

# Simultaneous Influencing and Mapping for Health Interventions

Leandro Soriano Marcolino<sup>1</sup>, Aravind Lakshminarayanan<sup>2</sup>, Amulya Yadav<sup>1</sup>, Milind Tambe<sup>1</sup>

<sup>1</sup>University of Southern California, Los Angeles, CA, 90089, USA

{sorianom, amulyaya, tambe}@usc.edu

<sup>2</sup>Indian Institute of Technology, Madras, Tamil Nadu, 600036, India

aravindsrinivas@gmail.com

## Abstract

Influence Maximization is an active topic, but it was always assumed full knowledge of the social network graph. However, the graph may actually be unknown beforehand. For example, when selecting a subset of a homeless population to attend interventions concerning health, we deal with a network that is not fully known. Hence, we introduce the novel problem of simultaneously influencing and mapping (i.e., learning) the graph. We study a class of algorithms, where we show that: (i) traditional algorithms may have arbitrarily low performance; (ii) we can effectively influence and map when the *independence of objectives* hypothesis holds; (iii) when it does not hold, the upper bound for the influence loss converges to 0. We run extensive experiments over four real-life social networks, where we study two alternative models, and obtain significantly better results in both than traditional approaches.

## 1 Introduction

Influencing a social network is an important technique, with great potential to positively impact society, as we can modify the behavior of a community. For example, we can increase the overall health of a population; Yadav et al. (2015), for instance, spread information about HIV prevention in homeless populations. However, although influence maximization has been extensively studied (Kempe, Kleinberg, and Tardos 2003; Cohen et al. 2014; Golovin and Krause 2010), their main motivation is viral marketing, and hence they assume that the social network graph is fully known, generally taken from some social media network (such as Facebook).

However, the graphs recorded in social media do not really represent all the people and all the connections of a population. Most critically, when performing interventions in real life, we deal with large degrees of lack of knowledge. Normally the social agencies have to perform several interviews in order to learn the social network graph (Marsden 2005). These highly unknown networks, however, are exactly the ones we need to influence in order to have a positive impact in the real world, beyond product advertisement.

Additionally, learning a social network graph is very valuable *per se*. Agencies also need data about a population, in

order to perform future actions to enhance their well-being, and better actuate in their practices (Marsden 2005). As mentioned, however, the works in influence maximization are currently ignoring this problem.

Each person in a social network actually knows other people, including the ones she cannot directly influence. Hence, each time we select someone for an intervention (to spread influence), we also have an opportunity to obtain knowledge from that person. Therefore, in this work we present for the first time the problem of simultaneously influencing and mapping a social network. We study the performance of the classical greedy influence maximization algorithm in this context, and show that it can be arbitrarily low. Hence, we study a class of algorithms for this problem, and show that we can effectively influence and map a network when *independence of objectives* holds. For the interventions where it does not hold, we give an upper bound in our loss, which converges to 0. We study an approximation of our main algorithm, that works as well but requiring fewer assumptions.

We perform a large scale experimentation using four real life social networks of homeless populations, where we show that our algorithm is competitive with previous approaches in terms of influence (even outperforming them in hard cases), and is significantly better in terms of mapping.

## 2 Related Work

The influence maximization problem has recently been a very popular topic of research. Normally, the main motivation is viral marketing in social media (like Facebook or MyPlace). Hence, previous works assume full knowledge of the social network graph. The classical result is Kempe, Kleinberg, and Tardos (2003), where they study the approximation guarantee of a greedy algorithm for the “independent cascade model”. Golovin and Krause (2010) further extended that result to the case where before picking each node we are able to observe which ones are already influenced or not. However, they still assume full knowledge of the graph. Cohen et al. (2014) focus on how to quickly estimate the potential spread of one node, since running simulations (as needed by the greedy algorithm) in large graphs is computationally expensive. A different view was studied by Yadav et al. (2015), where they analyze a model where nodes try to influence their neighbors multiple times, and each edge has an existence probability. We deal here with a different kind of uncertainty, as in our

model whole portions of the graph are completely unknown.

This work also relates to the problem of submodular optimization (influence is a submodular function) by selecting elements without knowledge of the whole set. Badanidiyuru et al. (2014) present an algorithm for selecting the best subset, giving elements arriving from a “stream”. Hence, given enough time the full set would be seen, and which elements are discovered does not depend on which ones are selected.

Our problem is also related to the classical max- $k$ -cover problem (Khuller, Moss, and Naor 1999), where we must pick  $k$  subsets that maximize our coverage of a set. In that case, however, the elements of each subset is known. There is also an online version of the max- $k$ -cover problem (Alon et al. 2009), where we pick one subset at a time, but an adversary fixes which element must be covered next. However, the set and the available subsets are still known in advance. Similarly, Grandoni et al. (2008) study the case of covering a set whose elements are randomly chosen from a universe of possible elements. Another important related problem is the submodular secretary problem (Batani, Hajiaghayi, and Zadimoghaddam 2013), where we again pick a subset to optimize a submodular function. However, in that case, we receive one element at a time, and we must make an irrevocable decision of either keeping it or not.

Finally, this work also relates to sequential decision making with multiple objectives (Rojers et al. 2013). Here we do not aim at computing an optimal policy (which is computationally expensive), but at studying a greedy method, similar to other works in influence maximization. Our algorithm is a scalarization over two objectives, but such method was never studied before to influence and map social networks.

### 3 Influencing and Mapping

We consider the problem of maximizing the influence in a social network. However, we start by knowing only a subgraph of the social network. Each time we pick a node to influence, it may teach us about subgraphs of the network. Our objective is to spread influence, at the same time learning the network graph (i.e., mapping). We call this problem as “Simultaneous Influencing and Mapping” (SIAM). In this paper, we consider a version of SIAM where we only need to map the nodes that compose the network. We assume that we always know all the edges between the nodes of the known subgraph. For clarity, we will define formally here only the version of SIAM that we handle in this work. Therefore, unless otherwise noted, henceforth by SIAM we mean the version of the problem that is formally defined below.

Let  $G := (V, E)$  be a graph with a set of nodes  $V$  and edges  $E$ . We perform  $\eta$  interventions, where at each one we pick one node. The selected node is used to spread influence and map the network. We assume we do not know the graph  $G$ , we only know a subgraph  $G_k = (V_k, E_k) \subset G$ , where  $k$  is the current intervention number.  $G_k$  starts as  $G_k := G_0 \subset G$ . For each node  $v_i$ , there is a subset of nodes  $V^i \subset V$ , which will be called “teaching list”. Each time we pick a node  $v_i$ , the known subgraph changes to  $G_k := (V_{k-1} \cup V^i, E_k)$ , where  $E_k$  contains all edges between the set of nodes  $V_{k-1} \cup V^i$  in  $G$ . Our objective is to maximize  $|V_k|$ , given  $\eta$  interventions.

For each node  $v_i$ , we assume we can observe a number  $\gamma_i$ , which indicates the size of its teaching list. We study two versions: in one  $\gamma_i$  is the number of nodes in  $V^i$  that are not yet in  $G_k$  (hence, the number of new nodes that will be learned when picking  $v_i$ ). We refer to this version as “perfect knowledge”. In the other,  $\gamma_i := |V^i|$ , and thus we cannot know how many nodes in  $V^i$  are going to be new or intersect with already known nodes in  $V_k$ . We refer to this version as “partial knowledge”. The *partial knowledge* version is more realistic, as by previous experience we may have estimations of how many people a person with a certain profile usually knows. We study both versions in order to analyze how much we may lose with *partial knowledge*. Note that we may also have nodes with empty teaching lists ( $\gamma_i = 0$ ). The teaching list of a node  $v_i$  is the set of nodes that  $v_i$  will teach us about once picked, and is not necessarily as complete as the true set of all nodes known by  $v_i$ . Some nodes could simply refuse to provide any information.

Additionally, note that we are assuming the teaching list and the neighbor list to be independent. That is, a node may teach us about nodes that it is not able to directly influence. For instance, it is common to know people that we do not have direct contact with, or we are not “close” enough to be able to influence. Similarly, a person may not tell us about all her close friends, due to limitations of an interview process, or even “shame” to describe some connections. However, some readers could argue that people would be more likely to teach us about their direct connections. Hence, we will handle the case where the independence does not hold in our empirical experiments in Section 4.2.

Simultaneously to the problem of mapping, we also want to maximize the spread of influence over the network. We consider here the traditional independent cascade model, with observation, as in Golovin and Krause (2010). That is, a node may be either influenced or uninfluenced. An uninfluenced node may change to influenced, but an influenced node will never change back to uninfluenced. Each time we pick a node for an intervention, it will change to influenced. When a node changes from uninfluenced to influenced, it will “spread” the influence to its neighbors with some probability. That is, at each edge  $e$  there is a probability  $p_e$ . When a node  $v_1$  changes to influenced, if there is an edge  $e = (v_1, v_2)$ , the node  $v_2$  will also change to influenced with probability  $p_e$ . Similarly, if  $v_2$  changes to influenced, it will spread the influence to its neighbors by the same process. Influence only spreads in the moment a node changes from uninfluenced to influenced. That is, a node may only “try” one time to spread influence to its neighbors. As in Golovin and Krause (2010), we consider that we have knowledge about whether a node is influenced or not (but in our case, we can only know about nodes in the current known subgraph  $G_k$ ). Let  $I_k$  be the number of influenced nodes after  $k$  interventions. Our objective is to maximize  $I_k$  given  $\eta$  interventions. Influence may spread beyond  $G_k$ . Hence, we consider  $I_k$  as the number of influenced nodes in the full graph  $G$ . We denote as  $\sigma_i$  the expected number of nodes that will be influenced when picking  $v_i$  (usually calculated by simulations).

As mentioned, we want to attend both objectives simultaneously. Hence, we must maximize both  $|V_k|$  and  $I_k$ . It is

easy to show that SIAM is an NP-Complete problem:

**Proposition 1.** *SIAM is NP-Complete.*

*Proof.* Let  $\kappa$  be an instance of the influence maximization problem, with graph  $G$ . Consider now a SIAM problem where no node carries information and  $G_0 := G$ . If we can solve this SIAM problem we can also solve the influence maximization problem  $\kappa$ . Therefore, SIAM is NP-Complete.  $\square$

As SIAM is NP-Complete, similarly to previous influence maximization works (Kempe, Kleinberg, and Tardos 2003; Golovin and Krause 2010), we study greedy solutions. Like the exploration vs exploitation dilemmas in online learning (Valizadegan, Jin, and Wang 2011), the fundamental problem of SIAM is whether to focus on influencing or mapping the network. Hence, we propose as a general framework to select the node  $v_i$  such that:

$$v_i = \operatorname{argmax}(c_1 \times \sigma_i + c_2 \times \gamma_i) \quad (1)$$

Constants  $c_1$  and  $c_2$  control the balance between influencing or mapping.  $c_1 = 1, c_2 = 0$  is the classical influence maximization algorithm (“*influence-greedy*”);  $c_1 = 0, c_2 = 1$ , on the other hand, only maximizes the knowledge-gain at each intervention (“*knowledge-greedy*”).  $c_1 = c_2 = 1$  is an algorithm where both objectives are equally balanced (“*balanced*”). Different weights may also be used.

Remember that we defined two versions for the  $\gamma$  values: *perfect knowledge*, where we know how many new nodes a node will teach us about; and *partial knowledge*, where we do not know how many nodes will be new. In order to better handle the *partial knowledge* case, we also propose the “*balanced-decreasing*” algorithm, where  $c_2$  constantly decreases until reaching 0. Hence, we define  $c_2$  as:

$$c_2 := \begin{cases} c'_2 - \frac{1}{d} \times c'_2 \times k & \text{if } k \leq d \\ 0 & \text{otherwise} \end{cases}, \quad (2)$$

where  $c'_2$  is the desired value for  $c_2$  at the very first iteration, and  $d$  controls how fast  $c_2$  decays to 0.

### 3.1 Analysis

We begin by studying *influence-greedy*. It was shown that when picking the node  $v$  which  $\operatorname{argmax}(\sigma_v)$  at each intervention, we achieve a solution that is a  $(1 - 1/e - \epsilon)$  approximation of the optimal solution, as long as our estimation of  $\sigma_v$  (by running simulations) is “good enough” (Kempe, Kleinberg, and Tardos 2003). However, even though the actual influence spread may go beyond the known graph  $G_k$ , we can only run simulations to estimate  $\sigma_v$  in the current  $G_k$ . Hence, the previous results are no longer valid. In fact, in the next observation we show that we can obtain arbitrarily low-performing solutions by using *influence-greedy*.

**Observation 1.** *The performance of influence-greedy can be arbitrarily low in a SIAM problem.*

We show with an example. Consider the graph in Figure 1, and assume we will run 2 interventions (i.e., pick 2 nodes). There is a probability 1 to spread influence in any edge. Our initial knowledge is  $V_0 = \{A, A', B, B', C\}$ . A and B can influence  $A'$  and  $B'$ , respectively. However, C cannot influence




Figure 1: A graph where the traditional greedy algorithm has arbitrarily low performance.

any node. A, B,  $A'$  and  $B'$  have empty teaching lists. C, on the other hand, can teach us about a connected graph of  $z$  nodes. *Influence-greedy*, by running simulations on the known graph, picks nodes A and B, since each can influence one more node. The optimal solution, however, is to pick node C, which will teach us about the connected graph of  $z$  nodes. Then, we can pick one node in that graph, and influence  $z + 1$  nodes in total. Hence, the *influence-greedy* solution is only  $\frac{4}{z+1}$  of the optimal. As  $z$  grows, *influence-greedy* will be arbitrarily far from the optimal solution.

If we make some assumptions about the distribution of the teaching lists across the nodes, however, *influence-greedy* eventually maps the full graph given enough interventions. Let  $n = |V|$ , and  $n_k = |V_k|$ . We show the expected number of interventions to learn all  $n$  nodes (subtracted by a small  $\epsilon$ , for numerical stability). We study the *partial knowledge* version. Assume the size of the teaching list of each node is drawn from a uniform distribution on the interval  $[0, u]$ , and any node is equally likely to be in a teaching list. We consider that there is a probability  $\varphi$  that a node will have a non-empty teaching list.

**Proposition 2.** *The expected number ( $k_{full}$ ) of interventions for influence-greedy to learn  $n - \epsilon$  nodes is  $\frac{\log(-\frac{\epsilon}{n_0 - n})}{\log(1 - \frac{\varphi \times u}{2 \times n})}$ .*

*Proof.* Since *influence-greedy* is not considering  $\gamma$ , it picks nodes arbitrarily in terms of knowledge-gain. Hence, on average it selects the expected value of the uniform distribution,  $u/2$ . For each node  $v$  in a teaching list, the probability that it is not yet known is  $\frac{n - n_k}{n}$ . Therefore, the expected number of nodes known at one iteration  $k$  is:  $E[n_k] = \varphi \times \frac{u}{2} \times \frac{n - E[n_{k-1}]}{n} + E[n_{k-1}]$ . Solving the recurrence gives:  $E[n_k] = n_0 \times (1 - \frac{\varphi \times u}{2 \times n})^k - n \times (1 - \frac{\varphi \times u}{2 \times n})^k + n$ . Solving for  $E[n_k] = n - \epsilon$  gives that the expected number of interventions is:  $k_{full} = \frac{\log(-\frac{\epsilon}{n_0 - n})}{\log(1 - \frac{\varphi \times u}{2 \times n})}$ .  $\square$

$k_{full}$  quickly increases as  $\varphi$  (or  $u$ ) decreases. In Section 4, we study experimentally the impact of  $\varphi$  on the performance of *influence-greedy*.

Now, let’s look at *balanced*. Clearly, it will learn the full graph with a lower number of expected interventions than *influence-greedy*. However, although intuitively *balanced* may seem reasonable, its performance may also quickly degrade if we assume *partial knowledge* (i.e.,  $\gamma_i = |V^i|$ ).

**Proposition 3.** *The performance of the balanced algorithm degrades as  $n_k \rightarrow n$ , if  $\gamma_i = |V^i|$ .*

*Proof.* Each node in the teaching list of a node  $v_i$  has probability  $\frac{n-n_k}{n}$  of being a yet unknown node. Hence, the expected number of unknown nodes that will be learned by picking a node with teaching list size  $\gamma_i$  is:  $E[\text{new}] = \gamma_i \times \frac{n-n_k}{n}$ . As  $n_k \rightarrow n$ ,  $E[\text{new}] \rightarrow 0$ . Hence, when  $n_k \rightarrow n$ , *balanced* picks a node  $v$  that maximizes  $\sigma_v + \gamma_v$ , thus missing to select nodes  $v_o$  (if available) with  $\sigma_o > \sigma_v$ ,  $\sigma_o + \gamma_o < \sigma_v + \gamma_v$ , with no actual gains in mapping.  $\square$

This problem does not happen in the *perfect knowledge* version. Since the  $\gamma$  values only include new nodes,  $\gamma \rightarrow 0$  as  $n_k \rightarrow n$ , for all  $\gamma$ . Hence, in the *perfect knowledge* version, *balanced* converges to the same behavior as *influence-greedy* as  $k$  increases. In order to approximate this behavior for the *partial knowledge* case, we propose the *balanced-decreasing* algorithm, where the constantly decreasing  $c_2$  ‘‘simulates’’ the decreasing  $\gamma$  values.

We now show that *balanced* can match the performance of *influence-greedy* in terms of influence, but at the same time mapping better the network. As we just discussed that *perfect knowledge* can be approximated by using *balanced-decreasing*, we focus here on the *perfect knowledge* case. We show that when the *independence of objectives* hypothesis holds (defined below), *balanced* plays the same as *influence-greedy* or better, while *influence-greedy* may still fail in terms of mapping. If the hypothesis does not hold, our influence loss at one intervention will be bounded by  $u/2 \rightarrow 0$ .

Let  $V_k^\sigma$  be a subset of  $V_k$  where each  $v \in V_k^\sigma$  maximizes  $\sigma$  in the current intervention  $k$ . Similarly, let  $V_k^\gamma \subset V_k$ , where each  $v \in V_k^\gamma$  maximizes  $\gamma$  in the current intervention  $k$ . As before, we consider that the teaching list size of a node is given by a uniform distribution, but since  $\gamma$  tends to decrease at each intervention, we denote the interval as  $[0, u_k]$ .

Clearly, any node in the set  $V_k^{\text{Good}} := V_k^\sigma \cap V_k^\gamma$  should be selected, as they maximize both objectives. Hence, when  $V_k^{\text{Good}} \neq \emptyset$  it is possible to simultaneously maximize both objectives, and thus we say that the *independence of objectives* hypothesis holds. Since we are studying greedy-algorithms, both *balanced* and *influence-greedy* lack optimality guarantees. Hence, we focus here on a ‘‘local’’ analysis, and show that given a set of  $k$  possible interventions (with the same graph state across both algorithms at each intervention), *balanced* is able to pick nodes that spread as much influence as *influence-greedy*. Moreover, when *balanced* picks a different node, our loss is bounded by  $u_k/2$ . As  $u_k \rightarrow 0$  with  $k \rightarrow \infty$ , our loss also converges to 0.

**Proposition 4.** *Balanced selects nodes that spread as much influence as influence-greedy, if  $|V_k^\sigma| > n_k/2$  and  $|V_k^\gamma| > n_k/2$ , or as  $k \rightarrow \infty$ . Influence-greedy, on the other hand, selects worse nodes than balanced in terms of mapping with probability  $1 - \frac{|V_k^\sigma \cap V_k^\gamma|}{|V_k^\sigma|}$ . Moreover, when balanced selects a node with worse  $\sigma$  than influence-greedy, the expected influence loss is bounded by  $u_k/2$ , which  $\rightarrow 0$  as  $k \rightarrow \infty$ .*

*Proof.* As *balanced* plays  $\text{argmax}(\sigma + \gamma)$ , if there is a node  $v \in V_k^{\text{Good}}$ , *balanced* picks  $v$ . *Influence-greedy*, however, selects an arbitrary node in  $V_k^\sigma$ . Hence, it picks a node  $v \in V_k^{\text{Good}}$  with probability  $\frac{|V_k^\sigma \cap V_k^\gamma|}{|V_k^\sigma|}$ . Therefore, for all in-

terventions where  $V_k^{\text{Good}} \neq \emptyset$ , *balanced* selects a node in  $V_k^{\text{Good}}$ , while *influence-greedy* makes a mistake in terms of mapping with probability  $1 - \frac{|V_k^\sigma \cap V_k^\gamma|}{|V_k^\sigma|}$ .

We consider now the probability of  $V_k^{\text{Good}} \neq \emptyset$  across  $k$  interventions. Clearly, if  $|V_k^\sigma| > n_k/2$ , and  $|V_k^\gamma| > n_k/2$ , we have  $V_k^{\text{Good}} \neq \emptyset$ . If not, note that as  $k \rightarrow \infty$ ,  $n_k \rightarrow n$ . Therefore,  $V_k^\gamma \rightarrow V_k$  (since all  $\gamma \rightarrow 0$ , all nodes will have the same teaching list size), thus  $V_k^{\text{Good}} \rightarrow V_k^\sigma \neq \emptyset$ . Hence, the probability of  $V_k^{\text{Good}} \neq \emptyset$  goes to 1 as  $k \rightarrow \infty$ .

Let’s study now the case when  $V_k^{\text{Good}} = \emptyset$ . Let  $v_1$  be the node in  $V_k^\sigma$  picked by *influence-greedy*, and  $v_2$  be the node in  $V_k \setminus V_k^\sigma$  with the largest  $\gamma_2$ . Since  $V_k^{\text{Good}} = \emptyset$ , we must have that  $\sigma_1 > \sigma_2$ , and  $\gamma_2 > \gamma_1$ . However, as long as  $\gamma_2 - \gamma_1 < \sigma_1 - \sigma_2$ , *balanced* still selects  $v_1$  (or an even better node). In the worst case, the expected value for  $\gamma_2$  is the expected maximum of the uniform distribution:  $E[\gamma_2] = u_k - u_k/(n_k + 1) \leq u_k \cdot \gamma_1$ , on the other hand, has the expected value of the uniform distribution  $E[\gamma_1] = u_k/2$ . Hence, as long as  $\sigma_1 - \sigma_2 > u_k/2$ , in expectation *balanced* still picks  $v_1$  (or an even better node). Moreover, when *balanced* does not pick  $v_1$ , our loss in terms of influence at intervention  $k$  is at most  $u_k/2$ . Since  $\gamma \rightarrow 0$  as  $n_k \rightarrow n$ ,  $u_k/2 \rightarrow 0$  as  $k \rightarrow \infty$ .  $\square$

Proposition 4 shows that we may experience loss in one intervention, when comparing *influence-greedy* with *balanced*. However, the loss is bounded by  $u_k/2$ , which goes to 0 as the number of interventions grows. Moreover, when we do not update the  $\gamma$  values, we can use the *balanced-decreasing* algorithm to simulate the same effect. Additionally, in Proposition 4 we considered the same graph states at each intervention across both algorithms. In practice, however, since *balanced* is able to map the network faster, any loss experienced in the beginning when  $k$  is low can be compensated by playing better later with full knowledge of the graph, while *influence-greedy* may still select nodes with lower  $\sigma$  due to lack of knowledge. As noted in Observation 1, lack of knowledge of the full graph can make *influence-greedy* play with arbitrarily low performance. In Section 4.2 we perform an empirical analysis assuming a power law model for the teaching lists, and we note here that our main results still hold.

## 4 Results

We run experiments using four real life social networks of the homeless population of Los Angeles, provided by us from Eric Rice, from the School of Social Work of the University of Southern California. All the networks are friendship-based social networks of homeless youth who visit a social agency. The first two networks (A, B) were created through surveys and interviews. The third and fourth networks (Facebook, MySpace) are online social networks of these youth created from their Facebook and MySpace profiles, respectively. Computation for the work described in this paper was supported by the University of Southern California’s Center for High-Performance Computing (hpc.usc.edu).

We run 100 executions per network. At the beginning of each execution, 4 nodes are randomly chosen to com-

pose our initial subgraph ( $G_0$ ). As mentioned, we consider that we always know the edges between the nodes of our current knowledge graph ( $G_k$ ). We noticed similar tendencies in the results across all four social networks. For clarity, due to space constraints, we plot here the results considering all networks simultaneously (that is, we average over all the 400 executions). In the appendix (available at <http://teamcore.usc.edu/people/sorianom/a16-ap.pdf>) we show the individual results for each network. In all graphs, the error bars show the confidence interval, with  $\rho = 0.01$ . When we say that a result is significantly better than another, we mean with statistical significance according to a  $t$ -test with  $\rho \leq 0.01$ , unless noted otherwise. The size of each network is: 142, 188, 33, 105; for A, B, Facebook and MySpace, respectively. We evaluate up to 40 interventions.

We measure the percentage of influence in the network (“Influence”) and percentage of known nodes (“Knowledge”) for *influence-greedy*, *knowledge-greedy*, *balanced* and *balanced-decreasing* (with  $c'_2 = 1.0$ , and  $d = 5$ ). In order to estimate the expected influence spread ( $\sigma_v$ ) of each node, we run 1000 simulations before each intervention. Estimating the expected influence through simulations is a common method in the literature. In our case, the simulations are run in the current known subgraph  $G_k$ , although the actual influence may go beyond  $G_k$ . Influenced nodes in  $G \setminus G_k$  will be considered when we measure Influence, but will not be considered in our estimation of  $\sigma_v$ . Concerning the teaching list size ( $\gamma_v$ ), we consider it to hold the number of *new* nodes that would be learned if  $v$  is selected, for *balanced* and *knowledge-greedy* (i.e., perfect knowledge). For *balanced-decreasing*, we consider  $\gamma_v$  to hold the full teaching list size, including nodes that are already known (i.e., partial knowledge). Therefore, we can evaluate if *balanced-decreasing* approximates well *balanced*, when perfect knowledge is not available.

We simulate the teaching lists, since there are no real world data available yet (we only have data about the connections in the four real life social networks). We study two models: (i) uniform, which follows the assumptions of our theoretical analysis; (ii) power law, which considers that nodes are more likely to teach us about others which are close to them in the social network graph. We present the second model to show that our conclusions hold irrespective of the uniform assumption. For each node, we decide whether it will have a non-empty teaching list according to a probability  $\varphi$ . We run experiments using different combinations of  $\varphi$ , probability of influence  $p$ , and  $c_1$  and  $c_2$  values.

#### 4.1 Uniform Model

Under the uniform model, if a node has a teaching list, we fix its size according to a uniform distribution from 0 to  $0.5 \times |V|$ . Each node in the graph is also equally likely to be in the teaching list of a node  $v_i$ . We consider here the teaching list and the neighbor list to be independent, as people may know others that they cannot influence, and they may also not tell us all their connections, as described before. The case where the teaching list and the neighbor list are not independent is considered in Section 4.2.

We run several parametrizations. Figure 2 shows the result at each intervention for  $\varphi = 0.5$  and  $p = 0.5$ . As we see




Figure 2: Results of 4 real world networks across many interventions, for  $p = 0.5$  and  $\varphi = 0.5$  (uniform distribution).

in Figure 2 (a), the Influence obtained by *influence-greedy*, *balanced*, and *balanced-decreasing* are very similar. In fact, out of all 40 interventions, their result is not significantly different in any of them (and they are significantly better than *knowledge-greedy* in around 75% of the interventions). This shows that *balanced* is able to successfully spread influence in the network, while at the same time mapping the graph. We can also notice that a perfect knowledge about the number of new nodes in the teaching lists is not necessary, as *balanced-decreasing* obtained close results to *balanced*.

Figure 2 (b) shows the results in terms of Knowledge. All algorithms clearly outperform *influence-greedy* with statistical significance. Moreover, the result for *knowledge-greedy*, *balanced* and *balanced-decreasing* are not significantly different in any of the interventions. This shows that we are able to successfully map the network (as well as *knowledge-greedy*), but at the same time spreading influence successfully over the network (as well as *influence-greedy*), even in the *partial knowledge* case. Hence, the *independence of objectives* hypothesis seems to hold at most interventions in the networks, since we could maximize both objectives simultaneously, as predicted in Proposition 4. Given enough interventions, however, *influence-greedy* is also able to map the network, as we discussed in Proposition 2.

It is also interesting to note that even though *influence-greedy* has much less information about the network (with significantly lower mapping performance in around 16 interventions), it is still able to perform as well as the other algorithms in terms of Influence. Observation 1, however, showed that its Influence performance can be arbitrarily low. As we discuss later, for some parametrizations we actually found that *influence-greedy* has significantly lower results than the other algorithms in terms of Influence as well.


In order to compare the results across different parametriza-

tions, we calculate the area under the curve (AUC) of the graphs. The closer the curves are to 1.0 the better, hence an AUC of 39 (that is, always at 1.0 across all 40 interventions) would be an “ideal” result. In Figure 3 (a) we show the results for a fixed influence probability value ( $p = 0.5$ ), but different teaching probability ( $\varphi$ ) values. First we discuss the results in terms of Influence (left-hand side of the graph). As we can see, except for *knowledge-greedy*, all algorithms obtain very similar results. However, for  $\varphi = 0.1$ , the Influence for *balanced* and *balanced-decreasing* is slightly better than *influence-greedy*, in the borderline of statistical significance ( $\rho = 0.101$  and  $0.115$ , respectively). Moreover, we can see that  $\varphi$  does impact the influence that we obtain over the network, although the impact is not big. For *influence-greedy*, from  $\varphi = 0.5$  to  $\varphi = 1.0$ , the difference is only statistically significant with  $\rho = 0.092$ . However, from  $\varphi = 0.1$  to  $\varphi = 0.5$  the difference is statistically significant with  $\rho = 3.26 \times 10^{-27}$ . Similarly, for all other algorithms there is a significant difference from  $\varphi = 0.1$  to  $\varphi = 0.5$ , while from  $\varphi = 0.5$  to  $\varphi = 1.0$  the difference is only significant with  $\rho < 0.1$  (except for *knowledge-greedy*, its difference is not significant between  $\varphi = 0.5$  and  $\varphi = 1.0$ ).


Let’s look at the results in terms of Knowledge, on the right-hand side of Figure 3 (a). We can see that  $\varphi$  has a much bigger impact in our mapping, as expected. *Knowledge-greedy*, *balanced* and *balanced-decreasing* are all significantly better than *influence-greedy*. However, we can notice that the difference between *influence-greedy* and the other algorithms decreases as  $\varphi$  increases. Similarly, when comparing *knowledge-greedy*, *balanced* and *balanced-decreasing*, we can notice that the difference between the algorithms also decreases as  $\varphi$  increases. For both  $\varphi = 0.1$  and  $\varphi = 0.5$ , however, the algorithms are not significantly different. Interestingly, when  $\varphi = 1$ , because of the lower variance, *knowledge-greedy* and *balanced* become significantly better than *balanced-decreasing*, even though the differences between the algorithms decreases.

In Figure 3 (b), we keep  $\varphi = 0.5$ , and change  $p$ . In the left-hand side we see the results for Influence. As expected, there is clearly a significant difference when  $p$  changes from 0.1 to 0.5. However, we can notice that the difference between the algorithms does not change significantly when  $p$  changes. In both cases, the differences between *influence-greedy*, *balanced* and *balanced-decreasing* are not significant. Additionally, in both cases all algorithms are significantly better than *knowledge-greedy*. In terms of Knowledge (right-hand side of the figure) we see that the influence probability has no impact in any algorithm, as it would be expected. For all algorithms, the difference between  $p = 0.1$  and  $p = 0.5$  is not statistically significant.

We also compare the *regret* obtained by the different algorithms at different influence probabilities and teaching probability values. First, we run the *influence-greedy* algorithm, but considering that we know the full graph (that is,  $G_k := G$ ). Although that solution is not optimal, it is the best known approximation of the optimal, hence we call it “perfect”. We calculate the AUC for *perfect*, and define the regret of an algorithm  $x$  as:  $AUC_{Perfect} - AUC_x$ . We analyze the regret in terms of Influence in Figure 4. Note that,



(a) Changing teaching probability



(b) Changing influence probability

Figure 3: Results of Influence and Knowledge for different teaching and influence probabilities (uniform distribution).




Figure 4: Regret for different teaching and influence probabilities (uniform distribution). Lower results are better.

in the figure, the lower the result the better. On the left-hand side we show the regret for  $p = 0.1$  and different  $\varphi$  values, while on the right-hand side we show for  $p = 0.5$ . All algorithms (except *knowledge-greedy*) have a similar regret, as it would be expected based on the previous results. However, we can notice here that the regret for *balanced* and *balanced-decreasing* when  $\varphi = 0.1$  and  $p = 0.5$  is actually lower than *influence-greedy*. The difference is statistically significant, with  $\rho = 0.069$  and  $\rho = 0.092$ , respectively. Hence, we can actually have a significantly better influence on the social network graph than the traditional greedy algorithm, when the teaching probability is low. However, when  $\varphi = 0.1$  and  $p = 0.1$ , even though the regret for *balanced* and *balanced-decreasing* is still lower than *influence-greedy*, it is not significant anymore (as there is larger variance when  $p = 0.1$ ). Additionally, we note that the difference in regret for *balanced* and *balanced-decreasing* is always not significant. It is also interesting to note that for some parametrizations, the regret of *influence-greedy* is actually close to 0, which means that in some cases lack of knowledge of the full graph does not significantly harm the influence performance. When  $p = 0.1$ , and  $\varphi = 0.5$  or  $\varphi = 1.0$ , the regret is not significant ( $\rho = 0.410$ , and  $\rho = 0.78$ , respectively). For  $p = 0.5$  and  $\varphi = 1.0$ , the regret is in the borderline of not being significant ( $\rho = 0.102$ ). In all other cases, the regret is significant.

We discussed 4 algorithms, but our framework can ac-

tually generate a variety of behaviors by using different  $c_1$  and  $c_2$  values. We tested 6 more combinations:  $\{(0.5, 2), (1, 0.5), (1, 1.5), (1, 2), (2, 0.5), (2, 1)\}$ , but we did not observe significant differences in comparison with the previous algorithms in the four social networks graphs. Hence, finding a good parametrization of  $c_1$  and  $c_2$  does not seem to be a crucial problem for the *balanced* algorithm.

## 4.2 Power Law Model

In order to show that our conclusions still hold under different models, we also run experiments considering a power law distribution for the teaching lists. The power law is a very suitable model for a range of real world phenomena. In fact, Andriani and McKelvey (2007), in a very comprehensive literature survey, lists 80 different kinds of phenomena which are modeled in the literature by power law distributions, and half of them are social phenomena. For example, it has been shown to be a good model for social networks, co-authorships networks, the structure of the world wide web and actor movie participation networks. A power law model also seems suitable in our case, as we can expect that a person will be very likely to teach us about the people who she has a direct connection with, and less and less likely to report people that are further away in the graph. Hence, when generating the teaching list of a node  $v_i$  in our experiments, each node  $v_o$  ( $v_o \neq v_i$ ) will be in its teaching list according to the following probability:  $p_o := (a - 1.0) \times h_o^{-a}$ , where  $1.0 < a \leq 2.0$ , and  $h_o$  is the shortest path distance between node  $v_i$  and  $v_o$ .  $a - 1.0$  represents the probability of a neighbor node  $v_o$  (i.e.,  $h_o = 1$ ) being selected. If node  $v_i$  and  $v_o$  are completely disconnected, we set  $h_o := |V|$ . Under this model the probability of a person teaching us about another is always strictly greater than 0, even though it may be very small if the respective nodes are very distant in the graph.

We fix  $a = 1.8$  (80% probability of each of a nodes' neighbors being in its teaching list). We show results for  $a = 1.2$  in the appendix, for the interested reader (and our conclusions still hold in the alternative parametrization). Similarly as before, Figure 5 shows the result at each intervention for  $\varphi = 0.5$  and  $p = 0.5$ . As we can see, our main conclusions still hold in the power law model. The Influence obtained by *influence-greedy*, *balanced*, and *balanced-decreasing* are very similar. Out of 40 interventions, their results are not significantly different in 39 ( $\rho \leq 0.05$ ), and they are significantly better than *knowledge-greedy* in around 60% of the interventions ( $\rho \leq 0.05$ ).

This time, however, *balanced-decreasing* obtained worse results than *balanced* in terms of Knowledge. Although up to iteration 4, *knowledge-greedy*, *balanced* and *balanced-decreasing* are not significantly different; both *knowledge-greedy* and *balanced* are significantly better than *balanced-decreasing* after that iteration. It is harder to obtain Knowledge under the power law model than under the uniform model (all algorithms converge slower to 1.0 than before). Hence, *balanced-decreasing* would require a slower decay speed (i.e., higher  $d$ ) in this case, in order to perform better.

We can also notice that all algorithms are significantly better than *influence-greedy* in all iterations, in terms of Knowledge. Note that under the uniform model, *influence-greedy*





Figure 5: Results of 4 real world networks across many interventions, for  $p = 0.5$  and  $\varphi = 0.5$  (power law distribution).


was not significantly worse than the other algorithms after iteration 20 ( $\rho \leq 0.1$ ). Hence, as expected, *influence-greedy* becomes relatively worse than the other algorithms when we assume a model where mapping is harder.

We calculate the AUC, in order to compare different parametrizations. Figure 6 (a) shows the result for a fixed influence probability value ( $p = 0.5$ ), and different teaching probability ( $\varphi$ ) values. As before, except for *knowledge-greedy*, all algorithms in general obtain similar results in terms of Influence. We notice, however, that for  $\varphi = 0.1$ , the result for *balanced* and *balanced-decreasing* is actually significantly better than *influence-greedy* ( $\rho = 0.064$  and  $0.038$ , respectively). Again, we also notice that  $\varphi$  significantly impacts the influence that we obtain over the network, although the impact is small. For all algorithms, the impact is significant from  $\varphi = 0.1$  to  $\varphi = 0.5$ . For *influence-greedy*, the difference is statistically significant with  $\rho = 1.39 \times 10^{-31}$ , while for *knowledge-greedy*, *balanced* and *balanced-decreasing*, there is statistically significant difference with  $\rho = 4.529 \times 10^{-6}$ ,  $3.258 \times 10^{-14}$  and  $8.429 \times 10^{-20}$ , respectively. However, the impact of  $\varphi$  increasing from 0.5 to 1 is not that significant for *knowledge-greedy*, *balanced* and *balanced-decreasing* ( $\rho = 0.41, 0.02, 0.03$ , respectively), while for *influence-greedy* the change has significant impact ( $\rho = 10^{-6}$ ). In terms of Knowledge, we can see that all algorithms are significantly better than *influence-greedy* for all  $\varphi$  values (with  $\rho \leq 3.464015 \times 10^{-19}$ ). However, this time we notice that *knowledge-greedy* and *balanced* are significantly better than *balanced-decreasing* for all  $\varphi$ . As mentioned, a different decay speed  $d$  is necessary in this case.

In Figure 6 (b), we show different values of  $p$  for  $\varphi = 0.5$ . As before, in terms of Influence the difference between *influence-greedy*, *balanced* and *balanced-decreasing* is not significant, and all algorithms are significantly better than



(a) Changing teaching probability



(b) Changing influence probability

Figure 6: Influence and Knowledge for different teaching and influence probabilities (power law distribution).




Figure 7: Regret for different teaching and influence probabilities (power law distribution). Lower results are better.

*knowledge-greedy*. In terms of Knowledge, the influence probability does not affect *knowledge-greedy*, *balanced* nor *balanced-decreasing* significantly, as expected. This time, however, *influence-greedy* obtains a significantly better result for  $p = 0.1$  than for  $p = 0.5$ . This may happen because for  $p = 0.1$  *influence-greedy* has a higher tendency of selecting nodes with a high number of neighbors, which also tends to be the ones with large teaching lists. Note that this does not happen when the teaching and neighbor lists are independent.

Figure 7 shows the regret (lower results are better). We can notice similar results as before: all algorithms have similar regret (except for *knowledge-greedy*), and the regret for *balanced* and *balanced-decreasing* when  $\varphi = 0.1$  and  $p = 0.5$  is again significantly lower than *influence-greedy* ( $\rho = 0.019$  and  $0.009$ , respectively). This time, however, we can notice that for some parametrizations *balanced-decreasing* is actually the algorithm with the lowest regret. For  $p = 0.5$  and  $\varphi = 0.5$ , *balanced-decreasing* is better than *balanced* with  $\rho = 4.7 \times 10^{-4}$ . Hence, even though *balanced-decreasing* performed relatively worse than under the uniform model in terms of Knowledge, it is actually the best algorithm in terms of Influence for some parametrizations.

## 5 Conclusion

We introduced the novel problem of simultaneously influencing and learning the graph of (i.e., mapping) a social network.

We show theoretically and experimentally that an algorithm which locally maximizes both influence and knowledge performs as well as an influence-only greedy algorithm in terms of influence, and as well as a knowledge-only greedy approach in terms of knowledge. We present an approximation of our algorithm that gradually decreases the weight given to knowledge-gain, which requires fewer assumptions. Results show that the approximation works well, and all algorithms can even significantly influence more nodes than the traditional greedy influence maximization algorithm when nodes have a low knowledge probability.

**Acknowledgments:** This research was supported by MURI grant W911NF-11-1-0332, and by IUSSTF. The authors would like to thank Eric Rice for the social networks data.

## References

- Alon, N.; Awerbuch, B.; Azar, Y.; Buchbinder, N.; and Naor, J. S. 2009. The online set cover problem. *SIAM Journal on Computing* 39(2).
- Andriani, P., and McKelvey, B. 2007. Beyond gaussian averages: Redirecting international business and management research toward extreme events and power laws. *Journal of International Business Studies* 38(7).
- Badanidiyuru, A.; Mirzasoleiman, B.; Karbasi, A.; and Krause, A. 2014. Streaming submodular maximization: Massive data summarization on the fly. In *KDD*.
- Batani, M.; Hajiaghayi, M.; and Zadimoghaddam, M. 2013. Submodular secretary problem and extensions. *TALG* 9(4).
- Cohen, E.; Delling, D.; Pajor, T.; and Werneck, R. F. 2014. Sketch-based influence maximization and computation: Scaling up with guarantees. Technical report, Microsoft Research.
- Golovin, D., and Krause, A. 2010. Adaptive submodularity: A new approach to active learning and stochastic optimization. In *COLT*.
- Grandoni, F.; Gupta, A.; Leonardi, S.; Miettinen, P.; Sankowski, P.; and Singh, M. 2008. Set covering with our eyes closed. In *FOCS*.
- Kempe, D.; Kleinberg, J.; and Tardos, E. 2003. Maximizing the spread of influence through a social network. In *KDD*.
- Khuller, S.; Moss, A.; and Naor, J. S. 1999. The budgeted maximum coverage problem. *Information Processing Letters* 70(1).
- Marsden, P. V. 2005. Recent developments in network measurement. In Carrington, P. J.; Scott, J.; and Wasserman, S., eds., *Models and methods in social network analysis*. Cambridge University Press.
- Rojers, D.; Vamplew, P.; Whiteson, S.; and Dazeley, R. 2013. A survey of multi-objective sequential decision-making. *Journal of Artificial Intelligence Research* 48.
- Valizadegan, H.; Jin, R.; and Wang, S. 2011. Learning to trade off between exploration and exploitation in multiclass bandit prediction. In *KDD*.
- Yadav, A.; Marcolino, L. S.; Rice, E.; Petering, R.; Winetrobe, H.; Rhoades, H.; Tambe, M.; and Carmichael, H. 2015. Preventing HIV spread in homeless populations using PSINET. In *IAAI*.



## A Appendix

### A.1 Results for each network

Since in Section 4 we presented our results across 4 different social network graphs, in this section we present those individually for each network, for the interested reader. Assuming a uniform model for the teaching lists, we can see the results for network A in Figures 8, 9, 10; for network B in Figures 11, 12, 13; for the Facebook network in Figures 14, 15, 16; and finally for the MySpace network in Figures 17, 18, 19. Assuming a power law model, we can see the results for network A in Figures 20, 21, 22; for network B in Figures 23, 24, 25; for the Facebook network in Figures 26, 27, 28; and finally for the MySpace network in Figures 29, 30, 31. As we can see, the results for each network show similar tendencies as the results across all networks presented in the main paper.




Figure 8: Results for network A across many interventions, for influence probability  $p = 0.5$ , teaching probability  $\varphi = 0.5$ , assuming uniform distribution.

### A.2 Additional results for power law distribution

In Section 4.2 we presented results for the power law distribution for  $a = 1.8$ . In this section, we present results for  $a = 1.2$  (that is, each neighbor has a 20% probability of being in a teaching list). We show in Figure 32 the result at each intervention for  $\varphi = 0.5$  and  $p = 0.5$ . In Figure 33, we show the AUC results for different parametrizations of  $p$  and  $\varphi$ . Finally, we present the regret for this case in Figure 34. As we can see, our main conclusions still hold for a different parametrization of  $a$ .




Figure 9: Results of Influence and Knowledge in network A for different teaching and influence probabilities, assuming uniform distribution.




Figure 10: Regret in network A for different teaching and influence probabilities, assuming uniform distribution.






Figure 11: Results for network B across many interventions, for influence probability  $p = 0.5$ , teaching probability  $\varphi = 0.5$ , assuming uniform distribution.




(a) Changing teaching probability




(b) Changing influence probability

Figure 12: Results of Influence and Knowledge in network B for different teaching and influence probabilities, assuming uniform distribution.



(a) Changing teaching probability



(b) Changing influence probability

Figure 15: Results of Influence and Knowledge in Facebook network for different teaching and influence probabilities, assuming uniform distribution.




Figure 13: Regret in network B for different teaching and influence probabilities, assuming uniform distribution.


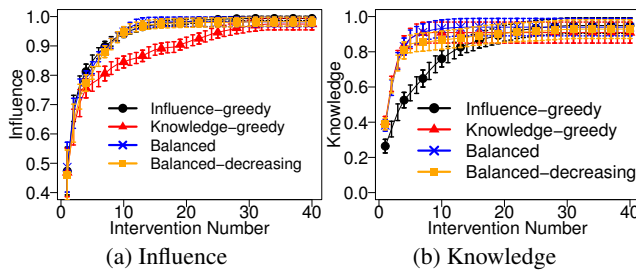



Figure 16: Regret in Facebook network for different teaching and influence probabilities, assuming uniform distribution.



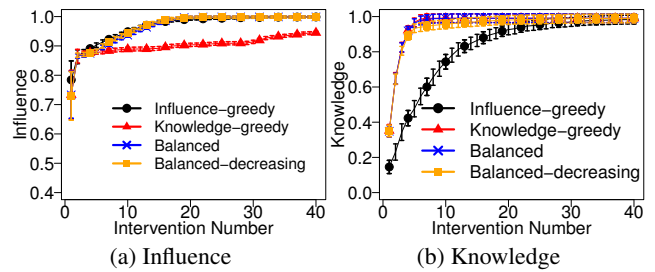
(a) Influence

(b) Knowledge




(c) Influence + Knowledge

Figure 14: Results for Facebook network across many interventions, for influence probability  $p = 0.5$ , teaching probability  $\varphi = 0.5$ , assuming uniform distribution.



(a) Influence

(b) Knowledge



(c) Influence + Knowledge

Figure 17: Results for MySpace network across many interventions, for influence probability  $p = 0.5$ , teaching probability  $\varphi = 0.5$ , assuming uniform distribution.





Figure 18: Results of Influence and Knowledge in MySpace network for different teaching and influence probabilities, assuming uniform distribution.




Figure 19: Regret in MySpace network for different teaching and influence probabilities, assuming uniform distribution.


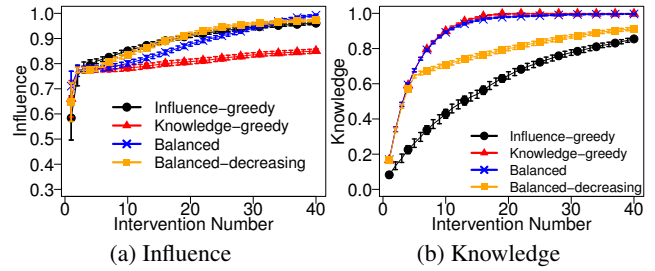


Figure 20: Results for network A across many interventions, for influence probability  $p = 0.5$ , teaching probability  $\varphi = 0.5$ , assuming power law distribution.





Figure 21: Results of Influence and Knowledge in network A for different teaching and influence probabilities, assuming power law distribution.




Figure 22: Regret in network A for different teaching and influence probabilities, assuming power law distribution.


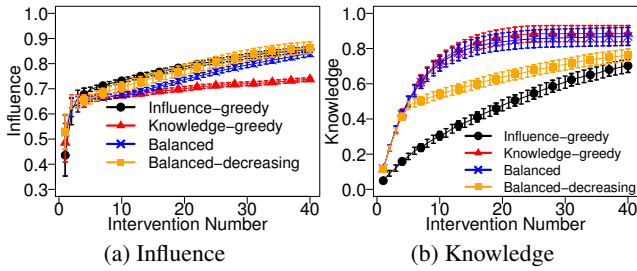


Figure 23: Results for network B across many interventions, for influence probability  $p = 0.5$ , teaching probability  $\varphi = 0.5$ , assuming power law distribution.


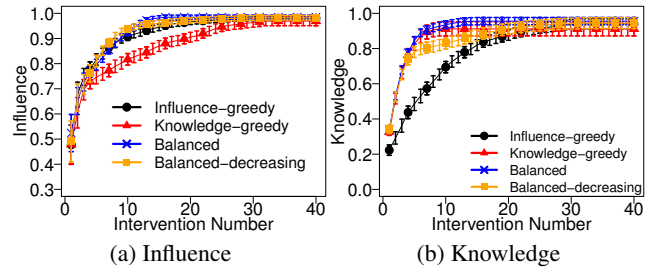


Figure 26: Results for Facebook network across many interventions, for influence probability  $p = 0.5$ , teaching probability  $\varphi = 0.5$ , assuming power law distribution.





Figure 24: Results of Influence and Knowledge in network B for different teaching and influence probabilities, assuming power law distribution.





Figure 27: Results of Influence and Knowledge in Facebook network for different teaching and influence probabilities, assuming power law distribution.




Figure 25: Regret in network B for different teaching and influence probabilities, assuming power law distribution.




Figure 28: Regret in Facebook network for different teaching and influence probabilities, assuming power law distribution.


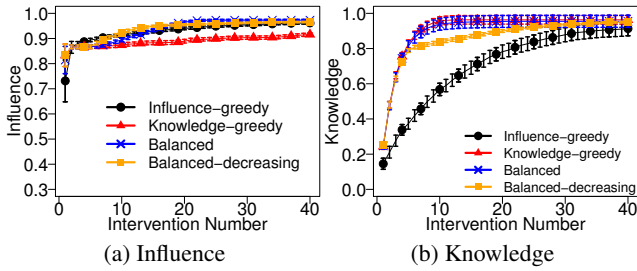


Figure 29: Results for MySpace network across many interventions, for influence probability  $p = 0.5$ , teaching probability  $\varphi = 0.5$ , assuming power law distribution.





Figure 30: Results of Influence and Knowledge in MySpace network for different teaching and influence probabilities, assuming power law distribution.




Figure 31: Regret in MySpace network for different teaching and influence probabilities, assuming power law distribution.


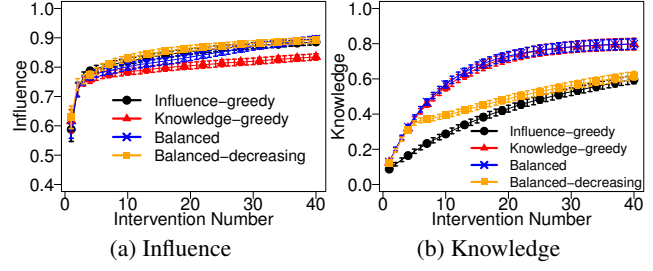


Figure 32: Results of 4 real world networks across many interventions, for  $p = 0.5$  and  $\varphi = 0.5$ , assuming power law distribution with  $a = 1.2$ .





Figure 33: Results of Influence and Knowledge for different teaching and influence probabilities, assuming power law distribution with  $a = 1.2$ .




Figure 34: Regret for different teaching and influence probabilities, assuming power law distribution with  $a = 1.2$ . Lower results are better.