithm, at the cost of invoking the so-called offline pair oracles instead of standard oracles. For the IC model, it remains unknown whether the same regret bound can be achieved and this has been an interesting open question in the field.

Our contribution. In this paper, we resolve the aforementioned open question and present the first $\tilde{O}(\text{poly}(|G|)\sqrt{T})$ -regret algorithm for OIM problem under the IC model with node-level feedback. Our algorithm also needs to invoke pair oracles since node-level feedback reveals less information. We compare our result with previous ones in Table 1.

In the technical part, our main contribution is a novel adaptation of the maximum likelihood estimation (MLE) approach which can learn the network parameters and their confidence ellipsoids based on the node-level feedback. We believe this technique is of independent interest and may inspire other results in the field. Besides, we prove the GOM bounded smoothness for the IC model, which is crucial for learning influence functions. The same property is also shown for the LT model in (Li et al. 2020).

Related work. The (offline) influence maximization problem has received great attentions in the past two decades. We refer interested readers to the surveys of (Chen, Lakshmanan, and Castillo 2013; Li et al. 2018) for an overall understanding.

The online influence maximization problem falls into the field of multi-armed bandits (MAB), a prosperous research area that dates back to 1933 (Thompson 1933). In the classical multi-armed *stochastic bandits* (Robbins 1952; Lai and

Robbins 1985), there is a set of n arms, each of which is associated with a reward specified by some *unknown* distribution. At each round t , the learner chooses an arm and receives a reward sampled from the corresponding distribution. The goal is to maximize the total expected rewards received over T rounds. The model was later generalized to the multi-armed *stochastic linear bandits* (Auer, Cesa-Bianchi, and Fischer 2002), where each arm is associated with a characteristic vector and its reward is given by the inner product of the vector and an unknown parameter vector. This model was extensively studied in the literature (Dani, Hayes, and Kakade 2008; Li et al. 2010; Rusmevichientong and Tsitsiklis 2010; Abbasi-Yadkori, Pál, and Szepesvári 2011). Further generalizations include *combinatorial* multi-armed bandits (CMAB) and CMAB with probabilistically triggered arms (CMAB-T) (Chen, Wang, and Yuan 2013; Chen et al. 2016; Wang and Chen 2017), where a subset of arms, called the *super-arm*, can be chosen, and the reward is defined over super-arms and may be non-linear. Besides, the arms beyond the chosen super-arm may also be triggered and observed. CMAB-T is a quite general bandits framework and indeed contains OIM with edge-level feedback as a special case. However, OIM with node-level feedback does not fit into the CMAB-T framework.

OIM has been studied extensively in the literature. For edge-level feedback, existing work (Chen, Wang, and Yuan 2013; Lei et al. 2015; Chen et al. 2016; Wang and Chen 2017; Wen et al. 2017; Wu et al. 2019) present both theoretical and heuristic results. The node-level feedback was first proposed in (Vaswani, Lakshmanan, and Schmidt 2016). However, only heuristic algorithms were presented. Very recently, an $\tilde{O}(\sqrt{T})$ -regret algorithm was presented for the LT model with node-level feedback using pair-oracles in (Li et al. 2020). However, it remains unknown whether the same result holds for the IC model.

2 Preliminaries

2.1 Notations

Given a vector $x \in \mathbb{R}^d$, its transpose is denoted by x^\top . The Euclidean norm of x is denoted by $\|x\|$. For a positive definite matrix $M \in \mathbb{R}^{d \times d}$, the weighted Euclidean norm of x is defined as $\|x\|_M = \sqrt{x^\top M x}$. The minimum eigenvalue of M is denoted by $\lambda_{\min}(M)$, and its determinant and trace are denoted by $\det[M]$ and $\text{tr}[M]$, respectively. For a real-valued function $\mu : \mathbb{R} \rightarrow \mathbb{R}$, its first and second derivatives are denoted by $\dot{\mu}$ and $\ddot{\mu}$, respectively.

2.2 Social Network

A social network is a weighted directed graph $G = (V, E)$ with a node set V of $n = |V|$ nodes and an edge set E

of $m = |E|$ edges. Each edge $e \in E$ is associated with a weight or probability $p(e) \in [0, 1]$. The edge probability vector is then denoted by $p = (p(e))_{e \in E}$, which describes the graph completely. For a node $v \in V$, let $N(v) = N^{in}(v)$ be the set of in-neighbors of v and $d_v = |N(v)|$ be its in-degree. The maximum in-degree of the graph is denoted by $D = \max_{v \in V} d_v$. In this paper, we use E_v to denote the set of incoming edges of v and $p_v = (p(e))_{e \in E_v} \in [0, 1]^{d_v}$ to denote the probability vector corresponding to these edges. The e -th entry of p_v is denoted by $p_v(e)$. Thus, $p(e)$ and $p_v(e)$ refers to the same edge probability and we will use them interchangeably throughout the paper. For an edge $e = (u, v) \in E_v$, we use e_{uv} to explicitly indicate e 's endpoints. Let $\chi(e_{uv}) \in \{0, 1\}^{d_v}$ be the characteristic vector of e_{uv} over E_v such that all entries of $\chi(e_{uv})$ are 0 except that the entry corresponding to e_{uv} is 1. The characteristic vector of a subset $E' \subseteq E_v$ is then defined as $\chi(E') := \sum_{e \in E'} \chi(e) \in \{0, 1\}^{d_v}$. For simplicity, we define $x_e := \chi(e)$.

2.3 Offline Influence Maximization

The input of the offline problem is a social network, over which the information spreads. A node $v \in V$ is called *active* if it receives the information and *inactive* otherwise. We first describe the independent cascade (IC) diffusion model.

In the IC model, the diffusion proceeds in discrete time steps $\tau = 0, 1, 2, \dots$. At the beginning of the diffusion ($\tau = 0$), there is an initially active set S_0 of nodes called *seeds*. For $\tau \geq 1$, the active node set S_τ after time τ is generated as follows. First, let $S_\tau = S_{\tau-1}$. Next, for each $v \in V \setminus S_{\tau-1}$, every node $u \in N(v) \cap (S_{\tau-1} \setminus S_{\tau-2})$ will try to *activate* v *independently* with probability $p(e_{uv})$ (let $S_{-1} = \emptyset$). Hence, v will be activated with probability $1 - \prod_{u \in N(v) \cap (S_{\tau-1} \setminus S_{\tau-2})} (1 - p(e_{uv}))$ and be added into S_τ once being activated. The diffusion terminates if $S_\tau = S_{\tau-1}$ for some τ and therefore it proceeds in at most n time steps. Let $(S_0, S_1, \dots, S_{n-1})$ be the sequence of the active node sets during the diffusion process, where S_τ denotes the active node set after time τ .

Given a seed set S_0 , the *influence spread* of S_0 is defined as $\sigma(S_0) = \mathbf{E}[|S_{n-1}|]$, i.e. the expected number of active nodes by the end of the diffusion. Here, $\sigma : 2^V \rightarrow \mathbb{R}_+$ is called the *influence spread function*. In this paper, we also use $\sigma(S, p)$ to state the edge probability vector p explicitly. The influence maximization (IM) problem takes as input the social network G and an integer $K \in \mathbb{N}_+$, and requires to find the seed set S^{opt} that gives the maximum influence spread with at most K seeds, i.e. $S^{\text{opt}} \in \arg\max_{S \subseteq V, |S| \leq K} \sigma(S)$. It is well-known that the IM problem admits a $(1 - 1/e - \epsilon)$ approximation under the IC model (Kempe, Kleinberg, and Tardos 2003), which is tight assuming $P \neq NP$ (Feige 1998).

2.4 Online Influence Maximization

In the online influence maximization problem (OIM) considered in this paper, there is an underlying social network $G = (V, E)$, whose edge parameter vector p^* is unknown initially. At each round t of total T rounds, the learner chooses a seed set S_t with cardinality at most K , observes

the cascade feedback, and updates her knowledge about the parameter p^* for later selections. The feedback considered in this paper is node-level feedback, which means that the learner observes a realization of the sequence of active nodes $(S_{t,0}, S_{t,1}, \dots, S_{t,n-1})$ after selecting $S_{t,0} = S_t$.

In order to solve OIM problem, oracles for offline IM problem are often invoked. Such an oracle takes as input the edge probability vector and outputs a good approximate solution for IM problem. However, when node-level feedback is used, both Li et al. (2020) and this paper can only guarantee that the true edge probability vector falls into a confidence region. Thus, we need the so-called *pair oracle* which takes as input the confidence region and can still find a good solution.

Formally, denote by p^* the true edge probability vector and by $\mathcal{C} \in \mathbb{R}^m$ a confidence region satisfying $p^* \in \mathcal{C}$. Let ORACLE be a pair-oracle which solves the problem $\max_{S: |S| \leq K, p \in \mathcal{C}} \sigma(S, p)$ and $(\tilde{S}, \tilde{p}) = \text{ORACLE}(G, K, \mathcal{C})$ be its output. Define $S^{\text{opt}} \in \arg\max_{S: |S| \leq K} \sigma(S, p^*)$ to be the optimal seed set. For $\alpha, \beta \in [0, 1]$, we say ORACLE is an (α, β) -pair-oracle if $\Pr[\sigma(\tilde{S}, \tilde{p}) \geq \alpha \cdot \sigma(S^{\text{opt}}, p^*)] \geq \beta$, where the probability is taking from the possible randomness of ORACLE. Note that ORACLE is hard to implement, but this paper mainly focuses on the effectiveness of OIM algorithms and the efficiency is not our concern.

Equipped with an (α, β) -pair-oracle, the objective of OIM is to minimize the cumulative $(\alpha\beta)$ -scaled regret over T rounds:

$$\begin{aligned} R(T) &= \mathbf{E} \left[\sum_{t=1}^T R_t \right] \\ &= \mathbf{E} \left[T\alpha\beta \cdot \sigma(S^{\text{opt}}, p^*) - \sum_{t=1}^T |S_{t,n-1}| \right]. \end{aligned}$$

Due to the additivity of expectation, it is equal to

$$R(T) = \mathbf{E} \left[T\alpha\beta \cdot \sigma(S^{\text{opt}}, p^*) - \sum_{t=1}^T \sigma(S_t, p^*) \right].$$

3 OIM Algorithm under the IC Model

In this section, we present an algorithm for OIM under the IC model with node-level feedback (Algorithm 1). Our algorithm adopts the canonical *upper confidence bound* (UCB) framework in the bandits problem. Under the UCB framework, at each round t , we first compute an estimate \hat{p}_{t-1} of p^* and a corresponding confidence region \mathcal{C}_{t-1} based on the feedback before round t . Then, a seed set S_t is selected by invoking an (α, β) -pair-oracle to obtain (S_t, \tilde{p}_{t-1}) , which satisfies that $\tilde{p}_{t-1} \in \mathcal{C}_{t-1}$ and $|S_t| \leq K$.

For OIM with node-level feedback, the key difficulty of applying the UCB framework lies in how to use the node-level feedback collected in the previous rounds to update the estimate of p^* . For each node $v \in V$, Algorithm 1 will estimate the probability vector $p_v^* \in [0, 1]^{d_v}$ of the incoming edges of v separately. Note that all p_v^* together form p^* .

We first explain how to extract information on p_v^* from the feedback $(S_{t,0}, S_{t,1}, \dots, S_{t,n-1})$ at round t . When $t > T_0$,

Algorithm 1: IC-UCB

Input: Graph $G = (V, E)$, seed set cardinality $K \in \mathbb{N}$, (α, β) -pair-oracle ORACLE, parameter $\gamma \in (0, 1)$ in Assumption 1.

- 1: Initialize $M_{0,v} \leftarrow \mathbf{0} \in \mathbb{R}^{d_v \times d_v}$ for all $v \in V$, $\delta \leftarrow 1/(3n\sqrt{T})$, $R \leftarrow \left\lceil \frac{512D}{\gamma^4} (D^2 + \ln(1/\delta)) \right\rceil$, $T_0 \leftarrow nR$ and $\rho \leftarrow \frac{3}{\gamma} \sqrt{\ln(1/\delta)}$.
 - 2: **for all** $u \in V$ **do**
 - 3: Choose $\{u\}$ as the seed set for the next R rounds and construct data pairs from observations (see the text in this section for details).
 - 4: **end for**
 - 5: **for** $t = T_0 + 1, T_0 + 2, \dots, T$ **do**
 - 6: $\{\hat{\theta}_{t-1,v}, \mathcal{C}'_{t-1,v}\}_{v \in V} =$
 Estimate($(S_{k,0}, S_{k,1}, \dots, S_{k,n-1})_{1 \leq k \leq t-1}$) (see Algorithm 2).
 - 7: Let $\mathcal{C}_{t-1,v} = \{p_v \in [0, 1]^{d_v} \mid \theta_v \in \mathcal{C}'_{t-1,v}\}$ and $\mathcal{C}_{t-1} = \{\mathcal{C}_{t-1,v}\}_{v \in V}$
 - 8: Choose $(S_t, \hat{p}_t) \in \text{ORACLE}(G, K, \mathcal{C}_{t-1})$ and observe node-level feedback $(S_{t,0}, S_{t,1}, \dots, S_{t,n-1})$.
 - 9: **end for**
-

the data is processed in a more economical way. Assume that node v remains inactive after time τ and some of its neighbors $(S_{t,\tau} \setminus S_{t,\tau-1}) \cap N(v) \neq \emptyset$ was newly activated in time τ . Then, these neighbors will try to activate node v in time $\tau + 1$. Let

$$E' := \{e_{uv} \in E_v \mid u \in (S_{t,\tau} \setminus S_{t,\tau-1}) \cap N(v)\}$$

be the set of edges which point from these neighbors to node v . By the diffusion rule of the IC model, the probability that v is activated by them in time $\tau + 1$ is

$$1 - \prod_{e \in E'} (1 - p^*(e)).$$

If node v did become active in time $\tau + 1$, then we use data pair $(\chi(E'), 1)$ to record this event. Otherwise, we use data pair $(\chi(E'), 0)$ to record the event that v remained inactive in time $\tau + 1$. By inspecting each time step of the diffusion till v became active or no new neighbors of v were activated, we are able to construct $J_{t,v}$ data pairs accordingly, denoted by $(X_{t,j,v}, Y_{t,j,v})$, $1 \leq j \leq J_{t,v}$. Here, $J_{t,v} \leq d_v$, since v has d_v neighbors and a new data pair is constructed only when some inactive neighbors of v become active. $X_{t,j,v} \in \{0, 1\}^{d_v}$ indicates the characteristic vector of the edges corresponding to the j -th batch of neighbors that were activated. $Y_{t,j,v} \in \{0, 1\}$ indicates if v was activated by these neighbors. It is easy to see when $j < J_{t,v}$, $Y_{t,j,v} = 0$, and only when $j = J_{t,v}$, it is possible for $Y_{t,j,v}$ to be 1. This is because v will remain active once it is activated. Though some newly active neighbors of v will still try to “activate” v thereafter, it is impossible to observe whether the attempt succeeds.

For the initial regularization phase (line 2 to line 4) where $t \leq T_0$, the process of extracting information is wasteful in that only the first-step activation is taken into account. In

Algorithm 2: Estimate. Note that the code is written as a computation from scratch in each round to accommodate the initialization period of Algorithm 1, and it can be easily adapted to the incremental computation form.

Input: All observations $(S_{k,0}, S_{k,1}, \dots, S_{k,n-1})_{1 \leq k \leq t}$ until round t .

- 1: **for all** $v \in V$ **do**
 - 2: Construct data pairs $(X_{k,j,v}, Y_{k,j,v})_{1 \leq k \leq t, 1 \leq j \leq J_{k,v}}$ from observations $(S_{k,0}, S_{k,1}, \dots, S_{k,n-1})_{1 \leq k \leq t}$ (see the text in this section for details).
 - 3: $L_{t,v}(\theta_v) \leftarrow \sum_{k=1}^t \sum_{j=1}^{J_{k,v}} [-\exp(-X_{k,j,v}^\top \theta_v) - (1 - Y_{k,j,v}) X_{k,j,v}^\top \theta_v]$, see eq. (5).
 - 4: $\hat{\theta}_{t,v} \leftarrow \operatorname{argmax}_{\theta_v} L_{t,v}(\theta_v)$.
 - 5: $M_{t,v} \leftarrow \sum_{k=1}^t \sum_{j=1}^{J_{k,v}} X_{k,j,v} X_{k,j,v}^\top$, see eq. (1).
 - 6: $\mathcal{C}'_{t,v} \leftarrow \{\theta \in [0, 1]^{d_v} \mid \|\theta - \hat{\theta}_{t,v}\|_{M_{t,v}} \leq \rho\}$.
 - 7: **end for**
-

this part, the algorithm chooses each node $u \in V$ as the seed set for R rounds, and then observes the activation of u 's all out-neighbors so as to gather information about its outgoing edges. More formally, let node u be chosen as the seed in round t . In the case $u \in N(v)$, $J_{t,v} = 1$ and we construct data pair $(\chi(e_{uv}), 1)$ if $v \in S_{t,1}$, or data pair $(\chi(e_{uv}), 0)$ if $v \notin S_{t,1}$. In the case $u \notin N(v)$, no data pair is constructed. By the regularization step, each edge will be observed exactly R times. Intuitively, this step leads to a coarse estimate of each individual probability $p(e)$ for $e \in E$. Technically, this step guarantees a lower bound of the minimum eigenvalue of the Gram matrix $M_{t,v}$ defined in eq. (1), which ensures the correctness of condition (6) in Theorem 1 in the analysis.

For each node $v \in V$, we can use all the feedback data $(S_{k,0}, S_{k,1}, \dots, S_{k,n-1})_{1 \leq k \leq t}$ in the first t rounds to construct data pairs $\{(X_{k,j,v}, Y_{k,j,v})\}_{1 \leq k \leq t, 1 \leq j \leq J_{k,v}}$. The Gram matrix $M_{t,v}$ of these data pairs is defined as

$$M_{t,v} := \sum_{k=1}^t \sum_{j=1}^{J_{k,v}} X_{k,j,v} X_{k,j,v}^\top \quad (1)$$

Now that we have explained how to extract information on p_v^* by constructing new data pairs from the node-level feedback, we next introduce how to use these data pairs to estimate p_v^* . Inspired by the network inference problem (Netrapalli and Sanghavi 2012; Narasimhan, Parkes, and Singer 2015; Pouget-Abadie and Horel 2015), we use the maximum likelihood estimation (MLE) to estimate p_v^* . Algorithm 2 provide a detailed estimation procedure, which has two important features described below.

Transformation of edge parameter p into parameter θ .

By the diffusion rule of the IC model, for each $v \in V$, given $X \in \{0, 1\}^{d_v}$, let $Y \in \{0, 1\}$ indicates whether v is activated in *one time step*. Then,

$$\mathbb{E}[Y \mid X] = 1 - \Pi_{e: X(e)=1} (1 - p(e)),$$

which a complex function of parameter $p(e)$. We therefore consider a transformation of edge probability vector p into a

new vector θ where

$$\theta(e) = -\ln(1 - p(e)) \text{ for each } e \in E. \quad (2)$$

Then,

$$p(e) = 1 - \exp(-\theta(e)) \text{ for each } e \in E, \quad (3)$$

and

$$\mathbf{E}[Y \mid X] = \mu(X^\top \theta_v),$$

where the *link function* $\mu : \mathbb{R} \rightarrow \mathbb{R}$ is defined as

$$\mu(x) := 1 - \exp(-x).$$

This indeed forms an instance of the *generalized linear bandit* (GLB) problem studied in (Filippi et al. 2010; Li, Lu, and Zhou 2017). They also use MLE to solve the GLB problem. Hence, we will analyze the regret of Algorithm 1 via their methods.

Pseudo log-likelihood function $L_{t,v}$.

During the update of the estimate of p^* (or θ^*), a standard log-likelihood function is often used:

$$\begin{aligned} \mathcal{L}_{t,v}^{std}(\theta_v) &= \sum_{k=1}^t \sum_{j=1}^{J_{k,v}} [Y_{k,j,v} \ln \mu(X_{k,j,v}^\top \theta_v) \\ &\quad + (1 - Y_{k,j,v}) \ln(1 - \mu(X_{k,j,v}^\top \theta_v))]. \end{aligned}$$

However, the analysis in (Filippi et al. 2010; Li, Lu, and Zhou 2017) requires that the gradient of the log-likelihood function has the form

$$\sum_{k=1}^t \sum_{j=1}^{J_{k,v}} [Y_{k,j,v} - \mu(X_{k,j,v}^\top \theta_v)] X_{k,j,v}. \quad (4)$$

Such requirement is met in Filippi et al. (2010); Li, Lu, and Zhou (2017) by assuming the distribution of Y conditioned on X falls into some sub-class of the exponential family of distributions, which is however not satisfied in our case. In this paper, we present an alternative way to overcome such technical difficulty. That is, we “integrate” the gradient in eq. (4) to obtain a *pseudo log-likelihood function* $L_{t,v}$:

$$\begin{aligned} L_{t,v}(\theta_v) &= \sum_{k=1}^t \sum_{j=1}^{J_{k,v}} [-\exp(-X_{k,j,v}^\top \theta_v) - (1 - Y_{k,j,v}) X_{k,j,v}^\top \theta_v]. \end{aligned} \quad (5)$$

This ensures that the gradient of $L_{t,v}$ has the form of eq. (4) and therefore the analysis of Filippi et al. (2010); Li, Lu, and Zhou (2017) can be used. Such an approach is of great independent interest and we leave it as an open problem to find a more intuitive explanation for it.

4 Regret Analysis

We now give an analysis of the regret of Algorithm 1. First, we need to show that for each $v \in V$, the estimate $\hat{\theta}_{t,v}$ is close to the true parameter θ_v^* . To ensure this, we require Assumption 1 below.

Assumption 1. *There exists a parameter $\gamma \in (0, 1)$ such that $\prod_{u \in N(v)} (1 - p^*(e_{uv})) \geq \gamma$ for all $v \in V$.*

Similar or even stronger assumptions are adopted in all previous approaches for network inference (Netrapalli and Sanghavi 2012; Narasimhan, Parkes, and Singer 2015; Pouget-Abadie and Horel 2015; Chen et al. 2021). Assumption 1 means that node $v \in V$ will remain inactive with probability at least γ even if all of its in-neighbors are simultaneously activated. It reflects the stubbornness of the agent (node). That is, the behavior of a node is partially determined by its intrinsic motivation, not by its neighbors. So, even when all its neighbors adopt a new behavior, there is a nontrivial probability that the node will still not adopt the new behavior.

Under Assumption 1, it is possible to show that $\hat{\theta}_{t,v}$ and θ_v^* are close to each other in all directions, from which we can obtain a confidence region for θ_v^* , as Theorem 1 states. The proof of Theorem 1 is similar to that of Theorem 1 in (Li, Lu, and Zhou 2017). For completeness, we include the proof in Appendix A.

Theorem 1. *Suppose that Assumption 1 holds. For each $v \in V$, $\hat{\theta}_{t,v}$ and $M_{t,v}$ are computed according to Algorithm 2. Given $\delta \in (0, 1)$, if*

$$\lambda_{\min}(M_{t,v}) \geq \frac{512d_v}{\gamma^4} \left(d_v^2 + \ln \frac{1}{\delta} \right). \quad (6)$$

Then, with probability at least $1 - 3\delta$, for any $x \in \mathbb{R}^{d_v}$, we have

$$|x^\top (\hat{\theta}_{t,v} - \theta_v^*)| \leq \frac{3}{\gamma} \sqrt{\ln(1/\delta)} \cdot \|x\|_{M_{t,v}^{-1}}.$$

Thus, by setting $x = M_{t,v}^\top (\hat{\theta}_{t,v} - \theta_v^)$, we obtain*

$$\|\hat{\theta}_{t,v} - \theta_v^*\|_{M_{t,v}} \leq \frac{3}{\gamma} \sqrt{\ln(1/\delta)}.$$

After we prove that $\hat{\theta}_{t,v}$ and θ_v^* are indeed close to each other, we need to show that the influence functions $\sigma(S, \hat{p})$ and $\sigma(S, p^*)$ induced by the corresponding probability vectors \hat{p}_t and p^* are also close. To this end, we prove the *group observation modulated* (GOM) bounded smoothness condition for the IC model. The condition is inspired by the GOM condition for the LT model (Li et al. 2020). We remark that for edge-level feedback, there is a related *triggering probability modulated* (TPM) bounded smoothness condition (Wang and Chen 2017; Wen et al. 2017). However, the TPM condition does not suffice for node-level feedback.

We now state the GOM condition formally. Given a seed set $S \subseteq V$ and a node $v \in V \setminus S$, we say node $u \in V \setminus S$ is *relevant* to node v if there is a path P from S to v such that $u \in P$. Let $V[S, v] \subseteq V$ be the set of nodes relevant to v given seed set S . Given diffusion cascade $(S_0 = S, S_1, \dots, S_{n-1})$, construct data pairs $\{(X_{j,v}, Y_{j,v})\}_{1 \leq j \leq J_v}$ according to the economical way described previously (not the wasteful way for the regularization phase). We have the following GOM condition for the IC model, whose proof is presented in Appendix B.

Lemma 1 (GOM bounded smoothness for the IC model). *Fix any seed set $S \subseteq V$. For any two edge-probability vectors $\tilde{p}, p^* \in [0, 1]^{|E|}$, let $\tilde{\theta}, \theta^*$ be the vectors defined as eq. (2). Then,*

$$\begin{aligned} & |\sigma(S, \tilde{p}) - \sigma(S, p^*)| \\ & \leq \sum_{v \in V \setminus S} \sum_{u \in V[S, v]} \mathbf{E} \left[\sum_{j=1}^{J_u} \left| X_{j,u}^\top (\tilde{\theta}_u - \theta_u^*) \right| \right], \end{aligned}$$

where the expectation is taken over the randomness of the diffusion cascade $(S_0, S_1, \dots, S_{n-1})$, which is generated with respect to parameter p^* .

Equipped with the aforementioned tools, we now set about presenting the analysis of Algorithm 1. Given a seed set $S \subseteq V$ and a node $u \in V \setminus S$, define

$$n_{S,u} := \sum_{v \in V \setminus S} \mathbf{1}\{u \in V[S, v]\}$$

to be the number of nodes that u is relevant to. Further, define

$$\zeta(G) := \max_{S: |S| \leq K} \sqrt{\sum_{u \in V} n_{S,u}^2} \leq O(n^{3/2}).$$

We present the regret of Algorithm 1 in Theorem 2.

Theorem 2. *When we use an (α, β) -pair-oracle in Algorithm 1, under Assumption 1, the $\alpha\beta$ -scaled regret of Algorithm 1 satisfies that*

$$R(T) = \tilde{O} \left(\frac{\zeta(G) D \sqrt{mT}}{\gamma} \right) = \tilde{O} \left(\frac{n^{7/2} \sqrt{T}}{\gamma} \right).$$

Proof. Let \mathcal{H}_t be the history of past rounds by the end of round t . For $t \leq T_0$, $\mathbf{E}[R_t] \leq n$, since there are n nodes in G . Now consider the case where $t > T_0$. By the definition of R_t ,

$$\mathbf{E}[R_t \mid \mathcal{H}_{t-1}] = \mathbf{E}[\alpha\beta \cdot \sigma(S^{\text{opt}}, p^*) - \sigma(S_t, p^*) \mid \mathcal{H}_{t-1}],$$

where the expectation is taken over the randomness of S_t .

For any $T_0 < t \leq T$ and $v \in V$, define event $\xi_{t-1,v}$ as

$$\xi_{t-1,v} := \{\|\hat{\theta}_{t-1,v} - \theta_v^*\|_{M_{t-1,v}} \leq \rho\},$$

and let $\bar{\xi}_{t-1,v}$ be its complement. By the choices of δ, R, T_0, ρ as in Algorithm 1, the fact that $\lambda_{\min}(M_{t-1,v}) \geq \lambda_{\min}(M_{T_0,v}) = R$ and Theorem 1, we have $\Pr[\bar{\xi}_{t-1,v}] \leq 3\delta$. Further define event $\xi_{t-1} := \bigwedge_{v \in V} \xi_{t-1,v}$ and let $\bar{\xi}_{t-1}$ be its complement. By union bound, $\Pr[\bar{\xi}_{t-1}] \leq 3\delta n$. Note that under event ξ_{t-1} , for all $v \in V$, $\theta_v^* \in \mathcal{C}'_{t-1,v}$. Hence, $p_v^* \in \mathcal{C}_{t-1,v}$ and $p^* \in \mathcal{C}_{t-1}$. Since (S_t, \tilde{p}_t) is obtained by invoking an (α, β) -pair-oracle ORACLE over \mathcal{C}_{t-1} , we have

$$\begin{aligned} \mathbf{E}[R_t] & \leq \Pr[\xi_{t-1}] \cdot \mathbf{E}[\alpha\beta \cdot \sigma(S^{\text{opt}}, p^*) - \sigma(S_t, p^*) \mid \xi_{t-1}] \\ & \quad + \Pr[\bar{\xi}_{t-1}] \cdot n \\ & \leq \mathbf{E}[\sigma(S_t, \tilde{p}_t) - \sigma(S_t, p^*) \mid \xi_{t-1}] + 3\delta n^2. \end{aligned}$$

Next, by the GOM bounded smoothness for the IC model in Lemma 1, we obtain that

$$\begin{aligned} & \mathbf{E}[R_t] - 3\delta n^2 \\ & \leq \mathbf{E} \left[\sum_{v \in V \setminus S_t} \sum_{u \in V[S_t, v]} \sum_{j=1}^{J_{t,u}} \left| X_{t,j,u}^\top (\tilde{\theta}_{t,u} - \theta_u^*) \right| \mid \xi_{t-1} \right]. \end{aligned}$$

By the Cauchy-Schwarz inequality, we have

$$|X_{t,j,u}^\top (\tilde{\theta}_{t,u} - \theta_u^*)| \leq \|X_{t,j,u}\|_{M_{t-1,u}^{-1}} \|\tilde{\theta}_{t,u} - \theta_u^*\|_{M_{t-1,u}}.$$

Besides, under event ξ_{t-1} , $\hat{\theta}_{t,u}, \theta_u^* \in \mathcal{C}'_{t-1,u}$ for all $u \in V$. Then, by the triangle inequality,

$$\begin{aligned} & \|\tilde{\theta}_{t,u} - \theta_u^*\|_{M_{t-1,u}} \\ & \leq \|\tilde{\theta}_{t,u} - \hat{\theta}_{t-1,u}\|_{M_{t-1,u}} + \|\hat{\theta}_{t-1,u} - \theta_u^*\|_{M_{t-1,u}} \leq 2\rho. \end{aligned}$$

Combining the above inequalities, we obtain that

$$\begin{aligned} & \mathbf{E}[R_t] - 3\delta n^2 \\ & \leq 2\rho \cdot \mathbf{E} \left[\sum_{v \in V \setminus S_t} \sum_{u \in V[S_t, v]} \sum_{j=1}^{J_{t,u}} \|X_{t,j,u}\|_{M_{t-1,u}^{-1}} \right] \\ & = 2\rho \cdot \mathbf{E} \left[\sum_{u \in V \setminus S_t} \sum_{j=1}^{J_{t,u}} \|X_{t,j,u}\|_{M_{t-1,u}^{-1}} \sum_{v \in V \setminus S_t} \mathbf{1}_{u \in V[S_t, v]} \right] \\ & = 2\rho \cdot \mathbf{E} \left[\sum_{u \in V \setminus S_t} n_{S_t, u} \sum_{j=1}^{J_{t,u}} \|X_{t,j,u}\|_{M_{t-1,u}^{-1}} \right]. \end{aligned}$$

Recall that the above derivation holds for $t > T_0$, and for $t \leq T_0$, $\mathbf{E}[R_t] \leq n$. We thus have

$$\begin{aligned} R(T) & \leq 2\rho \cdot \mathbf{E} \left[\sum_{t=T_0+1}^T \sum_{v \in V \setminus S_t} n_{S_t, v} \sum_{j=1}^{J_{t,v}} \|X_{t,j,v}\|_{M_{t-1,v}^{-1}} \right] \\ & \quad + 3\delta n^2(T - T_0) + nT_0. \end{aligned}$$

To further simplify the above inequality, we prove the following lemma, whose proof is presented in Appendix C.

Lemma 2. *For any $v \in V$,*

$$\begin{aligned} & \sum_{t=T_0+1}^T \sum_{v \in V \setminus S_t} n_{S_t, v} \sum_{j=1}^{J_{t,v}} \|X_{t,j,v}\|_{M_{t-1,v}^{-1}} \\ & \leq \zeta(G) D \sqrt{(m+n)(T - T_0) \ln(R + (T - T_0)D)}. \end{aligned}$$

By Lemma 2, we have

$$\begin{aligned}
& R(T) \\
& \leq 2\rho \cdot \mathbf{E} \left[\sum_{t=T_0+1}^T \sum_{v \in V \setminus S_t} n_{S_t, v} \sum_{j=1}^{J_{t,v}} \|X_{t,j,v}\|_{M_{t-1,v}^{-1}} \right] \\
& + 3\delta n^2(T - T_0) + nT_0 \\
& \leq 2\rho\zeta(G)D\sqrt{(m+n)(T-T_0)\ln(R+(T-T_0)D)} \\
& + 3\delta n^2(T - T_0) + nT_0 \\
& \leq \frac{6\zeta(G)D}{\gamma} \sqrt{(m+n)T\ln(TD)\ln(3nT)} + n\sqrt{T} \\
& + \frac{512Dn^2}{\gamma^4} (D^2 + \ln(3nT)) + 1 \\
& = \tilde{O} \left(\frac{\zeta(G)D\sqrt{mT}}{\gamma} \right).
\end{aligned}$$

The last inequality is obtained by plugging $\delta = 1/(3n\sqrt{T})$, $R = \left\lceil \frac{512D}{\gamma^4} (D^2 + \ln(1/\delta)) \right\rceil$, $T_0 = nR$ and $\rho = \frac{3}{\gamma} \sqrt{\ln(1/\delta)}$ into the formula. \square

We remark that the worst-case regret for the IC model with edge-level feedback is $\tilde{O}(n^3\sqrt{T})$ in (Wang and Chen 2017). Thus, our regret bound under node-level feedback matches the previous ones under edge-level feedback in the worst case, up to a $n^{1/2}/\gamma$ factor.

To get an intuition about γ 's value, assume that each edge probability $\leq 1 - c$ for some constant $c \in (0, 1)$. Then, $\gamma = O(c^D)$, where D is the maximum in-degree of the graph. Thus, in the worst case, $1/\gamma$ is exponential in n . But when $D = O(\log n)$, $1/\gamma$ is polynomial in n and so is the regret bound. We think $D = O(\log n)$ is reasonable in practice, since a person only has a limited attention and cannot pay attention to too many people in the network.

5 Conclusion

In this paper, we investigate the OIM problem under the IC model with node-level feedback. We presents an $\tilde{O}(\sqrt{T})$ -regret OIM algorithm for the problem, which almost matches the optimal regret bound as well as the state-of-the-art regret bound with edge-level feedback. Our novel adaptation of MLE to fit the GLB model is of great independent interest, which might be combined with the GLB model to handle rewards generated from a broader classes of distributions.

There still remain several open problems in the node-level feedback setting. An immediate one is to either remove Assumption 1 for the edge probability vector or at least the assumption parameter from the regret bound. Besides, one can also study if there exists optimal-regret OIM algorithms with node-level feedback that use standard offline oracles. Finally, it is interesting to develop a general bandit framework which includes OIM with node-level feedback as a special case, just like CMAB-T containing OIM with edge-level feedback.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China Grants No. 61832003, 61872334, the Strategic Priority Research Program of Chinese Academy of Sciences under Grant No. XDA27000000.

References

- Abbasi-Yadkori, Y.; Pál, D.; and Szepesvári, C. 2011. Improved Algorithms for Linear Stochastic Bandits. In *Advances in Neural Information Processing Systems 24: 25th Annual Conference on Neural Information Processing Systems 2011.*, 2312–2320.
- Abrahamo, B. D.; Chierichetti, F.; Kleinberg, R.; and Panconesi, A. 2013. Trace complexity of network inference. In *the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2013*, 491–499. ACM.
- Auer, P.; Cesa-Bianchi, N.; and Fischer, P. 2002. Finite-time Analysis of the Multiarmed Bandit Problem. *Mach. Learn.*, 47(2-3): 235–256.
- Borgs, C.; Brautbar, M.; Chayes, J. T.; and Lucier, B. 2014. Maximizing Social Influence in Nearly Optimal Time. In Chekuri, C., ed., *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014*, 946–957. SIAM.
- Chen, K.; Hu, I.; and Ying, Z. 1999. Strong consistency of maximum quasi-likelihood estimators in generalized linear models with fixed and adaptive designs. *The Annals of Statistics*, 27(4): 1155 – 1163.
- Chen, W.; Lakshmanan, L. V. S.; and Castillo, C. 2013. *Information and Influence Propagation in Social Networks*. Synthesis Lectures on Data Management. Morgan & Claypool Publishers.
- Chen, W.; Sun, X.; Zhang, J.; and Zhang, Z. 2021. Network Inference and Influence Maximization from Samples. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021*, 1707–1716. PMLR.
- Chen, W.; Wang, Y.; and Yuan, Y. 2013. Combinatorial Multi-Armed Bandit: General Framework and Applications. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013*, 151–159. JMLR.
- Chen, W.; Wang, Y.; Yuan, Y.; and Wang, Q. 2016. Combinatorial Multi-Armed Bandit and Its Extension to Probabilistically Triggered Arms. *J. Mach. Learn. Res.*, 17: 50:1–50:33.
- Daneshmand, H.; Gomez-Rodriguez, M.; Song, L.; and Schölkopf, B. 2014. Estimating Diffusion Network Structures: Recovery Conditions, Sample Complexity & Soft-thresholding Algorithm. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2014*, 793–801. JMLR.
- Dani, V.; Hayes, T. P.; and Kakade, S. M. 2008. Stochastic Linear Optimization under Bandit Feedback. In *the 21st Annual Conference on Learning Theory, COLT 2008*, 355–366. Omnipress.

- Du, N.; Liang, Y.; Balcan, M.; and Song, L. 2014. Influence Function Learning in Information Diffusion Networks. In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014*, 2016–2024. JMLR.
- Du, N.; Song, L.; Gomez-Rodriguez, M.; and Zha, H. 2013. Scalable Influence Estimation in Continuous-Time Diffusion Networks. In *Advances in Neural Information Processing Systems 26 (NIPS 2013)*, 3147–3155.
- Du, N.; Song, L.; Smola, A. J.; and Yuan, M. 2012. Learning Networks of Heterogeneous Influence. In *Advances in Neural Information Processing Systems 25 (NIPS 2012)*, 2789–2797.
- Feige, U. 1998. A Threshold of $\ln n$ for Approximating Set Cover. *J. ACM*, 45(4): 634–652.
- Filippi, S.; Cappé, O.; Garivier, A.; and Szepesvári, C. 2010. Parametric Bandits: The Generalized Linear Case. In *Advances in Neural Information Processing Systems 23 (NIPS 2010)*, 586–594. Curran Associates, Inc.
- Gomez-Rodriguez, M.; Balduzzi, D.; and Schölkopf, B. 2011. Uncovering the Temporal Dynamics of Diffusion Networks. In *Proceedings of the 28th International Conference on Machine Learning, ICML 2011*, 561–568. Omnipress.
- Gomez-Rodriguez, M.; Leskovec, J.; and Krause, A. 2010. Inferring networks of diffusion and influence. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2010*, 1019–1028. ACM.
- He, X.; Xu, K.; Kempe, D.; and Liu, Y. 2016. Learning Influence Functions from Incomplete Observations. In *Advances in Neural Information Processing Systems 29 (NIPS 2016)*, 2065–2073.
- Kempe, D.; Kleinberg, J. M.; and Tardos, É. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2003*, 137–146. ACM.
- Lai, T. L.; and Robbins, H. 1985. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1): 4–22.
- Lei, S.; Maniu, S.; Mo, L.; Cheng, R.; and Senellart, P. 2015. Online Influence Maximization. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2015*, 645–654. ACM.
- Li, L.; Chu, W.; Langford, J.; and Schapire, R. E. 2010. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th International Conference on World Wide Web, WWW 2010*, 661–670. ACM.
- Li, L.; Lu, Y.; and Zhou, D. 2017. Provably Optimal Algorithms for Generalized Linear Contextual Bandits. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017*, 2071–2080. PMLR.
- Li, S.; Kong, F.; Tang, K.; Li, Q.; and Chen, W. 2020. Online Influence Maximization under Linear Threshold Model. In *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*.
- Li, Y.; Fan, J.; Wang, Y.; and Tan, K.-L. 2018. Influence maximization on social graphs: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 30(10): 1852–1872.
- Myers, S. A.; and Leskovec, J. 2010. On the Convexity of Latent Social Network Inference. In *Advances in Neural Information Processing Systems 23 (NIPS 2010)*, 1741–1749. Curran Associates, Inc.
- Narasimhan, H.; Parkes, D. C.; and Singer, Y. 2015. Learnability of Influence in Networks. In *Advances in Neural Information Processing Systems 28 (NIPS 2015)*, 3186–3194.
- Netrapalli, P.; and Sanghavi, S. 2012. Learning the graph of epidemic cascades. In *ACM SIGMETRICS/PERFORMANCE Joint International Conference on Measurement and Modeling of Computer Systems, SIGMETRICS 2012*, 211–222. ACM.
- Pollard, D. 1990. Empirical Processes: Theory and Applications. *NSF-CBMS Regional Conference Series in Probability and Statistics*, 2: i–86.
- Pouget-Abadie, J.; and Horel, T. 2015. Inferring Graphs from Cascades: A Sparse Recovery Framework. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015*, 977–986. JMLR.
- Robbins, H. 1952. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58(5): 527–535.
- Rusmevichientong, P.; and Tsitsiklis, J. N. 2010. Linearly Parameterized Bandits. *Math. Oper. Res.*, 35(2): 395–411.
- Tang, Y.; Shi, Y.; and Xiao, X. 2015. Influence Maximization in Near-Linear Time: A Martingale Approach. In *International Conference on Management of Data, SIGMOD 2014*, 1539–1554. ACM.
- Tang, Y.; Xiao, X.; and Shi, Y. 2014. Influence maximization: near-optimal time complexity meets practical efficiency. In *International Conference on Management of Data, SIGMOD 2014*, 75–86. ACM.
- Thompson, W. R. 1933. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4): 285–294.
- Vaswani, S.; Lakshmanan, L. V. S.; and Schmidt, M. 2016. Influence Maximization with Bandits. arXiv:1503.00024.
- Wang, Q.; and Chen, W. 2017. Improving Regret Bounds for Combinatorial Semi-Bandits with Probabilistically Triggered Arms and Its Applications. In *Advances in Neural Information Processing Systems 30 (NIPS 2017)*, 1161–1171.
- Wen, Z.; Kveton, B.; Valko, M.; and Vaswani, S. 2017. Online Influence Maximization under Independent Cascade Model with Semi-Bandit Feedback. In *Advances in Neural Information Processing Systems 30 (NIPS 2017)*, 3022–3032.
- Wu, Q.; Li, Z.; Wang, H.; Chen, W.; and Wang, H. 2019. Factorization Bandits for Online Influence Maximization. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2019*, 636–646. ACM.

Appendix

A Proof of Theorem 1

OIM with node-level feedback under the IC model is closely related to the *generalized linear bandits* (GLB) problem, introduced by (Filippi et al. 2010) and further investigated in (Li, Lu, and Zhou 2017). To see this, we introduce the GLB model. Assume there are T rounds in total. At round t , an action $X_t \in \mathbb{R}^d$ with $X_t \neq \mathbf{0}$ is chosen. Assume there is an unknown parameter $\theta^* \in \mathbb{R}^d$ and a *link function* $\mu : \mathbb{R} \rightarrow \mathbb{R}$. The reward $Y_t \in \mathbb{R}$ given X_t satisfies that

$$\mathbf{E}[Y_t | X_t] = \mu(X_t^\top \theta^*).$$

The goal is to minimize the cumulative regret over T rounds.

Specific to our problem, $X_t \in \{0, 1\}^d$ and $Y_t \in \{0, 1\}$. The *link function* $\mu : \mathbb{R} \rightarrow \mathbb{R}$ is defined as $\mu(x) = 1 - \exp(-x)$. Besides, $\theta^* \in \Theta := \{\theta \in \mathbb{R}^d \mid \theta(e) \geq 0, \forall e \in [d], \sum_{e \in [d]} \theta(e) \leq \ln(1/\gamma)\}$, which is induced by Assumption 1. However, our problem is more difficult in that multiple data pairs $(X_{t,j,v}, Y_{t,j,v})$ may be generated for each $v \in V$ at each round t , and the reward is specified by a more involved influence function.

The following lemma gives some properties about μ , which are easy to verify and useful for our analysis.

Lemma 3. *For any $\theta \in \Theta$ and $X \in \{0, 1\}^d$ with $X \neq \mathbf{0}$, it satisfies that*

$$\dot{\mu}(X^\top \theta) \leq 1, \dot{\mu}(X^\top \theta) \geq \gamma, \text{ and } |\ddot{\mu}(X^\top \theta)| \leq 1.$$

The analysis in (Li, Lu, and Zhou 2017) cannot be applied here directly for (X_k, Y_k) , since it requires that the gradient of the log-likelihood function has the form $\sum_{k=1}^t [Y_k - \mu(X_k^\top \theta)] X_k$. We thus define a pseudo log-likelihood function by ‘‘integrating’’ the gradient to apply the analysis. As in the main text, we gain an estimate $\hat{\theta}_t$ of θ^* at round $t + 1$ by maximizing the following pseudo log-likelihood function:

$$\hat{\theta}_t = \operatorname{argmax} L_t(\theta), \tag{7}$$

where

$$L_t(\theta) = \sum_{k=1}^t [-\exp(-X_k^\top \theta) - (1 - Y_k) X_k^\top \theta].$$

The following theorem characterizes the confidence intervals of θ^* induced by $\hat{\theta}_t$, which is the same as Theorem 1. We will prove it instead of Theorem 1.

Theorem 3. *Assume that $\theta^* \in \Theta$. Define $M_t = \sum_{k=1}^t X_k X_k^\top$ and let $\hat{\theta}_t$ be defined as in Eq. (7). Given $\delta \in (0, 1)$, assume that*

$$\lambda_{\min}(M_t) \geq \frac{512d}{\gamma^4} \left(d^2 + \ln \frac{1}{\delta} \right).$$

Then, with probability at least $1 - 3\delta$, for any $x \in \mathbb{R}^d$, it satisfies that

$$|x^\top (\hat{\theta}_t - \theta^*)| \leq \frac{3}{\gamma} \sqrt{\ln(1/\delta)} \cdot \|x\|_{M_t^{-1}}.$$

Proof. The proof consists of two steps. In the first step, it is proved that $\hat{\theta}_t$ falls into the η -neighborhood \mathcal{B}_η of θ^* w.r.t. ℓ_2 -norm for some η . Since $\dot{\mu}(X^\top \theta^*) \geq \gamma$, for any $\theta \in \mathcal{B}_\eta$, $\dot{\mu}(X^\top \theta)$ also has a lower bound denoted by κ_η . The values of η and κ_η will be determined later. In the second step, it is proved that $\hat{\theta}_t$ and θ^* are close in any direction $x \in \mathbb{R}^d$.

First note that $\hat{\theta}_t$ satisfies that $\nabla L_t(\hat{\theta}_t) = 0$, where the gradient $\nabla L_t(\theta)$ is

$$\nabla L_t(\theta) = \sum_{k=1}^t [\exp(-X_k^\top \theta) - (1 - Y_k)] X_k = \sum_{k=1}^t [Y_k - \mu(X_k^\top \theta)] X_k.$$

Define $G(\theta) := \sum_{k=1}^t (\mu(X_k^\top \theta) - \mu(X_k^\top \theta^*)) X_k$. Then, we have

$$G(\theta^*) = 0 \text{ and } G(\hat{\theta}_t) = \sum_{k=1}^t \epsilon_k X_k,$$

where ϵ_k is defined as $\epsilon_k := Y_k - \mu(X_k^\top \theta^*)$. Note that $\mathbf{E}[\epsilon_k | X_k] = 0$ and $\epsilon_k = Y_k - \mu(X_k^\top \theta^*) \in [-1, 1]$ since $Y_k \in \{0, 1\}$ and $\mu(X_k^\top \theta^*) = \Pr[Y_k = 1 | X_k] \in [0, 1]$. Therefore, ϵ_k is 1-sub-Gaussian, i.e. $\mathbf{E}[\exp(\lambda \epsilon_k) | X_k] \leq \exp(\lambda^2/2), \forall \lambda \in \mathbb{R}$. Further, define $Z := G(\hat{\theta}_t) = \sum_{k=1}^t \epsilon_k X_k$ for convenience.

Step 1: Consistency of $\hat{\theta}_t$. We first prove the consistency of $\hat{\theta}_t$. For any $\theta_1, \theta_2 \in \mathbb{R}_+^d := \{\theta \in \mathbb{R}^d \mid \theta(e) \geq 0, \forall e \in [d]\}$, by the mean value theorem, there is some $\bar{\theta} = s\theta_1 + (1-s)\theta_2$ with $0 < s < 1$ such that

$$G(\theta_1) - G(\theta_2) = \left[\sum_{k=1}^t \dot{\mu}(X_k^\top \bar{\theta}) X_k X_k^\top \right] (\theta_1 - \theta_2) := F(\bar{\theta})(\theta_1 - \theta_2).$$

Since $\dot{\mu}(x) = \exp(-x)$, for $\bar{\theta} \in \mathbb{R}_+^d$, $\dot{\mu}(X^\top \bar{\theta}) > 0$. Together with $\lambda_{\min}(M_t) > 0$, we have $\lambda_{\min}(F(\bar{\theta})) > 0$. Therefore, for any $\theta_1 \neq \theta_2$,

$$(\theta_1 - \theta_2)^\top (G(\theta_1) - G(\theta_2)) = (\theta_1 - \theta_2)^\top F(\bar{\theta})(\theta_1 - \theta_2) > 0.$$

Consequently, $G(\theta)$ is an injection from \mathbb{R}^d to \mathbb{R}^d and therefore G^{-1} is well-defined. We thus have $\hat{\theta} = G^{-1}(Z)$.

Let $\mathcal{B}_\eta := \{\theta \mid \|\theta - \theta^*\| \leq \eta\}$ be the η -neighborhood of θ^* and $\partial\mathcal{B}_\eta := \{\theta \mid \|\theta - \theta^*\| = \eta\}$. Define $\kappa_\eta := \inf_{\theta \in \mathcal{B}_\eta, X \neq 0} \dot{\mu}(X^\top \theta) > 0$. The following lemma shows that if $G(\theta)$ and $G(\theta^*)$ are close, then θ and θ^* are also close. Its proof is presented in Appendix A.1.

Lemma 4. $\{\theta \mid \|G(\theta)\|_{M_t^{-1}} \leq \kappa_\eta \eta \sqrt{\lambda_{\min}(M_t)}\} \subseteq \mathcal{B}_\eta$.

Next, in the following lemma, we give an upper bound of $\|Z\|_{M_t^{-1}} = \|G(\hat{\theta}_t)\|_{M_t^{-1}}$, which shows that $G(\hat{\theta}_t)$ and $G(\theta^*)$ are indeed close. Its proof is presented in Appendix A.2.

Lemma 5. For any $\delta > 0$, define the following event:

$$\mathcal{E}_G := \{\|Z\|_{M_t^{-1}} \leq 4\sqrt{d + \ln(1/\delta)}\}.$$

Then, \mathcal{E}_G holds with probability at least $1 - \delta$.

By the above lemmas, when \mathcal{E}_G holds, for any η , $\eta \geq \frac{4}{\kappa_\eta} \sqrt{\frac{d + \ln(1/\delta)}{\lambda_{\min}(M_t)}}$ implies that $\|\hat{\theta}_t - \theta^*\| \leq \eta$.

It remains to determine appropriate η and κ_η for the second step. Note that since $\|\hat{\theta}_t - \theta^*\|_1 \leq \sqrt{d} \cdot \|\hat{\theta}_t - \theta^*\| \leq \sqrt{d}\eta$, we have $\sum_{e \in [d]} \hat{\theta}(e) \leq \ln(1/\gamma) + \sqrt{d}\eta$. Therefore, we choose $\eta = \ln(1 + \epsilon)/\sqrt{d}$ for some ϵ to be determined later. In this case, $\kappa_\eta = \gamma/(1 + \epsilon) := \bar{\gamma}$.

To summarize, when $\lambda_{\min}(M_t) \geq \frac{16(d + \ln(1/\delta))}{\kappa_\eta^2 \eta^2} = \frac{16d(d + \ln(1/\delta))}{\bar{\gamma}^2 \ln^2(1 + \epsilon)}$, $\|\hat{\theta}_t - \theta^*\| \leq \ln(1 + \epsilon)/\sqrt{d}$ with probability at least $1 - \delta$.

Step 2: Normality of $\hat{\theta}$. In the following, we assume that $\hat{\theta}_t$ falls in the η -neighborhood \mathcal{B}_η of θ^* , where $\eta = \ln(1 + \epsilon)/\sqrt{d}$, $\kappa_\eta = \bar{\gamma} = \gamma/(1 + \epsilon)$ and ϵ is set to be the largest value such that $32^2(1 + \epsilon)^6 \leq 512(3 - \sqrt{2}(1 + \epsilon))^2$. Define $\Delta := \hat{\theta}_t - \theta^*$. The previous argument shows that there exists a $s \in [0, 1]$ such that

$$Z = G(\hat{\theta}_t) - G(\theta^*) = (H + E)\Delta,$$

where $\bar{\theta} = s\theta^* + (1-s)\hat{\theta}_t \in \mathcal{B}_\eta$, $H := F(\theta^*) = \sum_{k=1}^t \dot{\mu}(X_k^\top \theta^*) X_k X_k^\top$ and $E := F(\bar{\theta}) - F(\theta^*)$. For any $x \in \mathbb{R}^d$,

$$x^\top (\hat{\theta}_t - \theta^*) = x^\top (H + E)^{-1} Z = x^\top H^{-1} Z - x^\top H^{-1} E (H + E)^{-1} Z.$$

Note that $(H + E)^{-1}$ exists since $H + E = F(\bar{\theta}) \succ \bar{\gamma} M_t \succ 0$. We now bound the two terms, respectively.

For the first term, define

$$D := (X_1, X_2, \dots, X_t)^\top \in \mathbb{R}^{n \times d}.$$

Note that $D^\top D = \sum_{k=1}^t X_k X_k^\top = M_t$. Since ϵ_k is 1-sub-Gaussian, by the Hoeffding inequality,

$$\Pr[|x^\top H^{-1} Z| \geq a] = \Pr\left[\left|\sum_{k=1}^t x^\top H^{-1} X_k \epsilon_k\right| \geq a\right] \leq \exp\left(-\frac{a^2}{2\|x^\top H^{-1} D^\top\|^2}\right).$$

Since $H \succeq \bar{\gamma} M_t$, we have

$$\|x^\top H^{-1} D^\top\|^2 = x^\top H^{-1} D^\top D H^{-1} x \leq \frac{1}{\bar{\gamma}^2} \|x\|_{M_t^{-1}}^2.$$

Thus we have

$$\Pr[|x^\top H^{-1} Z| \geq a] \leq \exp\left(-\frac{a^2 \bar{\gamma}^2}{2\|x\|_{M_t^{-1}}^2}\right).$$

By choosing an appropriate a , we obtain that with probability at least $1 - 2\delta$,

$$|x^\top H^{-1}Z| \leq \frac{\sqrt{2\ln(1/\delta)}}{\bar{\gamma}} \|x\|_{M_t^{-1}}.$$

For the second term,

$$\begin{aligned} |x^\top H^{-1}E(H+E)^{-1}Z| &\leq \|x\|_{H^{-1}} \|H^{-1/2}E(H+E)^{-1}Z\| \\ &\leq \|x\|_{H^{-1}} \|H^{-1/2}E(H+E)^{-1}H^{1/2}\| \|Z\|_{H^{-1}} \\ &\leq \frac{1}{\bar{\gamma}} \|x\|_{M_t^{-1}} \|H^{-1/2}E(H+E)^{-1}H^{1/2}\| \|Z\|_{M_t^{-1}}. \end{aligned}$$

The first inequality is due to Cauchy-Schwarz inequality. The second inequality holds since $\|AB\| \leq \|A\|\|B\|$. The last inequality holds since $H \succeq \bar{\gamma}M_t$. Next, since $(H+E)^{-1} = H^{-1} - H^{-1}E(H+E)^{-1}$, we have

$$\begin{aligned} \|H^{-1/2}E(H+E)^{-1}H^{1/2}\| &= \|H^{-1/2}E(H^{-1} - H^{-1}E(H+E)^{-1})H^{1/2}\| \\ &= \|H^{-1/2}EH^{-1/2} - H^{-1/2}EH^{-1}E(H+E)^{-1}H^{1/2}\| \\ &\leq \|H^{-1/2}EH^{-1/2}\| + \|H^{-1/2}EH^{-1/2}\| \|H^{-1/2}E(H+E)^{-1}H^{1/2}\| \end{aligned}$$

To complete our proof, we need the following technical lemma, whose proof is presented in Appendix A.3.

Lemma 6.

$$\|H^{-1/2}EH^{-1/2}\| \leq \frac{4}{\bar{\gamma}^2} \sqrt{\frac{d(d + \ln 1/\delta)}{\lambda_{\min}(M_t)}}.$$

Specifically, when $\lambda_{\min}(M_t) \geq 64d(d + \ln(1/\delta))/\bar{\gamma}^4$,

$$\|H^{-1/2}EH^{-1/2}\| \leq 1/2.$$

Therefore, we have

$$\|H^{-1/2}E(H+E)^{-1}H^{1/2}\| \leq \frac{\|H^{-1/2}EH^{-1/2}\|}{1 - \|H^{-1/2}EH^{-1/2}\|} \leq 2\|H^{-1/2}EH^{-1/2}\| \leq \frac{8}{\bar{\gamma}^2} \sqrt{\frac{d(d + \ln 1/\delta)}{\lambda_{\min}(M_t)}}.$$

Therefore, together with Lemma 5, we have

$$|x^\top H^{-1}E(H+E)^{-1}Z| \leq \frac{32\sqrt{d}(d + \ln 1/\delta)}{\bar{\gamma}^3 \sqrt{\lambda_{\min}(M_t)}} \|x\|_{M_t^{-1}}.$$

Combining the above inequalities, we have

$$|x^\top (\hat{\theta}_t - \theta^*)| \leq \left(\frac{\sqrt{2\ln(1/\delta)}}{\bar{\gamma}} + \frac{32\sqrt{d}(d + \ln 1/\delta)}{\bar{\gamma}^3 \sqrt{\lambda_{\min}(M_t)}} \right) \|x\|_{M_t^{-1}} \leq \frac{3\sqrt{\ln(1/\delta)}}{(1 + \epsilon)\bar{\gamma}} \|x\|_{M_t^{-1}} = \frac{3\sqrt{\ln(1/\delta)}}{\gamma} \|x\|_{M_t^{-1}}.$$

By the choice of ϵ , the last inequality holds when

$$\lambda_{\min}(M_t) \geq \frac{512d(d + \ln(1/\delta))^2}{\gamma^4 \ln(1/\delta)}.$$

The proof is completed. \square

A.1 Proof of Lemma 4

This lemma is a direct application of Lemma A of (Chen, Hu, and Ying 1999). For completeness, we restate it in the lemma below.

Lemma 7 ((Chen, Hu, and Ying 1999)). *Let H be a smooth injection from \mathbb{R}^d to \mathbb{R}^d with $H(x_0) = y_0$. Define $\mathcal{B}_\delta(x_0) := \{x \in \mathbb{R}^d \mid \|x - x_0\| \leq \delta\}$ and $\partial\mathcal{B}_\delta(x_0) := \{x \in \mathbb{R}^d \mid \|x - x_0\| = \delta\}$. Then $\inf_{x \in \partial\mathcal{B}_\delta(x_0)} \|H(x) - y_0\| \geq r$ implies*

1. $\mathcal{B}_r(y_0) := \{y \in \mathbb{R}^d \mid \|y - y_0\| \leq r\} \subseteq H(\mathcal{B}_\delta(x_0))$.
2. $H^{-1}(\mathcal{B}_r(y_0)) \subseteq \mathcal{B}_\delta(x_0)$.

For any $\theta \in \partial\mathcal{B}_\eta$, there is some $\bar{\theta} = s\theta + (1-s)\theta^* \in \mathcal{B}_\eta$ with $0 < s < 1$ such that $G(\theta) - G(\theta^*) = F(\bar{\theta})(\theta - \theta^*)$, where $F(\bar{\theta}) = \sum_{k=1}^t \dot{\mu}(X_k^\top \bar{\theta}) X_k X_k^\top \succeq \kappa_\eta M_t$, and

$$\|G(\theta)\|_{M_t^{-1}}^2 = \|G(\theta) - G(\theta^*)\|_{M_t^{-1}}^2 = (\theta - \theta^*)^\top F(\bar{\theta}) M_t^{-1} F(\bar{\theta})(\theta - \theta^*) \geq \kappa_\eta^2 \lambda_{\min}(M_t) \|\theta - \theta^*\|^2 = \kappa_\eta^2 \eta^2 \lambda_{\min}(M_t).$$

By Lemma A of (Chen, Hu, and Ying 1999), the proof is completed.

A.2 Proof of Lemma 5

Let $\langle \cdot, \cdot \rangle$ denote the inner product. Note that

$$\|Z\|_{M_t^{-1}} = \|M_t^{-1/2}Z\| = \sup_{\|y\| \leq 1} \langle y, M_t^{-1/2}Z \rangle.$$

Let $\hat{\mathbb{B}}$ be a $1/2$ -net of the unit ball $\mathbb{B}^d = \{y \in \mathbb{R}^d \mid \|y\| \leq 1\}$. Then $|\hat{\mathbb{B}}| \leq 6^d$ (Pollard 1990), and for any $x \in \mathbb{B}^d$, there is a $\hat{x} \in \hat{\mathbb{B}}$ such that $\|\hat{x} - x\| \leq 1/2$. Thus,

$$\begin{aligned} \langle x, M_t^{-1/2}Z \rangle &= \langle \hat{x}, M_t^{-1/2}Z \rangle + \langle x - \hat{x}, M_t^{-1/2}Z \rangle \\ &= \langle \hat{x}, M_t^{-1/2}Z \rangle + \|x - \hat{x}\| \left\langle \frac{x - \hat{x}}{\|x - \hat{x}\|}, M_t^{-1/2}Z \right\rangle \\ &\leq \langle \hat{x}, M_t^{-1/2}Z \rangle + \frac{1}{2} \sup_{\|y\| \leq 1} \langle y, M_t^{-1/2}Z \rangle. \end{aligned}$$

By taking supremum on both sides, we obtain that

$$\sup_{\|y\| \leq 1} \langle y, M_t^{-1/2}Z \rangle \leq 2 \max_{\hat{x} \in \hat{\mathbb{B}}} \langle \hat{x}, M_t^{-1/2}Z \rangle.$$

Finally, define $D := (X_1, X_2, \dots, X_t)^\top \in \mathbb{R}^{t \times d}$. Then, $D^\top D = M_t$. We have

$$\begin{aligned} \Pr[\|Z\|_{M_t^{-1}} > a] &\leq \Pr[\max_{\hat{x} \in \hat{\mathbb{B}}} \langle \hat{x}, M_t^{-1/2}Z \rangle > a/2] \\ &\leq \sum_{\hat{x} \in \hat{\mathbb{B}}} \Pr[\langle \hat{x}, M_t^{-1/2}Z \rangle > a/2] \\ &\leq \sum_{\hat{x} \in \hat{\mathbb{B}}} \exp\left(-\frac{a^2}{8\|\hat{x}^\top M_t^{-1/2}D^\top\|^2}\right) \\ &\leq \exp(-a^2/8 + d \ln 6) \leq \delta. \end{aligned}$$

The second to last inequality holds due to Hoeffding inequality. The last inequality holds by choosing $a = 4\sqrt{d + \ln(1/\delta)}$.

A.3 Proof of Lemma 6

By the mean value theorem,

$$E = \sum_{k=1}^t (\dot{\mu}(X_k^\top \bar{\theta}) - \dot{\mu}(X_k^\top \theta^*)) X_k X_k^\top = \sum_{k=1}^t \ddot{\mu}(r_k) X_k^\top \Delta X_k X_k^\top.$$

for some $r_k \in \mathbb{R}$. Since $|\ddot{\mu}| \leq 1$, for any $x \in \mathbb{R}^d \setminus \{0\}$, we have

$$\begin{aligned} x^\top H^{-1/2} E H^{-1/2} x &= \sum_{k=1}^t \ddot{\mu}(r_k) X_k^\top \Delta \|x^\top H^{-1/2} X_k\|^2 \\ &\leq \sqrt{d} \|\Delta\| \left(x^\top H^{-1/2} \left(\sum_{k=1}^t X_k X_k^\top \right) H^{-1/2} x \right) \\ &\leq \frac{\sqrt{d}}{\bar{\gamma}} \|\Delta\| \|x\|^2. \end{aligned}$$

The first inequality is due to Cauchy-Schwarz inequality. The second inequality holds since $\|X_k\| \leq \sqrt{d}$. The last inequality holds since $H \succ \bar{\gamma} M_t$. Therefore, by the definition of spectral norm of a matrix,

$$\|H^{-1/2} E H^{-1/2}\| \leq \frac{\sqrt{d}}{\bar{\gamma}} \|\Delta\| \leq \frac{4}{\bar{\gamma}^2} \sqrt{\frac{d(d + \ln 1/\delta)}{\lambda_{\min}(M_t)}}.$$

When $\lambda_{\min}(M_t) \geq 64d(d + \ln(1/\delta))/\bar{\gamma}^4$, we have

$$\|H^{-1/2} E H^{-1/2}\| \leq 1/2.$$

B Proof of Lemma 1

In this section, we present the proof of Lemma 1, the GOM bounded smoothness for the IC model. The proof is similar to that for the LT model in (Li et al. 2020).

Fix any seed set $S \subseteq V$ throughout the proof. For edge probability vector p^* , let $\sigma(S, p^*, v)$ be the probability that v is activated under p^* . Let G' be a random sub-graph of the original graph G such that each edge $e \in E$ of G appears in G' independently with probability $p^*(e)$. By the diffusion rule of the IC model, $\sigma(S, p^*, v)$ also means the probability that there is a path from S to v in G' . For two distinct edge probability vectors p^* and \tilde{p} , we couple them in the following way. Define a random vector $r \sim U[0, 1]^m$ which satisfies that $r(e) \sim U[0, 1]$ for each edge $e \in E$, where $U[0, 1]$ denotes the uniform distribution over interval $[0, 1]$. Given edge probability vector p^* , for $e \in E$, e appears in G' if and only if $r(e) \leq p^*(e)$. It is easy to see that the probability that e appears in G' is $p^*(e)$. For \tilde{p} , the random experiment can be executed using the same r . In this way, p^* and \tilde{p} are coupled.

For node $v \in V$, define event

$$\mathcal{E}_{0,v} := \mathbf{1}\{v \text{ is activated under } \tilde{p}\} \neq \mathbf{1}\{v \text{ is activated under } p^*\}.$$

According to our notations,

$$|\sigma(S, \tilde{p}) - \sigma(S, p^*)| = \left| \sum_{v \in V \setminus S} \sigma(S, \tilde{p}, v) - \sigma(S, p^*, v) \right| \leq \sum_{v \in V \setminus S} |\sigma(S, \tilde{p}, v) - \sigma(S, p^*, v)| = \sum_{v \in V \setminus S} \Pr_{r \sim U[0,1]^m} [\mathcal{E}_{0,v}].$$

Recall that $V[S, v]$ is defined to be the set of nodes relevant to v given seed set S , namely for each $u \in V[S, v]$, there is a path P from S to v such that $u \in P$. Thus, the occurrence of $\mathcal{E}_{0,v}$ means there is a node $u \in V[S, v]$ such that its activation state is different under \tilde{p} and p^* . More strictly, let $\Phi(p^*, r) := (S_0 = S, S_1, \dots, S_{n-1})$ be the sequence of active nodes under p^* and r . Let $\Phi_\tau(p^*, r) = S_\tau$ be the set of active nodes immediately after time τ . For node $u \in V[S, v]$, define $\mathcal{E}_1(u)$ to describe that u is the first node that possesses distinct activation states under \tilde{p} and p^* , namely

$$\mathcal{E}_1(u) := \{r \mid \exists \tau, \forall \tau' < \tau, \Phi_{\tau'}(p^*, r) = \Phi_{\tau'}(\tilde{p}, r), u \in (\Phi_\tau(p^*, r) \setminus \Phi_\tau(\tilde{p}, r)) \cup (\Phi_\tau(\tilde{p}, r) \setminus \Phi_\tau(p^*, r))\}.$$

Then, by the above argument,

$$\mathcal{E}_{0,v} \subseteq \bigcup_{u \in V[S, v]} \mathcal{E}_1(u).$$

Therefore, by union bound,

$$|\sigma(S, \tilde{p}) - \sigma(S, p^*)| \leq \sum_{v \in V \setminus S} \sum_{u \in V[S, v]} \Pr_{r \sim U[0,1]^m} [\mathcal{E}_1(u)]. \quad (8)$$

To further determine the probability that $\mathcal{E}_1(u)$ occurs, for each $0 \leq \tau \leq n-1$, we define the following event:

$$\mathcal{E}_{2,0}(u, \tau) := \{r \mid \forall \tau' < \tau, \Phi_{\tau'}(p^*, r) = \Phi_{\tau'}(\tilde{p}, r), u \notin \Phi_{\tau-1}(p^*, r)\}.$$

$$\mathcal{E}_{2,1}(u, \tau) := \{r \mid \forall \tau' < \tau, \Phi_{\tau'}(p^*, r) = \Phi_{\tau'}(\tilde{p}, r), u \in (\Phi_\tau(p^*, r) \setminus \Phi_\tau(\tilde{p}, r)) \cup (\Phi_\tau(\tilde{p}, r) \setminus \Phi_\tau(p^*, r))\}.$$

$$\mathcal{E}_{3,1}(u, \tau) := \{r \mid u \in (\Phi_\tau(p^*, r) \setminus \Phi_\tau(\tilde{p}, r)) \cup (\Phi_\tau(\tilde{p}, r) \setminus \Phi_\tau(p^*, r))\}.$$

Note that event $\mathcal{E}_{2,1}(u, \tau)$ means that event $\mathcal{E}_1(u)$ occurs in time τ . Hence,

$$\Pr_{r \sim U[0,1]^m} [\mathcal{E}_1(u)] = \sum_{\tau=0}^{n-1} \Pr_{r \sim U[0,1]^m} [\mathcal{E}_{2,1}(u, \tau)].$$

Next, we estimate the value of $\Pr_{r \sim U[0,1]^m} [\mathcal{E}_{2,1}(u, \tau)]$. For $u \in V$, let $r_u \in [0, 1]^{d_u}$ denote the sub-vector which consists of r 's entries over the incoming edges of u . Let $r_{-u} \in [0, 1]^{m-d_u}$ denote the vector consisting of the remaining entries. By fixing r_{-u} , we define sub-event $\mathcal{E}_{2,1}(u, \tau, r_{-u}) \subseteq \mathcal{E}_{2,1}(u, \tau)$ to be the restriction of event $\mathcal{E}_{2,1}(u, \tau)$ when r_{-u} is fixed. Similarly, we can define $\mathcal{E}_{2,0}(u, \tau, r_{-u}) \subseteq \mathcal{E}_{2,0}(u, \tau)$ and $\mathcal{E}_{3,1}(u, \tau, r_{-u}) \subseteq \mathcal{E}_{3,1}(u, \tau)$. Note that

$$\mathcal{E}_{2,1}(u, \tau) = \mathcal{E}_{2,0}(u, \tau) \cap \mathcal{E}_{3,1}(u, \tau).$$

$$\mathcal{E}_{2,1}(u, \tau, r_{-u}) = \mathcal{E}_{2,0}(u, \tau) \cap \mathcal{E}_{3,1}(u, \tau, r_{-u}).$$

Therefore,

$$\Pr_{r \sim U[0,1]^m} [\mathcal{E}_{2,1}(u, \tau)] = \Pr_{r \sim U[0,1]^m} [\mathcal{E}_{2,0}(u, \tau)] \cdot \Pr_{r \sim U[0,1]^m} [\mathcal{E}_{3,1}(u, \tau) \mid \mathcal{E}_{2,0}(u, \tau)] \leq \Pr_{r \sim U[0,1]^m} [\mathcal{E}_{3,1}(u, \tau) \mid \mathcal{E}_{2,0}(u, \tau)],$$

and

$$\begin{aligned} \Pr_{r_u \sim U[0,1]^{d_u}} [\mathcal{E}_{2,1}(u, \tau, r_{-u})] &= \Pr_{r_u \sim U[0,1]^{d_u}} [\mathcal{E}_{2,0}(u, \tau, r_{-u})] \cdot \Pr_{r_u \sim U[0,1]^{d_u}} [\mathcal{E}_{3,1}(u, \tau, r_{-u}) \mid \mathcal{E}_{2,0}(u, \tau, r_{-u})] \\ &\leq \Pr_{r_u \sim U[0,1]^{d_u}} [\mathcal{E}_{3,1}(u, \tau, r_{-u}) \mid \mathcal{E}_{2,0}(u, \tau, r_{-u})]. \end{aligned} \quad (9)$$

By the definition of event $\mathcal{E}_{2,0}(u, \tau, r_{-u})$, by the end of time $\tau - 1$, the sets of active nodes are the same under p^* and \tilde{p} , and u is inactive at the time. Besides, since r_{-u} is fixed, the set of active nodes by the end of time $\tau - 1$ is also fixed. We use $\Phi_{\tau'}(\mathcal{E}_{2,0}(u, \tau, r_{-u}))$ to denote the set of active nodes by the end of time τ' under event $\mathcal{E}_{2,0}(u, \tau, r_{-u})$. Consider the probability that event $\mathcal{E}_{3,1}(u, \tau, r_{-u})$ occurs conditioned on event $\mathcal{E}_{2,0}(u, \tau, r_{-u})$. Define $Q(u, \tau, r_{-u})$ to be the neighbors of u which were just activated in time $\tau - 1$, namely

$$Q(u, \tau, r_{-u}) := (\Phi_{\tau-1}(\mathcal{E}_{2,0}(u, \tau, r_{-u})) \setminus \Phi_{\tau-2}(\mathcal{E}_{2,0}(u, \tau, r_{-u}))) \cap N(u).$$

Let $Z = Z(u, \tau, r_{-u}) \in \{0, 1\}^{d_u}$ be the characteristic vector of $Q(u, \tau, r_{-u})$. By the diffusion rule of the IC model,

$$\Pr_{r_u \sim U[0,1]^{d_u}} [u \in \Phi_{\tau}(p^*, r) \mid \mathcal{E}_{2,0}(u, \tau, r_{-u})] = 1 - \prod_{u' \in Q} (1 - p^*(e_{u'u})) = \mu(Z^\top \theta_u^*),$$

where the definition of θ_u^* can be found in eq. (2) and $\mu : \mathbb{R} \rightarrow \mathbb{R}$ is the link function which has the form $\mu(x) = 1 - \exp(-x)$. Therefore,

$$\Pr_{r_u \sim U[0,1]^{d_u}} [\mathcal{E}_{3,1}(u, \tau, r_{-u}) \mid \mathcal{E}_{2,0}(u, \tau, r_{-u})] \leq |\mu(Z^\top \theta_u^*) - \mu(Z^\top \tilde{\theta}_u)| \leq |Z^\top (\theta_u^* - \tilde{\theta}_u)|.$$

The last inequality holds since μ is 1-Lipschitz. Combining with eq. (9), we obtain that

$$\Pr_{r_u \sim U[0,1]^{d_u}} [\mathcal{E}_{2,1}(u, \tau, r_{-u})] \leq |Z(u, \tau, r_{-u})^\top (\theta_u^* - \tilde{\theta}_u)|. \quad (10)$$

Event $\mathcal{E}_{2,0}(u, \tau, r_{-u})$ depends on both p^* and \tilde{p} , which is hard to analyze. For this reason, we define event

$$\mathcal{E}_{4,0}(u, \tau, r_{-u}) := \{\theta = (r_{-u}, r_u) \mid u \notin \Phi_{\tau-1}(p^*, r)\}.$$

Clearly, $\mathcal{E}_{2,0}(u, \tau, r_{-u}) \subseteq \mathcal{E}_{4,0}(u, \tau, r_{-u})$. Besides, when $\mathcal{E}_{2,0}(u, \tau, r_{-u}) \neq \emptyset$, since r_{-u} is fixed in events $\mathcal{E}_{2,0}(u, \tau, r_{-u})$ and $\mathcal{E}_{4,0}(u, \tau, r_{-u})$ and u is inactive by the end of time $\tau - 1$, the active nodes by the end of time $\tau - 1$ are the same under the two events. Define $P(u, \tau, r_{-u})$ be the set of u 's neighbors which are just activated in time $\tau - 1$ under event $\mathcal{E}_{4,0}(u, \tau, r_{-u})$, namely

$$P(u, \tau, r_{-u}) := (\Phi_{\tau-1}(\mathcal{E}_{4,0}(u, \tau, r_{-u})) \setminus \Phi_{\tau-2}(\mathcal{E}_{4,0}(u, \tau, r_{-u}))) \cap N(u).$$

Then, $Q(u, \tau, r_{-u}) = P(u, \tau, r_{-u})$. Let $X(u, \tau, r_{-u}) \in \{0, 1\}^{d_u}$ be the characteristic vector of $P(u, \tau, r_{-u})$. By the above argument and eq. (10),

$$\Pr_{r_u \sim U[0,1]^{d_u}} [\mathcal{E}_{2,1}(u, \tau, r_{-u})] \leq |X(u, \tau, r_{-u})^\top (\theta_u^* - \tilde{\theta}_u)|.$$

When $\mathcal{E}_{2,0}(u, \tau, r_{-u}) = \emptyset$, the left-hand side of the above inequality is 0. Thus, the inequality still holds. To sum up,

$$\begin{aligned} \Pr_{r \in U[0,1]^m} [\mathcal{E}_1(u)] &= \int_{r_{-u} \sim [0,1]^{m-d_u}} \sum_{\tau=0}^{n-1} \Pr_{r_u \sim U[0,1]^{d_u}} [\mathcal{E}_{2,1}(u, \tau, r_{-u})] dr_{-u} \\ &\leq \int_{r_{-u} \in [0,1]^{m-d_u}} \sum_{\tau=0}^{n-1} |X(u, \tau, r_{-u})^\top (\theta_u^* - \tilde{\theta}_u)| dr_{-u} \\ &= \mathbf{E}_{r_{-u} \in U[0,1]^{m-d_u}} \left[\sum_{\tau=0}^{n-1} |X(u, \tau, r_{-u})^\top (\theta_u^* - \tilde{\theta}_u)| \right]. \end{aligned}$$

By plugging the above inequality into eq. (8), we obtain that

$$\begin{aligned} |\sigma(S, \tilde{p}) - \sigma(S, p^*)| &\leq \sum_{v \in V \setminus S} \sum_{u \in V[S, v]} \mathbf{E}_{r_{-u} \in U[0,1]^{m-d_u}} \left[\sum_{\tau=0}^{n-1} |X(u, \tau, r_{-u})^\top (\theta_u^* - \tilde{\theta}_u)| \right] \\ &\leq \sum_{v \in V \setminus S} \sum_{u \in V[S, v]} \mathbf{E} \left[\sum_{j=1}^{J_u} |X_{j,u}^\top (\tilde{\theta}_u - \theta_u^*)| \right]. \end{aligned}$$

The last inequality holds since by definition, $X_{j,u} \in \{0, 1\}^{d_u}$ indicates the characteristic vector of the edges corresponding to the j -th batch of neighbors that were activated.

C Proof of Lemma 2

Fix $v \in V$. For simplicity, let $z_{t,j} = \|X_{t,j,v}\|_{M_{t-1,v}^{-1}}$. Recall the definition of $M_{t,v}$ (see eq. (1)), we have

$$M_{t,v} = M_{t-1,v} + \sum_{j=1}^{J_{t,v}} X_{t,j,v} X_{t,j,v}^\top.$$

Basic algebra gives us that

$$\begin{aligned} \det[M_{t,v}] &\geq \det[M_{t-1,v} + X_{t,j,v} X_{t,j,v}^\top] \\ &= \det\left[M_{t-1,v}^{1/2} \left(I + M_{t-1,v}^{-1/2} X_{t,j,v} X_{t,j,v}^\top M_{t-1,v}^{-1/2}\right) M_{t-1,v}^{1/2}\right] \\ &= \det[M_{t-1,v}] \det\left[I + M_{t-1,v}^{-1/2} X_{t,j,v} X_{t,j,v}^\top M_{t-1,v}^{-1/2}\right] \\ &= \det[M_{t-1,v}] (1 + X_{t,j,v}^\top M_{t-1,v}^{-1} X_{t,j,v}) \\ &= \det[M_{t-1,v}] (1 + z_{t,j}^2). \end{aligned}$$

Thus, it can be deduced that

$$\det[M_{t,v}]^{J_{t,v}} \geq \det[M_{t-1,v}]^{J_{t,v}} \prod_{j=1}^{J_{t,v}} (1 + z_{t,j}^2).$$

Next, by $\det[M_{t,v}] \geq \det[M_{t-1,v}]$ and $J_{t,v} \leq d_v$, we have

$$\det[M_{t,v}]^{d_v} \geq \det[M_{t-1,v}]^{d_v} \prod_{j=1}^{J_{t,v}} (1 + z_{t,j}^2).$$

Therefore, we have

$$\det[M_{T,v}]^{d_v} \geq \det[M_{T_0,v}]^{d_v} \prod_{t=T_0+1}^T \prod_{j=1}^{J_{t,v}} (1 + z_{t,j}^2) = R^{d_v^2} \prod_{t=T_0+1}^T \prod_{j=1}^{J_{t,v}} (1 + z_{t,j}^2).$$

On the other hand,

$$\text{tr}[M_{T,v}] = \text{tr}\left[M_{T_0,v} + \sum_{t=T_0+1}^T \sum_{j=1}^{J_{t,v}} X_{t,j,v} X_{t,j,v}^\top\right] = \text{tr}[M_{T_0,v}] + \sum_{t=T_0+1}^T \sum_{j=1}^{J_{t,v}} \|X_{t,j,v}\|^2 \leq R d_v + (T - T_0) d_v^2.$$

By the trace-determinant inequality, we have

$$(R + (T - T_0) d_v)^{d_v} \geq \left(\frac{1}{d_v} \text{tr}[M_{T,v}]\right)^{d_v} \geq \det[M_{T,v}] \geq R^{d_v^2} \prod_{t=T_0+1}^T \prod_{j=1}^{J_{t,v}} (1 + z_{t,j}^2).$$

By taking logarithm on both sides, we obtain that

$$d_v \ln(R + (T - T_0) d_v) \geq \sum_{t=T_0+1}^T \sum_{j=1}^{J_{t,v}} \ln(1 + z_{t,j}^2) + d_v^2 \ln R \geq \sum_{t=T_0+1}^T \sum_{j=1}^{J_{t,v}} \frac{z_{t,j}^2}{d_v + 1} + d_v \ln R^{d_v}.$$

The last inequality holds since $\ln(1 + x) \geq \frac{x}{d_v + 1}$ for $x \in [0, d_v]$, and

$$z_{t,j}^2 = \|X_{t,j,v}\|_{M_{t-1,v}^{-1}}^2 \leq \|X_{t,j,v}\|^2 \leq d_v.$$

By rearranging the above inequality, we have

$$\sum_{t=T_0+1}^T \sum_{j=1}^{J_{t,v}} \|X_{t,j,v}\|_{M_{t-1,v}^{-1}}^2 = \sum_{t=T_0+1}^T \sum_{j=1}^{J_{t,v}} z_{t,j}^2 \leq d_v (d_v + 1) \cdot \ln\left(\frac{R + (T - T_0) d_v}{R^{d_v}}\right).$$

Finally, by the Cauchy-Schwarz inequality,

$$\begin{aligned}
\sum_{t=T_0+1}^T \sum_{v \in V \setminus S_t} n_{S_t, v} \sum_{j=1}^{J_{t, v}} \|X_{t, j, v}\|_{M_{t-1, v}^{-1}} &\leq \sqrt{\left(\sum_{t=T_0+1}^T \sum_{v \in V \setminus S_t} \sum_{j=1}^{J_{t, v}} n_{S_t, v}^2 \right) \left(\sum_{t=T_0+1}^T \sum_{v \in V \setminus S_t} \sum_{j=1}^{J_{t, v}} \|X_{t, j, v}\|_{M_{t-1, v}^{-1}}^2 \right)} \\
&\leq \sqrt{\left(\sum_{t=T_0+1}^T \sum_{v \in V} d_v n_{S_t, v}^2 \right) \left(\sum_{v \in V} d_v (d_v + 1) \cdot \ln \left(\frac{R + (T - T_0) d_v}{R^{d_v}} \right) \right)} \\
&\leq \zeta(G) D \sqrt{(m + n)(T - T_0) \ln(R + (T - T_0) D)}.
\end{aligned}$$

The last inequality holds since by the definition,

$$\sqrt{\sum_{v \in V} n_{S_t, v}^2} \leq \zeta(G).$$