

Scalable influence maximization for independent cascade model in large-scale social networks

Chi Wang · Wei Chen · Yajun Wang

Received: 2 May 2011 / Accepted: 16 February 2012
© The Author(s) 2012

Abstract Influence maximization, defined by Kempe et al. (SIGKDD 2003), is the problem of finding a small set of seed nodes in a social network that maximizes the spread of influence under certain influence cascade models. The scalability of influence maximization is a key factor for enabling prevalent viral marketing in large-scale online social networks. Prior solutions, such as the greedy algorithm of Kempe et al. (SIGKDD 2003) and its improvements are slow and not scalable, while other heuristic algorithms do not provide consistently good performance on influence spreads. In this article, we design a new heuristic algorithm that is easily scalable to millions of nodes and edges in our experiments. Our algorithm has a simple tunable parameter for users to control the balance between the running time and the influence spread of the algorithm. Our results from extensive simulations on several real-world and synthetic networks demonstrate that our algorithm is currently the best scalable solution to the influence maximization problem: (a) our algorithm scales beyond million-sized graphs where the greedy algorithm becomes infeasible, and (b) in all size ranges, our algorithm performs consistently well in influence spread—it is always among the best algorithms, and in most cases it significantly outperforms all other scalable heuristics to as much as 100–260% increase in influence spread.

Responsible editor: Fei Wang, Hanghang Tong, Philip Yu, Charu Aggarwal.

C. Wang (✉)
University of Illinois at Urbana-Champaign, Urbana, IL, USA
e-mail: chiwang1@illinois.edu

W. Chen · Y. Wang
Microsoft Research Asia, Beijing, China
e-mail: weic@microsoft.com

Y. Wang
e-mail: yajunw@microsoft.com

Keywords Influence maximization · Social networks · Independent cascade model · Viral marketing

1 Introduction

Word-of-mouth or viral marketing differentiates itself from other marketing strategies because it is based on trust among individuals' close social circle of families, friends, and co-workers. Research shows that people trust the information obtained from their close social circle far more than the information obtained from general advertisement channels such as TV, newspaper and online advertisements (Nail 2004). Thus many people believe that word-of-mouth marketing is the most effective marketing strategy (e.g. Misner 1999).

The increasing popularity of many online social network sites, such as Facebook, Myspace and Twitter, presents new opportunities for enabling large-scale and prevalent viral marketing online. Consider the following hypothetical scenario as a motivating example. A small company develops an online application and wants to market it through an online social network. It has a limited budget such that it can only select a small number of initial users in the network to use it (by giving them gifts or payments). The company wishes that these initial users would love the application and start influencing their friends on the social network to use it, and their friends would influence their friends' friends and so on, and thus through the word-of-mouth effect a large population in the social network would adopt the application. The problem is whom to select as the initial users so that they eventually influence the largest number of people in the network.

The above problem, called *influence maximization*, is first formulated as a discrete optimization problem by Kempe et al. (2003) as follows: A social network is modeled as a graph with nodes representing individuals and edges representing connections or relationship between two individuals. Influence is propagated in the network according to a stochastic cascade model, such as the following independent cascade (IC) model¹: Each edge (u, v) in the graph is associated with a *propagation probability* $pp(u, v)$, which is the probability that node u independently activates (a.k.a. influences) node v at step $t + 1$ if u is activated at step t . Given a social network graph, the IC model, and a small number k (e.g., set according to the budget for viral marketing), the influence maximization problem is to find k nodes in the graph (referred to as *seeds*) such that under the influence cascade model, the expected number of nodes activated by the k seeds (referred to as the *influence spread*) is the largest possible. Kempe et al. prove that the optimization problem is NP-hard, and present a greedy approximation algorithm guaranteeing that the influence spread is within $(1 - 1/e - \epsilon)$ of the optimal influence spread, where e is the base of natural logarithm, and ϵ depends on the accuracy of their Monte-Carlo estimate of the influence spread given a seed set.

However, their algorithm has a serious drawback—it is not scalable to large networks. A key element of their greedy algorithm is to compute the influence spread

¹ Other models are also introduced in Kempe et al. (2003), but in this article we focus on the independent cascade model.

given a seed set, which turns out to be a difficult task (in fact, as we point out in Sect. 2 the computation is #P-hard). Instead of finding an exact algorithm, Monte-Carlo simulations of the influence cascade model are run for a large number of times in order to obtain an accurate estimate of the influence spread. Consequently, even with the recent optimizations (Leskovec et al. 2007; Chen et al. 2009) that could achieve hundreds of times speedup, it still takes hours on a modern server to select 50 seeds in a moderate sized graph (15K nodes and 31K edges) while it becomes completely infeasible for larger graphs (e.g. more than 500K edges). Given that online social networks are typically of large-scale, we believe that the scalability issue of the greedy algorithm will be a fatal obstacle preventing it from supporting prevalent viral marketing activities in large-scale online social networks.

1.1 Our contribution

In this article, we first show that computing influence spread in the independent cascade model is #P-hard, which closes an open question posed by Kempe et al. (2003). It indicates that the greedy algorithm of Kempe et al. (2003) may have intrinsic difficulties to be made scalable for large graphs.

We then address the scalability issue by proposing a new heuristic algorithm that is several orders of magnitude faster than all existing greedy algorithms while matching the influence spread of the greedy algorithms. Our heuristic gains efficiency by restricting computations on the *local influence regions* of nodes. Moreover, by tuning the size of local influence regions, our heuristic is able to achieve tunable tradeoff between efficiency (in terms of running time) and effectiveness (in terms of influence spread). Our heuristic can easily scale up to handle networks with millions of nodes and edges, and at this scale it beats all other existing heuristics of similar scalability in terms of the influence spread.

The main idea of our heuristic scheme is to use local arborescence² structures of each node to approximate the influence propagation. We first compute *maximum influence paths* (MIP) between every pair of nodes in the network via Dijkstra's shortest-path algorithm, and ignore MIPs with probability smaller than an influence threshold θ , effectively restricting influence to a local region. We then union the MIPs starting or ending at each node into the arborescence structures, which represent the local influence regions of each node. We only consider influence propagated through these local arborescences, and we refer to this model as the *maximum influence arborescence* (MIA) model.

We show that the influence spread in the MIA model is submodular (i.e. having a diminishing marginal return property), and thus the simple greedy algorithm that selects one node in each round with the maximum marginal influence spread can guarantee an influence spread within $(1 - 1/e)$ of the optimal solution in the MIA model, while any higher ratio approximation is NP-hard. The greedy algorithm on

² An arborescence is a tree in a directed graph where all edges are either pointing toward the root (in-arborescence) or pointing away from the root (out-arborescence).

the MIA model is very efficient because (a) computation of the marginal influence spread on the arborescence structures can be done by efficient recursion; and (b) after selecting one seed with the largest influence spread, we only need to update local arborescence structures related to this seed for the selection of the next seed, and we further design a batch update scheme to speed up the update process.

We conduct extensive experiments on several real-world and synthetic networks of different scale and features, and under different types of the IC model. We compare our heuristic with both the greedy algorithm (Kempe et al. 2003; Leskovec et al. 2007; Chen et al. 2009) and several existing heuristics including the degree discount heuristics of Chen et al. (2009), the shortest-path based heuristics of Kimura and Saito (2006), and the popular PageRank algorithm (Brin and Page 1998) for ranking web pages. Our simulation results show that: (a) the greedy algorithm of Kempe et al. (2003); Leskovec et al. (2007); Chen et al. (2009) and the shortest-path based heuristic (Kimura and Saito 2006) have poor scalability: they take hours or days to select 50 seeds when the graph size reaches a few hundred thousand and become infeasible for larger sized graphs, while in the same range MIA heuristic can finish in seconds (more than three orders of magnitude speedup), and it continues to scale up beyond graphs with millions of edges, (b) comparing with the greedy algorithm and the shortest-path based heuristic in real graphs in which they are feasible to run, MIA heuristic has influence spread matching or being very close to that of the other two algorithms, (c) comparing with the rest heuristics, MIA algorithm is always among the best in influence spread, and in most cases it significantly outperforms the rest heuristics, with a margin as much as 100–260% increase in influence spread. Moreover, we show that by tuning the threshold θ , we can adjust the tradeoff between efficiency and effectiveness at different balance points on a spectrum.

To summarize, our main contribution is the design and evaluation of a scalable and tunable heuristic that handles the influence maximization problem for large-scale social networks. We demonstrate that our heuristic is currently the best one that could handle large-scale networks with more than a million edges, while even for moderate sized networks it is a very competitive alternative to much slower algorithms. The balanced efficiency and effectiveness of our heuristic make it suitable as a generic solution to influence maximization for many large-scale online social networks encountered in practice.

1.2 Related work

Domingos and Richardson (2001); Richardson and Domingos (2002) are the first to study influence maximization as an algorithmic problem. Their methods are probabilistic, however. Kempe et al. (2003) are the first to formulate the problem as a discrete optimization problem. Besides what we mentioned above already, they also study a number of other topics such as generalizations of influence cascade models and mixed marketing strategies in influence maximization. As pointed out, the main drawback of their work is the scalability of their greedy algorithms.

Several recent studies aimed at addressing this issue. In Leskovec et al. (2007), the authors present a “lazy-forward” optimization in selecting new seeds, which greatly reduces the number of evaluations on the influence spread of nodes and results in

as much as 700 times speedup demonstrated by their experimental results. However, even though the “lazy-forward” optimization is significant, it still takes hours to find 50 most influential nodes in a network with a few tens of thousands of nodes, as shown in [Chen et al. \(2009\)](#).

[Kimura and Saito \(2006\)](#) propose shortest-path based influence cascade models and provide efficient algorithms to compute influence spread under these models. The key differences between their work and ours are (a) instead of using maximum influence paths, they use simple shortest paths on the graph, which are not related to propagation probabilities, and (b) they do not utilize local structures such as our arborescences and thus in every round they need global computations to select the next seed. Therefore, their algorithms are not as efficient as ours.

This article is the continuation of [Chen et al. \(2009\)](#) in the pursuit of efficient and scalable influence maximization algorithms. In [Chen et al. \(2009\)](#), we explore two directions in improving the efficiency: one is to further improve the greedy algorithm of [Kempe et al. \(2003\)](#), and the other is to design new heuristic algorithms. The first direction shows improvement but is not significant enough, indicating that this direction could be difficult to continue. The second direction leads to new degree discount heuristics that are very efficient and generate reasonably good influence spread. The major issue is that the degree discount heuristics are derived from the *uniform* IC model where propagation probabilities on all edges are the same, which is rarely the case in reality. Our current work is a major step in overcoming this limitation—our new heuristic algorithm works for the general IC model while still maintains good balance between efficiency and effectiveness. We conduct many more experiments than in [Chen et al. \(2009\)](#) on more and larger scale graphs, and our results show that the MIA heuristic performs consistently better than the degree discount heuristic in all graphs.

This article is an extended version of our conference paper presented at SIGKDD2010 ([Chen et al. 2010a](#)). We add more detailed algorithm descriptions and proofs, give online error bound analysis and algorithms, and report more comprehensive experimental results, including additional datasets and algorithms, new independent cascade models and case study. Adding these contents is necessary to consolidate our methodology. First, the online error bound analysis provides a tighter lower bound on the optimality of our solution (76–82%) than the theoretical bound 63%, and thus should increase the confidence in the practical effectiveness of our algorithm. Second, the two additional datasets make the experiments more comprehensive: one was used in previous work and should be included for comprehensive comparison; the other is smaller but more familiar to data mining researchers, and is easier for case study of the selected seeds. Third, the newly introduced cascade models validate the robustness of our algorithm in different scenarios, e.g., when the overall influence probability in the network is higher, and when the influence between most pairs of nodes is highly asymmetric.

Several studies follow the design principle of the algorithm presented this article, and develop algorithms under other influence propagation assumptions. [Chen et al. \(2010b\)](#) study the linear threshold model and design an LDAG algorithm. It uses local directed acyclic graph (DAG) structures instead of local tree structures as used in our algorithm MIA. [Chen et al. \(2011\)](#) study an extension of IC model where negative opinions may emerge and propagate, and propose a MIA-N algorithm to solve it.

Another line of research aims to design data mining or machine learning algorithms to extract influence cascade model parameters from real datasets (Cha et al. 2009; Bakshy et al. 2009; Tang et al. 2009; Goyal et al. 2010; Rodriguez et al. 2010; Cui et al. 2011), which can be used to generate influence graphs studied in this article.

Paper organization. Section 2 provides preliminaries on the IC model and the greedy algorithm, and also points out that computing the exact influence spread given a seed set is #P-hard. Section 3 presents our MIA model and the algorithm for this model as well as its extension, the PMIA model. Sections 4 and 5 show our experimental results. We discuss future directions in Sect. 6.

2 IC model and greedy algorithm

We consider a directed graph $G = (V, E)$ with edge labels $pp : E \rightarrow [0, 1]$. For every edge $(u, v) \in E$, $pp(u, v)$ denotes the propagation probability of the edge, which is the probability that v is activated by u through the edge in the next step after u is activated.

Given a seed set $S \subseteq V$, the independent cascade (IC) model works as follows. Let $S_t \subseteq V$ be the set of nodes that are activated at step $t \geq 0$, with $S_0 = S$. At step $t + 1$, every node $u \in S_t$ may activate its out-neighbors $v \in V \setminus \cup_{0 \leq i \leq t} S_i$ with an independent probability of $pp(u, v)$. The process ends at a step t with $S_t = \emptyset$. Note that each activated node only has one chance to activate its out-neighbors at the step right after itself is activated, and each node stays as an activated node after it is activated. The *influence spread* of S , which is the expected number of activated nodes given seed set S , is denoted as $\sigma_I(S)$.

Given an input k , the influence maximization problem in the IC model is to find a subset $S^* \subseteq V$ such that $|S^*| = k$ and $\sigma_I(S^*) = \max\{\sigma_I(S) \mid |S| = k, S \subseteq V\}$. It is shown in Kempe et al. (2003) that this problem is NP-hard, but a constant-ratio approximation algorithm is available.

We say that a non-negative real valued function f on subsets of V is *submodular* if $f(S \cup \{v\}) - f(S) \geq f(T \cup \{v\}) - f(T)$, for all $v \in V$ and all pairs of subsets S and T with $S \subseteq T \subseteq V$. Intuitively, this means that f has diminishing marginal return. Moreover, we say that f is *monotone* if $f(S) \leq f(T)$ for all $S \subseteq T$. For any submodular and monotone function f with $f(\emptyset) = 0$, the problem of finding a set S of size k that maximizes $f(S)$ can be approximated by a simple greedy algorithm shown as Algorithm 1. The algorithm iteratively selects a new seed u that maximizes the incremental change of f into the seed set S until k seeds are selected. It is shown in Nemhauser et al. (1978) that the algorithm guarantees the approximation ratio

Algorithm 1 Greedy(k, f)

```

1: initialize  $S = \emptyset$ 
2: for  $i = 1$  to  $k$  do
3:   select  $u = \arg \max_{w \in V \setminus S} (f(S \cup \{w\}) - f(S))$ 
4:    $S = S \cup \{u\}$ 
5: end for
6: output  $S$ 

```

$f(S)/f(S^*) \geq 1 - 1/e$, where S is the output of the greedy algorithm and S^* is the optimal solution.

In [Kempe et al. \(2003\)](#), it is shown that function $\sigma_I(\cdot)$ is submodular and monotone with $\sigma_I(\emptyset) = 0$. Therefore, algorithm **Greedy**(k, σ_I) solves the influence maximization problem with an approximation ratio of $1 - 1/e$.

One important issue, however, is that there is no efficient way to compute $\sigma_I(S)$ given a set S . Although [Kempe et al. \(2003\)](#) claim that finding an efficient algorithm for computing $\sigma_I(S)$ is open, we point out that the computation is actually #P-hard, by showing a reduction from the counting problem of s - t connectness in a graph.

Theorem 1 *Computing the influence spread $\sigma_I(S)$ given a seed set S is #P-hard.*

Proof We prove the theorem by a reduction from the counting problem of s - t connectness in a directed graph ([Valiant 1979](#)). An instance of s - t connectness is a directed graph $G = (V, E)$ and two vertices s and t in the graph. The problem is to count the number of subgraphs of G in which s is connected to t . It is straightforward to see that this problem is equivalent to computing the probability that s is connected to t when each edge in G has an independent probability of $1/2$ to be connected, and another $1/2$ to be disconnected. We reduce this problem to the influence spread computation problem as follows. Let $\sigma_I(S, G)$ denote the influence spread in G given a seed set S . First, let $S = \{s\}$, and let $pp(e) = 1/2$ for all $e \in E$, and compute $I_1 = \sigma_I(S, G)$. Next, we add a new node t' and a directed edge from t to t' to G , obtaining a new graph G' , and let $pp(t, t') = 1$. Then we compute influence spread $I_2 = \sigma_I(S, G')$. Let $p(S, v, G)$ denote the probability that v is influenced by seed set S in G . It is easy to see that $I_2 = \sigma_I(S, G) + p(S, t, G) \cdot pp(t, t')$. Therefore, $I_2 - I_1$ is the probability that s is connected to t , and thus we solve the s - t connectness counting problem. It is shown in [Valiant \(1979\)](#) that s - t connectness is #P-complete, and thus the influence spread computation problem is #P-hard. □

The above theorem shows that computing exact influence spread is hard. Moreover, finding an efficient approximation algorithm for computing the probability of s - t connectivity is a long-standing open problem ([Vazirani 2004](#)). Together with the fact that several improvements ([Leskovec et al. 2007](#); [Chen et al. 2009](#)) of the original greedy algorithm of [Kempe et al. \(2003\)](#) are still not efficient, we believe that we need to look for alternative ways, such as heuristic algorithms, to tackle the efficiency problem in influence maximization.

3 MIA model and its algorithm

3.1 Basic MIA model and greedy algorithm

For a path $P = \langle u = p_1, p_2, \dots, p_m = v \rangle$, we define the *propagation probability* of the path, $pp(P)$, as

$$pp(P) = \prod_{i=1}^{m-1} pp(p_i, p_{i+1}).$$

Intuitively the probability that u activates v through path P is $pp(P)$, because it needs to activate all nodes along the path. To approximate the actual expected influence within the social network, we propose to use the *maximum influence path (MIP)* to estimate the influence from one node to another. Let $\mathcal{P}(G, u, v)$ denote the set of all paths from u to v in a graph G .

Definition 1 (Maximum influence path) For a graph G , we define the *maximum influence path* $MIP_G(u, v)$ from u to v in G as

$$MIP_G(u, v) = \arg \max_P \{pp(P) \mid P \in \mathcal{P}(G, u, v)\}.$$

Ties are broken in a predetermined and consistent way, such that $MIP_G(u, v)$ is always unique, and any subpath in $MIP_G(u, v)$ from x to y is also the $MIP_G(x, y)$. If $\mathcal{P}(G, u, v) = \emptyset$, we denote $MIP_G(u, v) = \emptyset$.

Note that for each edge (u, v) in the graph, if we translate the propagation probability $pp(u, v)$ to a distance weight $-\log pp(u, v)$ on the edge, then $MIP_G(u, v)$ is simply the shortest path from u to v in the weighted graph G . Therefore, the maximum influence paths and the later maximum influence arborescences directly correspond to shortest paths and shortest-path arborescences, and thus they permit efficient algorithms such as Dijkstra's algorithm to compute them.

For a given node v in the graph, we propose to use the *maximum influence in-arborescence (MIIA)*, which is the union of the maximum influence paths to v ,³ to estimate the influence to v from other nodes in the network. We use an *influence threshold* θ to eliminate MIPs that have too small propagation probabilities. Symmetrically, we also define *maximum influence out-arborescence (MIOA)* to estimate the influence of v to other nodes.

Definition 2 (MAXIMUM INFLUENCE IN(OUT)- ARBORESCENCE) For an influence threshold θ , the *maximum influence in-arborescence* of a node $v \in V$, $MIIA(v, \theta)$, is

$$MIIA(v, \theta) = \cup_{u \in V, pp(MIP_G(u, v)) \geq \theta} MIP_G(u, v).$$

The *maximum influence out-arborescence* $MIOA(v, \theta)$ is:

$$MIOA(v, \theta) = \cup_{u \in V, pp(MIP_G(v, u)) \geq \theta} MIP_G(v, u).$$

Intuitively, $MIIA(v, \theta)$ and $MIOA(v, \theta)$ give the local influence regions of v , and different values of θ control the size of these local influence regions.

Given a set of seeds S in G and the in-arborescence $MIIA(v, \theta)$ for some $v \notin S$, we approximate the IC model by assuming that the influence from S to v is only propagated through edges in $MIIA(v, \theta)$. With this approximation, we can calculate the probability that v is activated given S exactly. Let the *activation probability* of any node u in $MIIA(v, \theta)$, denoted as $ap(u, S, MIIA(v, \theta))$, be the probability that

³ Since we break ties in maximum influence paths consistently, the union of maximum influence paths to a node do not have undirected cycles, and thus it is indeed an arborescence.

Algorithm 2 $ap(u, S, MIIA(v, \theta))$

```

1: if  $u \in S$  then
2:    $ap(u) = 1$ 
3: else if  $N^{in}(u) = \emptyset$  then
4:    $ap(u) = 0$ 
5: else
6:    $ap(u) = 1 - \prod_{w \in N^{in}(u)} (1 - ap(w) \cdot pp(w, u))$ 
7: end if

```

u is activated when the seed set is S and influence is propagated in $IIA(v, \theta)$. Let $N^{in}(u, MIIA(v, \theta))$ be the set of in-neighbors of u in $IIA(v, \theta)$. In the above notations, $IIA(v, \theta)$ and S may be dropped when it is clear from the context. Then $ap(u, S, MIIA(v, \theta))$ can be computed recursively as given in Algorithm 2.

Line 6 in Algorithm 2 corresponds to the independent cascade assumption: every neighbor of a node u activates u independently. Thus the probability that u is not activated is equal to the product of the probability that u is not activated by each of its neighbors w , whose negation is recursively computed as the joint probability of two events: w is activated, and u is activated by w . The recursion closes at the leaves of the arborescence (line 4) and the seed nodes (line 2). Because $IIA(v, \theta)$ is an in-arborescence, there are no multiple paths between any pair of nodes in $IIA(v, \theta)$, and thus there is no dependency issue in the calculation of the activation probability and the calculation in Algorithm 2 exactly matches the IC model restricted onto $IIA(v, \theta)$.

In our MIA model we assume that seeds in S influence every individual node v in G through its $IIA(v, \theta)$. Let $\sigma_M(S)$ denote the influence spread of S in our MIA model, then we have

$$\sigma_M(S) = \sum_{v \in V} ap(v, S, MIIA(v, \theta)). \tag{3.1}$$

Even though activating multiple nodes from the same set of seeds in the MIA model are correlated events, Eq. 3.1 is still correct due to the linearity of the expectation over the sum of random variables.

We are interested in finding a set of seeds S of size k such that $\sigma_M(S)$ is maximized. It is not surprising that this optimization problem is NP-hard. In fact, the same reduction from set cover problem in Kempe et al. (2003) together with Theorem 5.3 of Feige (1998) is sufficient to show the following.

Theorem 2 *It is NP-hard to compute a set of nodes S of size k such that $\sigma_M(S)$ is maximized. Furthermore, it is NP-hard to approximate within a factor of $1 - 1/e + \epsilon$ for any $\epsilon > 0$.*

It is straight forward to verify the following result, which means we have an approximation algorithm.

Theorem 3 *Function σ_M is submodular and monotone and $\sigma_M(\emptyset) = 0$. Therefore, Greedy(k, σ_M) of Algorithm 1 achieves $1 - 1/e$ approximation ratio for the influence maximization problem in the basic MIA model.*

Note that the recursive computation of $ap(u)$ in Algorithm 2 can be transformed into an iterative form such that all $ap(u)$'s with u in $IIA(v, \theta)$ can be computed by

one traverse of the arborescence $MIIA(v, \theta)$ from leaves to the root. Thus, computing $\sigma_M(S)$ using Eq. 3.1 and Algorithm 2 is polynomial-time. Together with Algorithm 1, we already have a polynomial-time approximation algorithm. However, we could further improve the efficiency of the algorithm, as shown in the next section.

3.2 More efficient greedy algorithm

The only important step in the greedy algorithm is to select the next seed that gives the largest incremental influence spread. Consider the maximum influence in-arborescence $MIIA(v, \theta)$ of size t and a given seed set S . To select the next seed u , we need to compute the activation probability $ap(v, S \cup \{w\}, MIIA(v, \theta))$ for every $w \in MIIA(v, \theta)$, which takes $O(t^2)$ time if we simply use Algorithm 2 to compute every $ap(v, S \cup \{w\}, MIIA(v, \theta))$. We now show a batch update scheme such that we could compute $ap(v, S \cup \{w\}, MIIA(v, \theta))$'s for all $w \in MIIA(v, \theta)$ in $O(t)$ time.

To do so, we utilize the linear relationship between $ap(u)$ and $ap(v)$ in $MIIA(v, \theta)$, as shown by the following lemma, which is not difficult to derive from line 6 of Algorithm 2.

Lemma 1 (Influence linearity) *Consider $MIIA(v, \theta)$ and a node u in it. If we treat the activation probability $ap(u)$ as an independent variable, $ap(v)$ as a dependent variable, and other $ap(w)$'s as constants for all w 's not on the path from u to v in $MIIA(v, \theta)$, then $ap(v) = \alpha(v, u) \cdot ap(u) + \beta(v, u)$, where $\alpha(v, u)$, $\beta(v, u)$ are constants independent of $ap(u)$.*

Based on the recursive computation of $ap(u, S, MIIA(v, \theta))$ as shown in line 6 of Algorithm 2, it is straightforward to derive a recursive computation of $\alpha(v, u)$, as shown in Algorithm 3. To see the intuition behind the equations, we remind that $\alpha(v, u)$ is the increment of the activation probability of v caused by unit increment of u 's activation probability. Therefore, the boundary of the recursive computation is natural: when $u = v$, unit increment of u 's activation probability results in the same increment of v 's activation probability; when u 's out-neighbor w is a seed node, the increment of its activation probability does not induce any change of v 's activation

Algorithm 3 Compute $\alpha(v, u)$ with $MIIA(v, \theta)$ and S , after $ap(u, S, MIIA(v, \theta))$ for all u in $MIIA(v, \theta)$ are known.

```

1: /* the following is computed recursively */
2: if  $u = v$  then
3:    $\alpha(v, u) = 1$ 
4: else
5:   set  $w$  to be the out-neighbor of  $u$ 
6:   if  $w \in S$  then
7:      $\alpha(v, u) = 0$  /*  $u$ 's influence to  $v$  is blocked by seed  $w$  */
8:   else
9:      $\alpha(v, u) = \alpha(v, w) \cdot pp(u, w) \cdot \prod_{u' \in N^{in}(w) \setminus \{u\}} (1 - ap(u') \cdot pp(u', w))$ 
10:  end if
11: end if

```

probability because w is already activated. In other cases, the effect of u on v is determined by the effect of its out-neighbor w on v and the chance that w is activated by its in-neighbor u exclusively. The product over w 's in-neighbors except for u in line 9 corresponds to the probability that w is not activated by other in-neighbors. Note that Algorithm 3 can be transformed into an iterative form such that all $\alpha(v, u)$'s can be computed by one traverse of $MIIA(v, \theta)$ from the root to the leaves.

Computing the linear coefficients $\alpha(v, u)$ as defined in Lemma 1 is crucial in computing the incremental influence spread of a node u . Let us consider again the maximum influence in-arborescence $MIIA(v, \theta)$ of size t and a given seed set S . For any $w \in MIIA(v, \theta)$, if we select w as the next seed, its $ap(w)$ increases from the current value to 1. Since $ap(w)$ and $ap(v)$ has a linear relationship with the linear coefficient $\alpha(v, w)$, the incremental influence of w on v is given by $\alpha(v, w) \cdot (1 - ap(w))$. Therefore, we only need one pass of $MIIA(v, \theta)$ to compute $ap(w)$'s for all $w \in MIIA(v, \theta)$, and a second pass of $MIIA(v, \theta)$ to compute $\alpha(v, w)$'s and $\alpha(v, w) \cdot (1 - ap(w))$'s for all $w \in MIIA(v, \theta)$. This reduces the running time of computing incremental influence spread of all nodes in $MIIA(v, \theta)$ from $O(t^2)$ to $O(t)$.

Our complete greedy algorithm for the basic MIA model is presented in Algorithm 4. Lines (2–11) evaluate the incremental influence spread $IncInf(u)$ for any node u when the current seed set is empty. The evaluation is exactly as we described above using the linear coefficients $\alpha(v, u)$.

Lines (15–30) update the incremental influences whenever a new seed is selected in line 14. Suppose u is selected as the new seed in an iteration. The influence of u in the MIA model only reaches nodes in $MIOA(u, \theta)$. Thus the incremental influence spread $IncInf(w)$ for some w needs to be updated if and only if w is in $MIIA(v, \theta)$ for some $v \in MIOA(u, \theta)$. This means that the update process is relatively local to u . The update is done by first subtracting $\alpha(v, w) \cdot (1 - ap(w, S, MIIA(v, \theta)))$ before adding u into the seed set (line 19), and then adding u into the seed set (line 22), recomputing the $ap(w, S, MIIA(v, \theta))$ and $\alpha(v, w)$ under the new seed set (lines 24–25), and adding $\alpha(v, w) \cdot (1 - ap(w, S, MIIA(v, \theta)))$ into $IncInf(w)$ (line 28).

The algorithm is illustrated graphically in Fig. 1.

Time and space complexity. Let $n_{i\theta} = \max_{v \in V} \{|MIIA(v, \theta)|\}$ and $n_{o\theta} = \max_{v \in V} \{|MIOA(v, \theta)|\}$. Computing $MIIA(v, \theta)$ can be done using efficient implementations of Dijkstra's shortest-path algorithm. Assume the maximum running time to compute $MIIA(v, \theta)$ for any $v \in V$ is $t_{i\theta}$. When $MIIA(v, \theta)$'s for all node $v \in V$ are available, $MIOA(v, \theta)$'s can be derived from $MIIA(v, \theta)$'s, therefore no extra running time for $MIOA(v, \theta)$'s is needed. Notice that $n_{i\theta} = O(t_{i\theta})$.

For every node $v \in V$, our algorithm stores $MIIA(v, \theta)$, $MIOA(v, \theta)$, and for every $u \in MIIA(v, \theta)$, $ap(u, S, MIIA(v, \theta))$ and $\alpha(v, u)$ are stored (note that $ap(u, S, MIIA(v, \theta))$ can reuse the same entry for different seed set S). We also use a max-heap to store and update $IncInf(v)$ for all $v \in V$. Therefore, the space complexity of the algorithm is $O(n(n_{i\theta} + n_{o\theta}))$.

During the initialization of Algorithm 4, it takes $O(nt_{i\theta})$ time to compute $MIIA(v, \theta)$ for all $v \in V$, $O(nn_{i\theta})$ time to compute all $\alpha(v, u)$'s and $IncInf(u)$'s, and $O(n)$ time to initialize the max-heap for storing $IncInf(u)$'s. Therefore, the total running time for initialization is $O(nt_{i\theta})$. During one iteration of the main loop, it takes constant time to select the new seed from the max-heap, $O(n_{o\theta}n_{i\theta} \log n)$

Algorithm 4 $MIA(G, k, \theta)$

```

1: /* initialization */
2: set  $S = \emptyset$ 
3: set  $IncInf(v) = 0$  for each node  $v \in V$ 
4: for each node  $v \in V$  do
5:   compute  $MIA(v, \theta)$  and  $MIOA(v, \theta)$ 
6:   set  $ap(u, S, MIA(v, \theta)) = 0, \forall u \in MIA(v, \theta)$  /* since  $S = \emptyset$  */
7:   compute  $\alpha(v, u), \forall u \in MIA(v, \theta)$  (Algorithm 3)
8:   for each node  $u \in MIA(v, \theta)$  do
9:      $IncInf(u) += \alpha(v, u) \cdot (1 - ap(u, S, MIA(v, \theta)))$ 
10:  end for
11: end for
12: /* main loop */
13: for  $i = 1$  to  $k$  do
14:   pick  $u = \arg \max_{v \in V \setminus S} \{IncInf(v)\}$ 
15:   /* update incremental influence spreads */
16:   for  $v \in MIOA(u, \theta) \setminus S$  do
17:     /* subtract previous incremental influence */
18:     for  $w \in MIA(v, \theta) \setminus S$  do
19:        $IncInf(w) -= \alpha(v, w) \cdot (1 - ap(w, S, MIA(v, \theta)))$ 
20:     end for
21:   end for
22:    $S = S \cup \{u\}$ 
23:   for  $v \in MIOA(u, \theta) \setminus S$  do
24:     compute  $ap(w, S, MIA(v, \theta)), \forall w \in MIA(v, \theta)$  (Algo. 2)
25:     compute  $\alpha(v, w), \forall w \in MIA(v, \theta)$  (Algo. 3)
26:     /* add new incremental influence */
27:     for  $w \in MIA(v, \theta) \setminus S$  do
28:        $IncInf(w) += \alpha(v, w) \cdot (1 - ap(w, S, MIA(v, \theta)))$ 
29:     end for
30:   end for
31: end for
32: return  $S$ 

```

time to update $IncInf(w)$'s on the max-heap, and $O(n_{o\theta}n_{i\theta})$ time to compute $ap(w, S, MIA(v, \theta, S))$'s and $\alpha(v, w)$'s after selecting the new seed. Thus, one iteration of the main loop takes $O(n_{o\theta}n_{i\theta} \log n)$ time. Together, the total running time of the algorithm is $O(nt_{i\theta} + kn_{o\theta}n_{i\theta} \log n)$. Note that without applying the improvement of utilizing the linear relationship, the time complexity would be $O(nt_{i\theta} + kn_{o\theta}n_{i\theta}(n_{i\theta} + \log n))$.

Therefore, the algorithm performs the fastest when $n_{i\theta}$, $n_{o\theta}$, and $t_{i\theta}$ are significantly smaller than n , that is, when the arborescences are small. This typically occurs for a reasonable range of θ values, when the graph is sparse and the propagation probabilities on edges are small, which is usually the case for social networks. Our experiments in Sects. 4 and 5 will demonstrate the efficiency of our algorithm.

3.3 Prefix excluding MIA model

In the basic MIA model, we only consider the maximum influence path from u to v for influence propagation. Consider the scenario of two seeds s_1 and s_2 such that $MIP_G(s_2, v) \subset MIP_G(s_1, v)$. The probability that v is activated in the basic MIA




Fig. 1 Illustration of MIA algorithm. **a** In the beginning MIAs are computed (the figure shows one MIA rooted at v) and the accumulated influence score of every node is recorded. **b** The node with the largest influence score is selected as a seed. **c** Once a seed u is selected, all the MIAs rooted at a node in u 's MIOA should be updated. The figure shows one such MIA rooted at v , and the influence score of all the other nodes $\{x_1, x_2, x_3, x_4, x_5, v\}$ in this MIA should be discounted. Dotted edges are the edges not used in corresponding computation

model is only determined by s_2 and is not affected by s_1 , or we can say that the influence of s_1 to v is blocked by s_2 in the middle.

To achieve a better approximation to the IC model, we prefer a MIA model in which the influence of a seed is not blocked by other seeds. A natural way to extend the basic MIA model is considering maximum influence paths avoiding other seeds. Let $S = \{s_1, s_2, \dots, s_m\}$ and $S^i = S \setminus \{s_i\}$. We define $G(S^i)$ be the subgraph of G induced by $V \setminus S^i$. Then, for each seed s_i and node $v \in V \setminus S$, we use the maximum influence path $MIP_{G(S^i)}(s_i, v)$ to estimate the influence from s_i to v . In other words, we consider maximum influence paths avoiding other seeds in calculating the influence spread.

The generic Algorithm 1 also works in this model. However, it is not clear how to implement it efficiently similar to the approach in Algorithm 4. In this section, we consider a variant of the above extension that allows an efficient greedy algorithm. We call this extension the *prefix excluding MIA* (PMIA) model.

Intuitively, in the PMIA model, the seeds have an order (as the order by which they are selected by the greedy algorithm). For any given seed s , its maximum influence paths to other nodes should avoid all seeds in the prefix before s . The major technical difference is the definition of the *maximum influence in(out)-arborescence* for the PMIA model, especially if we want to design an efficient greedy algorithm in the framework of Algorithm 4.

Let $S = \langle s_1, s_2, \dots, s_m \rangle$ be a sequence of seeds. Define $S_i = \langle s_1, s_2, \dots, s_{i-1} \rangle$ and $S_1 = \emptyset$. Let $G(S')$ be the subgraph of G induced by $V \setminus S'$ for any sequence S' . We first define *ineffective seeds* with respect to a node v , which are those seeds whose influence to v are blocked by some other subsequent seeds in sequence S .

Definition 3 (Ineffective seeds) For a given node $v \in V \setminus S$, we define the set of *ineffective seeds* for v as:

$$IS(v, S) = \{s_i \in S \mid \exists j > i, s.t., s_j \in MIP_{G(S_i)}(s_i, v)\}.$$

Now consider the maximum influence in-arborescence (MIIA) of a node v in the PMIA model. First, for the maximum influence path from a seed s_i to v , it should be defined as $MIP_{G(S_i)}(s_i, v)$ to avoid seeds in its prefix. Second, for the case where the MIP from seed s_i to v is blocked by a subsequent seed s_j , we need to give a special treatment in order to use the influence linearity of Lemma 1 for an efficient computation of incremental influence spread. Consider a node $u \notin S$ located on the MIP from s_i to s_j . If u is selected as a seed later, then its MIP to v should avoid all seeds in S , and thus to compute its incremental influence spread correctly using the linearity property, we need to compute the MIP from u to v in the graph $G(S)$. Moreover, we need to remove the ineffective seed s_i and its MIP to v because otherwise s_i would have two different paths to v , violating the arborescence definition.

For out-arborescence from $v \notin S$, we need to consider all MIPs from v that avoid all seeds in S . This is because we only need to compute the out-arborescence of a node v when v is just selected as the next seed. In this case, the paths in the above computed out-arborescence of v match the paths in the corresponding in-arborescences used to compute the incremental influence of v (since those paths avoid all seeds already in S). Therefore, we have the following formal definitions.

Definition 4 (MIIA(MIOA) for the PMIA Model) The *maximum influence in-arborescence* of v in the PMIA model for $v \notin S$ is:

$$\begin{aligned} PMIA(v, \theta, S) = & \\ & (\cup\{MIP_{G(S_i)}(s_i, v) \mid s_i \in S \setminus IS(v, S), \\ & \quad pp(MIP_{G(S_i)}(s_i, v)) \geq \theta\}) \\ & \cup(\cup\{MIP_{G(S)}(u, v) \mid u \in V \setminus S, \\ & \quad pp(MIP_{G(S)}(u, v)) \geq \theta\}). \end{aligned}$$

The *maximum influence out-arborescence* of v in the PMIA model for $v \notin S$ is:

$$\begin{aligned} PMIOA(v, \theta, S) = & \cup\{MIP_{G(S)}(v, u) \mid u \in V \setminus S, \\ & \quad pp(MIP_{G(S)}(v, u)) \geq \theta\}. \end{aligned}$$

Given the above definition, we can have activation probabilities $ap(u, S, PMIA(v, \theta, S))$ computed by Algorithm 2. Then, similar to Eq. 3.1, we can define $\sigma_P(S)$ as the influence spread given a seed sequence S , which is computed using the following equation:

$$\sigma_P(S) = \sum_{v \in V} ap(v, S, PMIA(v, \theta, S)). \quad (3.2)$$

Notice that different sequences S of the same set of seeds may generate different values of $\sigma_P(S)$. Therefore, the submodularity defined on set functions previously does not apply to σ_P . Fortunately, we can define *sequence submodularity* in a similar way, which also leads to the greedy algorithm with an approximation ratio of $1 - 1/e$.

Sequence submodularity. We now define sequence submodularity, which is implicitly used by Streeter and Golovin in Streeter and Golovin (2007). Let \mathcal{S} be the set of all sequences of V , including the empty sequence \emptyset . Let \oplus be the binary operator that concatenates two sequences into one. We say that a non-negative function f defined on \mathcal{S} is *sequence submodular* if $f(S_1 \oplus S_2 \oplus \{t\}) - f(S_1 \oplus S_2) \leq f(S_1 \oplus \{t\}) - f(S_1)$ for all sequences $S_1, S_2 \in \mathcal{S}$. Moreover, f is *prefix monotone* if $f(S_1) \leq f(S_2 \oplus S_1)$ for all $S_1, S_2 \in \mathcal{S}$. An important result that matches the one for set submodular functions is that if f is sequence submodular and prefix monotone and $f(\emptyset) = 0$, then the greedy algorithm of Algorithm 1 (with set union \cup replaced by sequence concatenation \oplus) finds a sequence S within $1 - 1/e$ of the optimal S^* . Since the original proof in Streeter and Golovin (2007) is presented in a different context, we rephrase the proof below.

Theorem 4 (Theorem 3 in Streeter and Golovin (2007)) *Let f be a sequence submodular, prefix monotone function with $f(\emptyset) = 0$. Define $S_0 = \emptyset$ and for $1 \leq i \leq k$, let $s_i = \arg \max_{s \in V} \{f(S_{i-1} \oplus \{s\})\}$ and $S_i = S_{i-1} \oplus \{s_i\}$. Let $S^* = \arg \max_{S' \in \mathcal{S}} \{f(S') \mid S' \in \mathcal{S} \text{ and } |S'| = k\}$. We have*

$$f(S_k) \geq (1 - 1/e) \cdot f(S^*).$$

Proof Let $\Delta_i = f(S^*) - f(S_i)$. By prefix monotonicity, we have $f(S^*) \leq f(S_i \oplus S^*)$. Let $S^* = \langle s_1^*, \dots, s_k^* \rangle$, and $S_i^* = \langle s_1^*, \dots, s_i^* \rangle$. By submodularity, for $1 \leq i \leq k$, we have

$$\begin{aligned} f(S_i \oplus S^*) &= f(S_i \oplus S_{k-1}^* \oplus \langle s_k^* \rangle) \\ &\leq f(S_i \oplus S_{k-1}^*) + f(S_i \oplus \langle s_k^* \rangle) - f(S_i) \\ &\leq f(S_i \oplus S_{k-1}^*) + f(S_{i+1}) - f(S_i), \end{aligned}$$

where the last inequality is due to the definition of S_{i+1} . Repeating the above derivation for k times, we have

$$\begin{aligned} f(S^*) \leq f(S_i \oplus S^*) &\leq f(S_i) + k \cdot (f(S_{i+1}) - f(S_i)) \\ &= f(S_i) + k \cdot (\Delta_i - \Delta_{i+1}). \end{aligned}$$

Therefore, $\Delta_i \leq k \cdot (\Delta_i - \Delta_{i+1})$ and $\Delta_{i+1} \leq (1 - \frac{1}{k})\Delta_i$. Hence

$$f(S^*) - f(S_k) = \Delta_k \leq (1 - \frac{1}{k})^k \Delta_0 \leq f(S^*)/e.$$

□

It is not difficult to verify the following result on σ_P , which means that the greedy algorithm works as an approximation algorithm.

Theorem 5 *Function σ_P is sequence submodular and prefix monotone and $\sigma_P(\emptyset) = 0$. Therefore, Greedy(k, σ_P) of Algorithm 1 (with set union \cup replaced by sequence*

concatenation \oplus) achieves $1 - 1/e$ approximation ratio for the influence maximization problem in the PMIA model.

Algorithm in the PMIA model. We now present the necessary changes needed to adapt Algorithm 4 to the PMIA model. The major issue is the computation of $PMIA(v, \theta, S)$ and $PMIOA(v, \theta, S)$. The computation of $PMIOA(v, \theta, S)$ is relatively simple, since we only need to remove S from the graph. Therefore, we can use Dijkstra's algorithm on graph $G(S)$ to compute $PMIOA(v, \theta, S)$.

To efficiently compute $PMIA(v, \theta, S)$, we maintain the set of ineffective seeds $IS(v, S)$ for each node $v \in V \setminus S$. Given $IS(v, S)$, $PMIA(v, \theta, S)$ can be calculated as follows. We start Dijkstra's algorithm from v traversing inward edges. Whenever the algorithm hits a seed node s , it stops this branch and does not go further onto the in-neighbors of s . After this variant of Dijkstra's algorithm completes, we remove all nodes $IS(v, S)$ from the computed in-arborescence.

When a new seed u is selected, we have to update $IS(v, S)$ for all nodes v in $PMIOA(u, \theta, S)$. This can be done by checking the set of *effective seeds* (those in $S \setminus IS(v, S)$) that are blocked by u in $PMIA(v, \theta, S)$. For completeness, we present Algorithm 5 for the efficient greedy algorithm in the PMIA model. Algorithm 5 essentially follows Algorithm 4, with all MIA 's and $MIOA$'s being replaced by $PMIA$'s and $PMIOA$'s, and these $PMIA$'s and $PMIOA$'s being recomputed whenever the seed set changes (lines 16 and 26).

3.4 Online bounds from sequence submodularity

We have proved the approximation guarantee of $(1 - 1/e)$. This approximation bound is *offline*, which can be stated before running the actual algorithm. When we have selected a set of nodes as seeds, we can estimate a tighter bound for that solution thanks to the sequence submodularity. This online bound is irrelevant with the algorithm to use, but relevant with the solution.

Theorem 6 *Let f be a sequence submodular, prefix monotone function with $f(\emptyset) = 0$. Let S be an arbitrary set of k nodes and S^* be the optimal set of k nodes, i.e., $S^* = \arg \max_{|S|=k} f(S) = \langle s_1^*, \dots, s_k^* \rangle$. For any node v not in S , define $d_v = f(S \oplus \langle v \rangle) - f(S)$. Sort all the nodes v not in S in the decreasing order of d_v , so that $d_{v_1} \geq d_{v_2} \geq \dots$. We have the online bound*

$$f(S^*) \leq f(S) + \sum_{i=1}^k d_{v_i}.$$

Proof Let $S_i^* = \langle s_1^*, \dots, s_i^* \rangle$. By prefix monotonicity, we have $f(S^*) \leq f(S \oplus S^*)$. By sequence submodularity, for $1 \leq i \leq k$, we have

$$\begin{aligned} f(S \oplus S_i^*) &= f(S \oplus S_{i-1}^* \oplus \langle s_i^* \rangle) \\ &\leq f(S \oplus S_{i-1}^*) + f(S \oplus \langle s_i^* \rangle) - f(S) \end{aligned}$$

Sum the above equations with i from 1 to k , we have

Algorithm 5 $PMIA(G, k, \theta)$

```

1: /* initialization */
2: set  $S = \emptyset$ 
3: set  $IncInf(v) = 0$  for each node  $v \in V$ 
4: for each node  $v \in V$  do
5:   compute  $PMIA(v, \theta, S)$ 
6:   set  $ap(u, S, PMIA(v, \theta, S)) = 0, \forall u \in PMIA(v, \theta, S)$  /* since  $S = \emptyset$  */
7:   compute  $\alpha(v, u), \forall u \in PMIA(v, \theta, S)$  (Algorithm 3)
8:   for each node  $u \in PMIA(v, \theta, S)$  do
9:      $IncInf(u) += \alpha(v, u) \cdot (1 - ap(u, S, PMIA(v, \theta, S)))$ 
10:  end for
11: end for
12: /* main loop */
13: for  $i = 1$  to  $k$  do
14:   pick  $u = \arg \max_{v \in V \setminus S} \{IncInf(v)\}$ 
15:   /* update incremental influence spreads */
16:   compute  $PMIOA(u, \theta, S)$ 
17:   for  $v \in PMIOA(u, \theta, S)$  do
18:     /* subtract previous incremental influence */
19:     for  $w \in PMIA(v, \theta, S) \setminus S$  do
20:        $IncInf(w) -= \alpha(v, w) \cdot (1 - ap(w, S, PMIA(v, \theta, S)))$ 
21:     end for
22:   end for
23:    $S = S \cup \{u\}$ 
24:   /* the following  $PMIOA(u, \theta, S \setminus \{u\})$  is the same as computed in line 16 */
25:   for  $v \in PMIOA(u, \theta, S \setminus \{u\}) \setminus \{u\}$  do
26:     compute  $PMIA(v, \theta, S)$ 
27:     compute  $ap(w, S, PMIA(v, \theta, S)), \forall w \in PMIA(v, \theta, S)$  (Algo. 2)
28:     compute  $\alpha(v, w), \forall w \in PMIA(v, \theta, S)$  (Algo. 3)
29:     /* add new incremental influence */
30:     for  $w \in PMIA(v, \theta, S) \setminus S$  do
31:        $IncInf(w) += \alpha(v, w) \cdot (1 - ap(w, S, PMIA(v, \theta, S)))$ 
32:     end for
33:   end for
34: end for
35: return  $S$ 

```

$$\begin{aligned}
 f(S^*) &\leq f(S \oplus S^*) \leq f(S) + \sum_{i=1}^k (f(S \oplus \{s_i^*\}) - f(S)) \\
 &\leq f(S) + \sum_{i=1}^k (d_{v_i}).
 \end{aligned}$$

The last inequality is because v_1, v_2, \dots, v_k are the k nodes with the largest marginal influence given S . □

Theorem 6 establishes a way of computing how far any given solution S is from the optimal solution, no matter it is obtained by which algorithm. We can turn this theorem into an algorithm, as formalized in Algorithm 6.

In Leskovec et al. (2007), a technique of lazy evaluations is used to improve the Greedy algorithm: when finding the next seed, we mark the marginal influence d_v of all nodes v as invalid, and go through the nodes in decreasing order of their marginal

Algorithm 6 OnlineBound(f, S)

```

1: for  $k = 1$  to  $|S|$  do
2:   initialize  $U = \emptyset, b_k = f(S_k) / * S_k$  as defined in Theorem 4 */
3:   for  $i = 1$  to  $k$  do
4:     select  $u = \arg \max_{v \in V \setminus S_k \setminus U} (f(S_k \oplus \langle v \rangle) - f(S_k))$ 
5:      $U = U \cup \{u\}$ 
6:      $b_k = b_k + f(S_k \oplus \langle u \rangle) - f(S_k)$ 
7:   end for
8: end for
9: output  $\{b_k\}$ 

```

Algorithm 7 Lazy_OnlineBound(k, f, S)

```

1: initialize  $d_v$  with all 0's
2: initialize a max-heap  $Q = \{v_n\}$  in the decreasing order of  $d_v$ 
3: for  $k = 1$  to  $|S|$  do
4:   initialize  $U = \emptyset, b_k = f(S_k), valid_v = false \forall v / * S_k$  as defined in Theorem 4 */
5:   for  $i = 1$  to  $k$  do
6:     while  $valid_{v_0} = false$  for the top element  $v_0 \in Q$  do
7:        $v = v_0$ 
8:       remove  $v_0$  from  $Q$ 
9:       compute  $d_v = (f(S_k \oplus \langle v \rangle) - f(S_k))$ 
10:      set  $valid_v = true$ , and reinsert  $v$  into  $Q$ 
11:     end while
12:      $u = v_0$ 
13:     remove  $v_0$  from  $Q$ 
14:      $U = U \cup \{u\}$ 
15:      $b_k = b_k + f(S_k \oplus \langle u \rangle) - f(S_k)$ 
16:   end for
17:   for  $u \in U$  do
18:     reinsert  $u$  into  $Q$ 
19:   end for
20: end for
21: output  $\{b_k\}$ 

```

influence d_v computed in last round. If d_v for the top node v is invalid, we recompute it, and insert it into the existing order of d_v 's, and mark it as valid. Thus all d_v 's can be stored in a priority queue, while some of them are valid and some invalid. We repeat the computation for the top node v in the priority queue and the insertion, until the top element has a valid d_v . The correctness of this lazy procedure follows directly from submodularity, and leads to far fewer evaluations of d_v 's because we do not need all d_v 's to become valid to obtain the largest d_v .

Similar techniques can be used in our online bounds evaluation algorithm. We can still maintain a max-heap and keep recomputation for the top element and insertion in each round. Instead of stopping once the top element is valid, we change the stopping criteria as the top k element is valid. For example, this can be achieved as in Algorithm 7.

4 Primary experiment

We conduct experiments with our algorithm as well as a number of other algorithms on several real-world and synthetic networks. This section contains primary results illustrating the performance of our algorithm from the following aspects: (a) its scalability comparing to other algorithms; (b) its influence spread comparing to other algorithms; and (c) the tuning of its control parameter θ .

4.1 Experiment setup

Datasets. We use four real-world networks and a synthetic dataset. The first one, denoted NetHEPT, is the same as used in [Chen et al. \(2009\)](#). It is an academic collaboration network extracted from “High Energy Physics-Theory” section of the e-print arXiv (<http://www.arXiv.org>), with nodes representing authors and edges representing coauthorship relations. The second is a much larger collaboration network, the DBLP Computer Science Bibliography Database maintained by Michael Ley. The other two datasets are published network data by Jure Leskovec. One is a Who-trust-whom network of Epinions.com⁴, where nodes are members of the site and a directed edge from u to v means v trust u (and thus u has influence to v). Another is the Amazon product co-purchasing network⁵ dated on March 2, 2003, where nodes are products and a directed edge from u to v means product v is often purchased with product u (and thus u has influence to v).⁶ We refer to these two datasets as Epinions and Amazon. We choose these networks since they cover a variety of networks with sizes ranging from 30K edges to 2M edges. Some basic statistics about these networks are given in Table 1 (Epinions and Amazon networks are treated as undirected graphs in the statistics). Finally, in the scalability test, we use the DIGG package available on the web (<http://digg.cs.tufts.edu>) to randomly generate power-law graphs of different sizes based on the model of [Aiello et al. \(2000\)](#).

Generating propagation probabilities. Since our algorithm is targeted at the general IC model with nonuniform propagation probabilities, we use the following two models to generate these nonuniform probabilities.

- WC model: This is the *weighted cascade* model proposed in [Kempe et al. \(2003\)](#). In this model, $pp(u, v)$ for an edge (u, v) is $1/deg(v)$, where $deg(v)$ is the in-degree of v . Thus even if the original graph is undirected, the model will generate asymmetric and nonuniform propagation probabilities.
- TRIVALENCY model: On every edge (u, v) , we uniformly at random select a probability from the set $\{0.1, 0.01, 0.001\}$, which corresponds to high, medium, and low influences.

⁴ <http://snap.stanford.edu/data/soc-Epinions1.html>.

⁵ <http://snap.stanford.edu/data/amazon0302.html>.

⁶ Although the Amazon dataset is for products, we still include it in our experiments to test a variant of social network. Moreover, it also makes sense to find top seed products that lead to the most co-purchasing behaviors.

Table 1 Statistics of four tested real-world networks

Dataset	NetHEPT	DBLP	Epinions	Amazon
#Node	15K	655K	76K	262K
#Edge	31K	2.0M	509K	1.2M
Average degree	4.12	6.1	13.4	9.4
Maximal degree	64	588	3079	425
#Connected Component	1781	73K	11	1
Largest component size	6794	517K	76K	262K
Average component size	8.6	9.0	6.9K	262K

Algorithms. We compare our MIA heuristic with both the greedy algorithm and several heuristics that appear in the literature. The following is a list of algorithms we evaluate in our experiments.

- **PMIA(θ):** Our Algorithm 4 for the PMIA model with influence threshold θ . The value of θ for a particular dataset is selected using the heuristic discussed in the “tuning of parameter θ ” part of Section 4.2.
- **Greedy:** The original greedy algorithm on the IC model (Kempe et al. 2003) with the lazy-forward optimization of Leskovec et al. (2007). For each candidate seed set S , 20000 simulations is run to obtain an accurate estimate of $\sigma_I(S)$.
- **DegreeDiscountIC:** The degree discount heuristic of Chen et al. (2009) developed for the uniform IC model with a propagation probability of $p = 0.01$, same as used in Chen et al. (2009).
- **SP1M:** The shortest-path based heuristic algorithm of Kimura and Saito (2006), also enhanced with the lazy-forward optimization of Leskovec et al. (2007).
- **PageRank:** The popular algorithm used for ranking web pages (Brin and Page 1998). Here the transition probability along edge (u, v) is $pp(v, u)/\rho_u$, where ρ_u is the sum of propagation probabilities on all incoming edges of u . Note that in the PageRank algorithm the transition probability of (u, v) indicates u 's “vote” to v 's ranking, and thus if $pp(v, u)$ is higher, v is more influential to u and thus u should vote v higher. We use 0.15 as the restart probability for PageRank, and we use the power method to compute the PageRank values. The stopping criteria is when two consecutive iterations differ for at most 10^{-4} in L_1 norm.
- **Random:** As a baseline comparison, simply select k random vertices in the graph.

We ignore other centrality measures, such as distance centrality and betweenness centrality (Freeman 1979) as heuristics, since we have shown in Chen et al. (2009) that distance centrality is very slow and has very poor influence spread, while betweenness centrality would be much slower than distance centrality.

To obtain the influence spread of the heuristic algorithms, for each seed set, we run the simulation on the networks 20000 times and take the average of the influence spread, which matches the accuracy of the greedy algorithms. The reported running time for PMIA includes the time for computing the bounded shortest path to build the arborescence. The experiments are run on a server with 2.33GHz Quad-Core Intel Xeon E5410 and 32G memory.

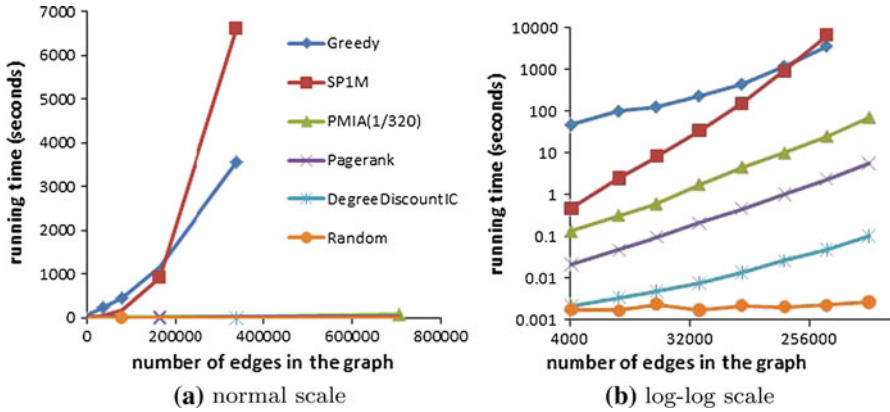


Fig. 2 Scalability of different algorithms in synthetic datasets. Each data point is an average of ten runs

4.2 Experiment results

Scalability on the synthetic dataset. To test scalability, we generate a family of graphs of increasing sizes using the DIGG package (<http://digg.cs.tufts.edu>), which applies the random power-law graph model of Aiello et al. (2000) to generate random graphs. We use graphs of doubling sizes— $2K$, $4K$, $8K$, . . . , up to $256K$ in the number of nodes, and a power-law exponent of 2.16. The average degree of these graphs is between 2 and 3 for these graphs, which is lower than the real networks in Table 1. We use the WC model for the graphs, and run PMIA algorithm with a fixed $\theta = 1/320$, as well as other algorithms, to find 50 seeds in every graph. The result is shown in Fig. 2, with normal scale shown in (a) and log-log scale of the same figure shown in (b) to differentiate these algorithms better.

The result in Fig 2 (a) clearly separate all algorithms into two groups. Algorithms Greedy and SP1M are not scalable: their running times are in the hour range with around $400K$ edge graphs and it becomes infeasible to run them in larger graphs since we want to take average of 10 runs of every algorithm. Note that we already choose graphs with low average degree so that they could run faster. Later reports on real graphs will show that they run even slower on those graphs. Our PMIA along with the rest heuristics can all scale up quite well. Figure 2 (b) differentiates the algorithms further. SP1M has the worst slope and is certainly not feasible for large-scale graphs. Greedy has the similar slope as other algorithms but its intercept is too large, because its Monte-Carlo simulation based estimation of incremental influence spread for every node is too slow. Our PMIA has both good slope and intercept, making it easily scalable to large graphs with millions of edges.

Influence spread and running time for the real-world datasets. We run tests on the four datasets and the two IC models to obtain influence spread results. The seed set size k ranges from 1 to 50. For ease of reading, in all influence spread figures (best viewed in color), the legend ranks the algorithms top-down in the same order as the influence spreads of the algorithms when $k = 50$. Moreover, if two curves are too close to each other, we group them together and show properly in the legend. All percentage difference reported below on influence spreads are the average of

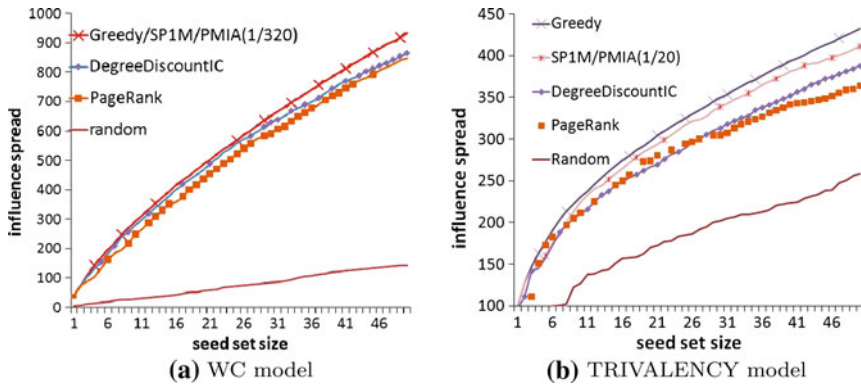


Fig. 3 Influence spread results on the NetHEPT dataset




Fig. 4 Influence spread results on the Epinions dataset

percentage differences from selecting one seed to selecting 50 seeds. Taking average is reasonable, since some algorithms may behave better when selecting the first few seeds while other algorithms behave better when selecting more seeds. The running time results are the time for selecting 50 seeds.

Figures 3, 4, 5 and 6 show the results on influence spreads for the four datasets on two IC models, while Fig. 7 shows the running time results of the four datasets on the WC model (results on the TRIVALENCY model are similar and omitted).

For the moderate sized graph NetHEPT where Greedy is still feasible to run, the influence results in Fig. 3 show that Greedy produces the best influence spread, but PMIA is very close to Greedy: its influence spread essentially matches that of Greedy for the WC model and is only 3.8% less than Greedy for the TRIVALENCY model. Comparing with other heuristics, PMIA performs quite well: it matches the influence spread of SP1M while outforms the rest heuristics in both models—in the WC model, PMIA is 3.9% and 11.4% better, while in the TRIVALENCY model, PMIA is 6.5% and 15.4% better, comparing to DegreeDiscountIC and PageRank respectively. Random has a much worse influence spread, indicating that a careful seed selection




Fig. 5 Influence spread results on the Amazon dataset




Fig. 6 Influence spread results on the DBLP dataset

is indeed important to effective viral marketing results. When looking at the running time in Fig. 7 for NetHEPT on WC, we clearly see that Greedy is already quite slow (1.3 hours), while PMIA only takes 1 second, more than three orders of magnitude better. PMIA is also more than one order of magnitude faster than SP1M, and is comparable with PageRank. DegreeDiscountC is the best in running time, because it is simple and specially tuned for the uniform IC model.

Figure 4 shows the result on the Epinions dataset, a large network with half a million edges. The graph is already too large for Greedy to run, so Greedy is out of the picture. For the WC model, PMIA still matches the influence spread of SPIM while it has a large winning margin over DegreeDiscountC and PageRank—PMIA is 96% and 115% better than DegreeDiscountC and PageRank, respectively. This demonstrates that DegreeDiscountC and PageRank are rather unstable heuristics while PMIA is very consistent in influence performance. For the TRIVALENCY model, we see that all heuristics, including Random reach a high level of influence spread after only a few seeds, while afterwards the increase in influence spread is slow. This behavior is quite different from the behavior of other test results we have seen so far, but it is very similar to a result presented in Kempe et al. (2003) for a graph when every edge





Fig. 7 Running time of different algorithms in four datasets

has a propagation probability of 0.1. Therefore, we believe that the explanation is also similar: in this test, after deleting the edges based on their propagation probabilities and only keeping the edges that will propagate influence, the resulting graph is likely to have a relatively large strongly connected component, and thus even random node selection would likely to hit this component after a few attempts, drastically increasing the influence spread. However, afterwards, additional seeds could only reach a small portion of still unaffected nodes, so further improvement in influence spread is small. But even in this case PMIA is still the best, outperforming the rest heuristics. For running time, we see that PMIA only takes 10s but SP1M now takes 2.1 h, more than 700 times slower than PMIA.

Next, for the one-million-edge graph Amazon, Fig. 5 shows that in the WC model PMIA again outperforms PageRank and DegreeDiscountIC with a large margin (99% and 266%, respectively), and in the TRIVALENCY model, it even outperforms SP1M significantly (14.1%, 23.9%, and 41.7% better than SP1M, PageRank, DegreeDiscountIC, respectively). Two unique features for this dataset are: (a) the influence spread is rather small, e.g., in TRIVALENCY, 50 seeds only generate a spread of around 80 nodes, and (b) the increase in influence spread is almost linear. The two features have the same reason— influence is very local and cannot propagate very far. It is probably because Amazon is a product co-purchasing network, not a social network. For running time, we now see that SP1M takes 30h, reaching its feasibility limit, while PMIA still only takes 10s, showing its superb scalability over SP1M.

Finally, for the two million edge DBLP dataset, Fig. 6 shows that this time PageRank and DegreeDiscountIC match PMIA and are slightly better than PMIA for the WC model. Among all test cases (including additional ones in Sect. 5), there are only a couple of cases where other scalable heuristics have matching influence spread as PMIA. This means that PMIA performs consistently well among the best scalable heuristics while others such as PageRank and DegreeDiscountIC are not stable—

Fig. 8 Running time and average arborescence size of PMIA vs. the threshold $1/\theta$ in the WC model, for NetHEPT dataset



there exist a few cases that they perform well but in most other cases they perform not as well and sometimes they perform poorly comparing to PMIA. For running time, even at two million edge range, PMIA only takes 3 min to run. Therefore, PMIA has very good scalability and can handle million-sized or even larger graphs well.

Overall, we see that PMIA can scale beyond millions of edges, while Greedy and SP1M become too slow for half million edges or above. In all size ranges, PMIA consistently performs among the best algorithms (including Greedy and SP1M), while in most cases it significantly outperforms the rest scalable heuristics to as much as 100–260% increase in influence spread.

Tuning of parameter θ . We investigate the effect of the tuning parameter θ on the running time and the influence spread of our algorithm. Figure 8 shows that the running time increases when the θ value decreases, as expected. More interestingly, the running time is almost linear to $1/\theta$. This can be roughly explained as follows. First, by the running time analysis of Sect. 3.2, we can see that when n and k are fixed and θ is varied, the dominant term is a quadratic term $n_{o\theta}n_i\theta$, which means the running time is proportional to the square of the average arborescence size. Figure 8 further shows that the average arborescence size is about $O(\sqrt{1/\theta})$. Therefore altogether the running time is close to a linear relationship with $1/\theta$.

Figure 9 shows the change of influence spread with respect to the running time of our algorithm for the NetHEPT dataset in the WC model. Since the relationship between running time and $1/\theta$ is linear, it does not matter much if we use running time or $1/\theta$ as x -axis. The result indicates that as running time increases (θ decreases), the influence spread also increases, meaning that we obtain better quality results. Comparing other algorithms also shown in the figure, we see that on one side, we can tune $1/\theta$ to a larger value so that our influence spread can match the one provided by SP1M with at least 10 times speedup, while on the other side we can tune $1/\theta$ to a small value to get close to the running time of PageRank with matching influence spread. Therefore, we can use one algorithm to achieve different efficiency-effectiveness tradeoff needs by properly tuning the parameters.

One noticeable result is the knee in the curve of our algorithm. It means that the increase in influence spread is no longer significant after we lower θ to a certain level. This is because as shown in Fig. 8, arborescence size increases in square root of $1/\theta$




Fig. 9 Maximal influence spread by 50 seeds w.r.t. running time, for the NetHEPT dataset in the WC model

(and thus in square root of running time), while influence spread may change much slower after the arborescence grows beyond a certain size. The knee point suggests a good tuning point for the algorithm. If we select θ such that the influence-time tradeoff is close to the knee point, we could obtain the best gain from both influence spread and running time. Correlating with Fig. 8, we found that the corresponding knee point is close to the point where the change of arborescence size slows down (the dot with $1/\theta = 320$). We observe similar situations in other dataset that we did not report here. Thus, this suggests the following way of tuning parameter θ . Given a new graph, randomly sample a small portion of nodes in the graph to compute the average arborescence sizes with varying $1/\theta$, and find a point where the change of arborescence size slows down, and use the θ value at that point for the *PMIA* algorithm. The θ values selected in our experiments are based on this method.

5 Additional experiment results

In this section, we report additional results of our experiments on additional datasets, new propagation probability setting for the IC model, additional heuristic algorithms, online bounds of the optimal influence, and qualitative study of the effectiveness of influence maximization.

Additional datasets. Two additional datasets are tested. The first one is from the full paper list of the “Physics” section of e-print arXive, doted as NetPHY, which contains 37,154 nodes and 231,584 edges, the same one used in [Chen et al. \(2009\)](#). The second dataset is obtained from the authors of [Tang et al. \(2009\)](#), which is another collaboration network extracted from the data mining research area in the ArnetMiner archive (<http://www.arnetminer.org>) with 679 nodes and 1687 edges, and is denoted as DM. Some basic statistics about these networks are given in Table 2.

Table 2 Statistics of NetPHY and DM

Dataset	NetPHY	DM
#Node	37K	679
#Edge	174K	1687
Average degree	12.5	4.97
Maximal degree	286	63
#Connected component	3883	1
Largest component size	19873	679
Average component size	9.57	679

Generating propagation probabilities. We use one more model to generate propagation probabilities, as described below. We also use a different set of values for the TRIVALENCY model.

- TAP model: This is a model developed recently in [Tang et al. \(2009\)](#), in which the authors develop a *topical affinity propagation* (TAP) algorithm to compute propagation probabilities of every edge based on structural and topical information available to the graph. The resulting propagation probabilities are also nonuniform. For the DM dataset, we use the propagation probabilities computed from the topical information available to the dataset. For the NetHEPT dataset, we use uniform topic distribution among nodes for TAP to compute propagation probabilities, since specific topical information is not available. The NetPHY dataset is too large for the TAP algorithm in a single machine, so we do not use it for this dataset.
- TRIVALENCY model: use probability values 0.2, 0.04, 0.008 instead of 0.1, 0.01 and 0.001 in Sect. 4.

Algorithms. We include the following additional algorithms for comparison.

- Degree: The simple heuristic that selects the k nodes with the largest out-degrees in the graph.
- WeightedDegree: The weighted degree of a node is the sum of propagation probabilities on all its outgoing edges. This heuristic selects the k nodes with the largest weighted degrees.
- SPM: The shortest-path based algorithm of [Kimura and Saito \(2006\)](#), also enhanced with the lazy-forward optimization of [Leskovec et al. \(2007\)](#). In this version, only the shortest paths from S to a node v are counted for influence. Note that SP1M is an enhanced version of SPM, in which both the shortest paths and paths one hop longer than the shortest paths from S to a node v are counted for influence.

Results on influence spread. Figures 10–14 show the results on influence spreads, where we also include results for algorithms we tested in Sect. 4. The results are mainly self-explanatory, and consistent with the finding we concluded in Sect. 4.2. Overall PMIA performs consistently well over all datasets and all propagation models, matching or very close to the performance of Greedy and SPM/SP1M while outperforming the rest heuristics, including the new ones we tested here. A special attention is on

Fig. 10 Influence spread for different algorithms in the WC model, for the NetHEPT dataset




Fig. 11 Influence spread for different algorithms in the WC model, for the NetPHY dataset

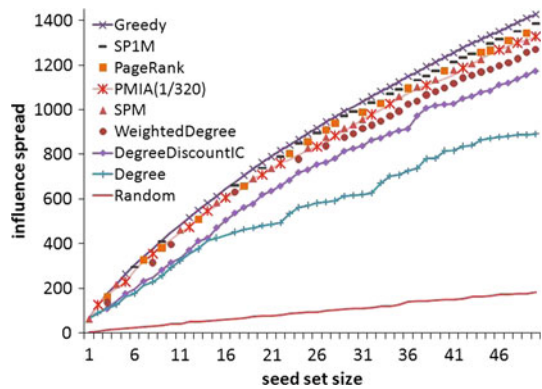


Fig. 12 Influence spread for different algorithms in the TAP model, for the NetHEPT dataset

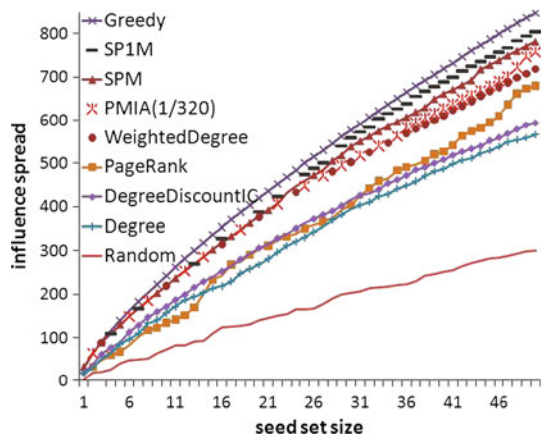


Fig. 13, which shows that Greedy performs visibly worse than PMIA. The reason is Greedy is too slow and we have to reduce the number of simulations for influence spread estimation from 20,000 to 200, causing it to lose accuracy on estimation (see the running time section for a reason why it is slow). This is also an indication that we

Fig. 13 Influence spread for different algorithms in the TRIVALENCY model with three probabilities 0.2, 0.04, 0.008, for the NetHEPT dataset





Fig. 14 Influence spread for different algorithms in the TAP model, for the DM dataset



cannot easily speed up Greedy by reducing the number of simulations. Another point worth explanation is that WeightedDegree performs quite well, closing to PMIA, in the two TAP model related tests (Figs. 12, 14). The reason is because WeightedDegree only considers influence propagated within one-step neighbors while the TAP model is likely to generate influence model in which most influences are indeed only propagated within one step. However, WeightedDegree performs not as well in other tests, showing that it has no consistent performance as PMIA.

Running time. Figure 15 shows the running time of different algorithms when selecting 50 seeds for 3 different tests: NetPHY using the WC model, NetHEPT using the TAP model, and NetHEPT using the TRIVALENCY model (with probabilities 0.2, 0.04, and 0.008). The result is again consistent with what we have seen in Sect. 4. Two specific points we would like to explain are as follows. First, Greedy is much slower in the TRIVALENCY model. This is because in this model after selecting a seed, the marginal influence spread for the next seed candidate decreases dramatically, causing a lot of re-evaluations of marginal influence spread for selecting the next seed and making the lazy forward optimization of Leskovec et al. (2007) much less effective




Fig. 15 Running time of different algorithms in 3 tests

than in other cases. Second, the running time of PMIA in the third test (NetHEPT on TRIVALENCY) is very fast (67ms). The reason that it is much faster than the other cases is because it uses a larger θ value of $1/20$, which generate smaller arborescences with depth at most 1. In this case, its running time is always close to that of the WeightedDegree, with the overhead only in the maintenance of the arborescence data structures and repeated updates due to seed selection. Thus we see that tuning θ could achieve much better running time. On the other hand, our PMIA is still better than WeightedDegree in influence spread (see Fig. 13), because it considers overlapping influences among seeds while WeightedDegree does not.

Online Bounds. We implement Algorithm 7 in Section 3.4 to generate online bounds for PMIA algorithm. The online approximation bound is from 76% to 82% when the number of seeds vary from 30 to 50, which is tighter than the offline bound $1 - 1/e \approx 63\%$.

Case Study. To have an intuition of the influence maximization results, we show an example on the collaboration network in data mining area, i.e., the DM dataset. Table 3 presents the discovered seed nodes using three different models to set the propagation probabilities. We see that different models lead to different results. In this article, we do not study the problem of model selection. It remains an open problem to decide which model fits the reality. We just give a simple insight based on the influence maximization results here. For each set of seed nodes, we calculate the *density* measure in network science, dividing the sum of coauthored papers by the number of different pairs between seeds, i.e. $\frac{10 \times 9}{2} = 45$ in our case. The larger the density is, the more redundancy the seed nodes have in the network. To maximize the influence spread, it is desirable to minimize the density. We see that TAP leads to a seed set with the minimal density. It can be observed that Philip S. Yu and Jian Pei are not selected as a top-10 seed in TAP model due to the large overlap of the nodes influenced by them and by other seeds. One explanation is that in Uniform IC and WC model,

Table 3 Discovered seed nodes in influence maximization by the greedy algorithm with different influence setting

No.	Uniform	WC	TAP
1	Philip S. Yu	Philip S. Yu	Jiawei Han
2	Jiawei Han	Jiawei Han	Qiang Yang
3	Christos Faloutsos	Wei Wang	Christos Faloutsos
4	Qiang Yang	Christos Faloutsos	Heikki Mannila
5	Heikki Mannila	Heikki Mannila	Vipin Kumar
6	Wei Wang	C. Lee Giles	C. Lee Giles
7	Jian Pei	Shusaku Tsumoto	Saso Dzeroski
8	Vipin Kumar	Jian Pei	Graham J. Williams
9	Bing Liu	Bing Liu	Myra Spiliopoulou
10	C. Lee Giles	Joost N. Kok	Eamonn J. Keogh
Density	0.4222	0.2444	0.1778

Uniform: influence=unique probability (0.01); *WC*: influence=inverse of in-degree; *TAP* influence = result of TAP

influence from neighbors to a node is independent, while in TAP, the correlation of influence between neighbors is reflected in the probability setting: one node has only a few strong influencers; and if a and b both influence c but a also influences b , a will tend to have larger influence on c . Whether one model is more desirable than another depends on the specific applications.

6 Future work

One possible future research is to further explore the advantages of our MIA heuristic. For example, we believe that MIA heuristic fits into the parallel computation framework better than the greedy algorithm and shortest-path based SP1M heuristic. This is because our computation are restricted on local arborescences around nodes, and thus the graph can be easily partitioned for parallel computation, with sharing data only needed for arborescences at the boundary. On the contrary, the greedy algorithm and the SP1M heuristic need simulations and computations among the whole graph, so graph partition is difficult, and parallel computation is only possible for different computation tasks that require sharing of the entire graph. Another future direction is to look for hybrid approaches that combine the advantages of different algorithms to further improve the efficiency and effectiveness of influence maximization.

Beyond influence maximization, one interesting direction that requires further research is the data mining of social influence from real online social network datasets. A few studies have started to address this issue for blogspace (Gruhl et al. 2004), academic collaboration network (Tang et al. 2009), and online social media (Goyal et al. 2010). In fact, we used a dataset from Tang et al. (2009) with propagation probabilities computed by their algorithm. We plan to study social influence mining in other social media and design appropriate algorithms for these social media. Social

influence mining and influence maximization together will form the key components that enable prevalent viral marketing in online social networks.

References

- Aiello W, Chung FRK, Lu L (2000) A random graph model for massive graphs. In: STOC '00
- Bakshy E, Karrer B, Adamic LA (2009) Social influence and the diffusion of user-created content. In: EC '09: Proc. 10th ACM Conf. Electronic Commerce
- Brin S, Page L (1998) The anatomy of a large-scale hypertextual web search engine. *Comput Netw* 30(1-7):107–117
- Cha M, Mislove A, Gummadi KP (2009) A measurement-driven analysis of information propagation in the flickr social network. In: WWW '09
- Chen W, Wang Y, Yang S (2009) Efficient influence maximization in social networks. In: KDD '09
- Chen W, Wang C, Wang Y (2010a) Scalable influence maximization for prevalent viral marketing in large-scale social networks. In: KDD '10
- Chen W, Yuan Y, Zhang L (2010b) Scalable influence maximization in social networks under the linear threshold model. In: ICDM '10
- Chen W, Collins A, Cummings R, Ke T, Liu Z, Rincon D, Sun X, Wang Y, Wei W, Yuan Y (2011) Influence maximization in social networks when negative opinions may emerge and propagate. In: SDM '11
- Cui P, Wang F, Liu S, Ou M, Yang S, Sun L (2011) Who should share what?: item-level social influence prediction for users and posts ranking. In: SIGIR '11
- Domingos P, Richardson M (2001) Mining the network value of customers. In: KDD '01
- Feige U (1998) A threshold of $\ln n$ for approximating set cover. *J ACM* 45(4):634–652
- Freeman L (1979) Centrality in social networks: conceptual clarification. *Soc Netw* 1:215–239
- Goyal A, Bonchi F, Lakshmanan LV (2010) Learning influence probabilities in social networks. In: WSDM '10
- Gruhl D, Guha RV, Liben-Nowell D, Tomkins A (2004) Information diffusion through blogspace. In: WWW '04
- Kempe D, Kleinberg JM, Tardos É (2003) Maximizing the spread of influence through a social network. In: KDD '03
- Kimura M, Saito K (2006) Tractable models for information diffusion in social networks. In: PKDD '06
- Leskovec J, Krause A, Guestrin C, Faloutsos C, VanBriesen J, Glance NS (2007) Cost-effective outbreak detection in networks. In: KDD '07
- Misner IR (1999) *The world's best known marketing secret: Building your business with word-of-mouth marketing*, 2nd edn. Bard Press, Austin
- Nail J (2004) *The consumer advertising backlash*. Forrester Research and Intelliseek Market Research Report
- Nemhauser G, Wolsey L, Fisher M (1978) An analysis of the approximations for maximizing submodular set functions. *Math Program* 14:265–294
- Richardson M, Domingos P (2002) Mining knowledge-sharing sites for viral marketing. In: KDD '02
- Rodriguez MG, Leskovec J, Krause A (2010) Inferring networks of diffusion and influence. In: KDD '10
- Streeter M, Golovin D (2007) An online algorithm for maximizing submodular functions. Technical Report CMU-CS-07-171, Carnegie Mellon University, Pittsburgh
- Tang J, Sun J, Wang C, Yang Z (2009) Social influence analysis in large-scale networks. In: KDD '09
- Valiant LG (1979) The complexity of enumeration and reliability problems. *SIAM J Comput* 8(3):410–421
- Vazirani VV (2004) *Approximation algorithms*. Springer, Berlin