

Robust, dynamic influence maximization

Bryan Wilder¹, Amulya Yadav¹, Nicole Immorlica², Eric Rice¹, Milind Tambe¹

¹ University of Southern California
{bwilder, amulyaya, ericr, tambe}@usc.edu
² Microsoft Research, New England
nicimm@gmail.com

Abstract. This paper focuses on new challenges in influence maximization inspired by non-profits' use of social networks to effect behavioral change in their target populations. Influence maximization is a multi-agent problem where the challenge is to select the most influential agents from a population connected by a social network. Specifically, our work is motivated by the problem of spreading messages about HIV prevention among homeless youth using their social network. We show how to compute solutions which are provably close to optimal when the parameters of the influence process are unknown. We then extend our algorithm to a dynamic setting where information about the network is revealed at each stage. Simulation experiments using real world networks collected by the homeless shelter show the advantages of our approach.

1 Introduction

Many behaviors are mediated by social influences, so a population's social network is essential to spreading a desired behavior. The influence maximization problem models a social network where a population of agents can influence each other via the links of the network. The challenge is to select the most influential nodes, a task with many applications. For instance, nonprofit organizations harness social networks to improve health in disadvantaged groups by conducting in-person social interventions. Such interventions have proven successful in many domains [17], including reducing HIV spread [13], improving nutrition [8], and reducing smoking [15]. In this paper, we use as an example the problem of spreading messages about HIV prevention among homeless youth. HIV/AIDS kills 10,000 people each year just in the U.S. [2], and the proportion of homeless youth who are HIV positive is over ten times that for populations with stable housing [12]. However, as is typical of nonprofits, homeless shelters only have sufficient resources to accommodate a few youth in each intervention [14]. They depend on word of mouth to expand the reach of their message.

While influence maximization has been well studied in online social networks, many domains (e.g., HIV prevention in homeless youth) raise three new challenges. First, organizations face *parameter uncertainty*. Models of influence spread [7] require numerous parameters which give the probability of influence spread along each edge. However, organizations acquire only a partial picture

of the network from field observations, Facebook contacts, etc., and there will inevitably be many links of uncertain strength [13]. Thus, we require algorithms which perform well with large uncertainty about the model parameters.

Second, the problem is *dynamic*. The organization (e.g., the homeless shelter) can ask participants about their immediate social circle during each intervention. Hence, updates about the true parameters are given at each stage, informing node selections for future interventions. We discuss the shortcomings of some related work which uses POMDPs for dynamic influence maximization in Section 3 [19, 18], but this related work has certainly failed to address the challenges of parameter and execution uncertainty.

Third, organizations face *execution uncertainty* because they cannot be sure that the youth (i.e., nodes in the social network) they select for an intervention will actually show up. This problem is particularly acute in our example domain because a number of events could prevent a youth’s participation. For instance, a youth might have been arrested or gone to stay with relatives.

While we illustrate these issues through health-related interventions, the underlying principles apply to many influence maximization domains. Take, for instance, a viral marketer who generates word of mouth by giving away a limited number of discounts. While the marketer likely will not know how influence propagates ahead of time, they can learn more during the campaign by surveying their users (e.g. “name five friends to receive this coupon”).

Overview of our approach: We start by addressing parameter uncertainty in the nondynamic setting by showing how to compute solutions whose value under *any* parameter setting is close to the optimum if the true parameters were known. We view the problem as zero sum game between the influencer and an adversary who chooses the worst case parameters, and present techniques which provably approximate the minimax strategies. The key challenge is that the algorithm’s strategy space is exponentially large, while the adversary’s is infinite. We give a polynomial time primal-dual algorithm, as well as DOSIM, a more practical algorithm based on a double oracle approach. DOSIM uses a greedy algorithm as the inner oracle for the influencer. DOSIM also handles execution uncertainty. We then extend these techniques to the dynamic setting by showing that DOSIM’s greedy oracle still gives excellent empirical performance. Our approach is evaluated on real world social networks of homeless youth in a large U.S. city. We show that it produces policies which perform near optimally regardless of the unknown parameters, while failing to consider this uncertainty results in substantially worse performance.

Our algorithm has been piloted by a homeless shelter in a major US city. A companion paper [20] shows results and detailed analysis from a pilot test of this and other algorithms at homeless shelters in a major US city, and shows the benefits of DOSIM in practical settings. A large-scale deployment of our algorithm with approximately 300 youth is scheduled to begin in Spring 2017.

2 Dynamic influence maximization

Problem overview: A social network is a directed graph $G = (V, E)$, with $|V| = n$ and $|E| = m$, where an edge from u to v indicates that u is a friend of v . An influencer does the following for T stages (where T is the total number of interventions). First, she selects a set of K nodes (participants) to attend the intervention and be influenced. Each node actually attends with an independent probability. Then, the influencer receives *observations* about the propagation probability on each edge outgoing from nodes that did attend, clarifying those nodes' immediate social circle. Finally, the influencer selects the next K nodes and the process repeats. The objective is to maximize the expected number of influenced nodes at the end of stage T .

Influence model: Influence propagates according to a variant of the standard independent cascade model [7]. In the independent cascade model, each edge $e = (u, v)$ is labeled with a probability p_e . The vector \mathbf{p} contains p_e for each $e \in E$. If u is influenced, it makes one attempt to influence v which succeeds with probability p_e . Each influenced node remains influenced for all stages. We use a variant where a node tries to influence its neighbor at each stage (until it succeeds), instead of just the first. This has been shown to better match empirical diffusion patterns [4].

It has been observed in many domains [1, 9] that the likelihood of influence spread depends largely on the kind of relationship. To model this phenomenon, each edge $e \in E$ has a type θ_e drawn from a set Θ . In our example, homeless shelter officials can annotate the edges as being either strong or weak ties. Each type $\theta \in \Theta$ is associated with a prior distribution over propagation probabilities and each edge e independently samples p_e from the prior for θ_e . This is a natural framework for our domain because organizations cannot provide accurately estimate each p_e , but they may classify the edges into types. We focus on uniform priors but our work can be extended to other distributions. We will assume a fixed width w for each uniform prior and overload notation to let θ_e give the center of the distribution. Hence, each edge samples $p_e \sim U[\theta_e - \frac{w}{2}, \theta_e + \frac{w}{2}]$, where U denotes a uniform distribution. The vector $\boldsymbol{\theta}$ contains θ_e for each $e \in E$.

Observation model: The influencer receives observations at each stage as she asks participants (nodes) about their immediate social circle. Importantly, it is not possible to ask participants for the exact p_e . Instead, the influencer only learns proxies for this probability, such as how frequently the participant interacts with the other person. We abstract this feedback as updates *relative* to the prior. For example, if $e = (u, v)$ and u and v interact more often than typical acquaintances, then we would infer that p_e is on the upper end of the prior θ_e . Accordingly, for each outgoing edge e , we assume that the algorithm observes the quantile of p_e relative to θ_e .

Example 1: If $\theta_e = 0.3$ and $w = 0.4$, then our prior distribution would be $p_e \sim U[0.1, 0.5]$. If frequent interaction was observed, then the influencer would infer that p_e lies in the upper quantile of the prior, and the updated posterior distribution would be $U[0.3, 0.5]$.

Objective formalization: Our algorithm starts with an input consisting of a graph G (with all nodes uninfluenced) and prior parameters θ . Unknown to the algorithm, the vector of edge propagation probabilities \mathbf{p} is drawn from the factorized prior described by θ . At each stage $t \in [1, T]$ the algorithm picks an action S_t from the set $\mathcal{A} = \{S \subset V \text{ s.t. } |S| = K\}$. Each node $v \in S_t$ attends with an independent probability q_v , which is known in advance. q_v models execution uncertainty. The algorithm then receives an observation O_t containing each outgoing edge from the subset of S_t that attended, which allows it to update its posterior distribution over the edge probabilities. The algorithm’s state is given by the actions taken and observations received so far, $\{(S_1, O_1) \dots (S_t, O_t)\}$. A policy π is a mapping from states to actions, and hence specifies the action to take in each stage of the problem. Let $f(\pi, \mathbf{p})$ be the expected number of nodes influenced by π by the end of stage T when the propagation probabilities are \mathbf{p} ³. f can be evaluated by averaging over random samples of the influence process [7, 16]. Our goal is to find a policy which maximizes $\mathbb{E}_{\mathbf{p} \sim \theta}[f(\pi, \mathbf{p})]$.

3 Uncertainty

Thus far, our model contains uncertainty about the realized value of the probability p_e given its expectation θ_e . This reflects variance in the distribution, and is captured via w (the width of the support of the prior over p_e). This uncertainty is important in the dynamic setting, because w controls the amount of information which can be gained via observations. For instance, if $w = 0$, then the prior for any edge is a point distribution and subsequent observations cannot reveal anything new. However, if $w = 0.4$, then the posterior mean on a given edge will shift by 0.1 if the algorithm observes whether p_e for that edge lies in the top or bottom quantile.

On its own, though, placing a distribution over p_e is inadequate in the following sense. As noted by He and Kempe [6], distributional uncertainty does not impact the average behavior of the model. This is because the following processes are analytically equivalent: (1) draw $p_e \sim U[\theta_e - \frac{w}{2}, \theta_e + \frac{w}{2}]$ and then propagate influence with probability p_e (2) propagate influence with probability θ_e . That is, in the classical influence maximization problem, placing a distribution over p_e is no different from setting p_e deterministically to its mean θ_e . Simply increasing the width w does not change this: the priors $p_e \sim U[0.4, 0.6]$ and $p_e \sim U[0.2, 0.8]$ both imply that edge e will propagate influence with probability $\mathbb{E}[p_e] = 0.5$. While distributional uncertainty does have a role in our dynamic problem because we allow the influencer to learn more about the realized values, it is not sufficient to capture uncertainty about the average likelihood that influence will spread. For instance, the influencer may know that p_e is drawn from the “strong tie” prior, but this will mean entirely different things if a strong tie has $\theta_e = 0.1$ than if $\theta_e = 0.9$. In domains such as HIV prevention, we will not know what θ_e should be set to because, e.g., we do not know the average probability that one friend will be able to convince another to get tested for HIV.

³ All of our techniques also apply to an infinite time horizon with discounted rewards.

We thus consider a second, higher-order uncertainty over the value of θ . We assume interval uncertainty, that is, $\theta \in [a_\theta, b_\theta] \forall \theta \in \Theta$. Since the influencer never learns the true value of θ , our goal is to produce an algorithm which simultaneously performs well for all possible values of θ . Robustness to this second kind of uncertainty is what was considered by Chen et al. [3] and He and Kempe [6]. The key difference from the first level of uncertainty is that we do not have a distribution over $[a_\theta, b_\theta]$. This leads to an adversarial model, requiring robustness to a worst-case choice of θ [6]. A worst-case approach is further motivated by the fact that organizations may not be able to quantify higher-order uncertainty, i.e., provide a prior over the prior parameters. *In short, we have a more expressive model to represent two levels of uncertainty.* The following example shows how the entire model works together, incorporating both levels of uncertainty and our observation model.

Example 2: Suppose that for some edge e , our interval uncertainty on θ gives $\theta_e \in [0.2, 0.8]$ and $w = 0.4$. If the actual value is $\theta_e = 0.2$, then this would imply that $p_e \sim U[0.0, 0.4]$. If instead the value is $\theta_e = 0.6$, then we would have $p_e \sim U[0.4, 0.8]$. Similar distributions are implied for every θ_e contained in our uncertainty interval $[0.2, 0.8]$. The influencer does not know which distribution is the true one. However, suppose she observes that p_e is in the upper quantile of the prior. Then in the case that $\theta_e = 0.2$, the updated posterior would be $p_e \sim U[0.2, 0.4]$, and in the case $\theta_e = 0.6$, the posterior would be $p_e \sim U[0.6, 0.8]$ (and so on). The influencer still does not know the true value of θ_e . But she now knows that (whatever it is) p_e lies in $U[\theta_e, \theta_e + \frac{w}{2}]$.

When our problem contains only the first level of uncertainty (about the realized value of p_e), we call this the *known parameter* case (as the key parameters θ are known). When the second level of uncertainty is also present, and we are hence uncertain about θ , we call this the *robust* case. Our objective for the robust case is formalized in Section 5.

4 Related Work

We discuss work on two related topics. First is work which deals with parameter uncertainty in classical influence maximization. Two recent papers [3, 6] analyze a similar model of parameter uncertainty, where the parameter for each edge is chosen adversarially within an interval. In their model, the edge parameters are totally independent. However, due to the domains we consider, we assume that edges belong to a small number of types. Therefore, our models are not directly comparable, and in particular the hardness result proved by He and Kempe [6] does not apply (we elaborate in Section 5). Further, neither of [3, 6] have a dynamic component, which is a key feature of our domain. Independently, several recent papers deal with learning influence networks [5, 11]. All of these papers assume that the algorithm observes many independent influence cascades, so our contributions are largely orthogonal to learning based approaches.

Second is work by Yadav et al. [19, 18] on applying dynamic influence maximization to HIV prevention. This is the most closely related work to ours, and

hence Section 7 empirically compares our algorithm and theirs. Yadav et al. formulate the problem as a POMDP. Our contributions extend their work in two ways. First, unlike their work, we address the robust case by giving an algorithm which is provably robust to parameter and execution uncertainty. Second, for known parameters, we show that dynamic influence maximization can be solved by a greedy algorithm instead of POMDP and that only greedy scales beyond small networks.

5 Robustness to parameter uncertainty

We start by showing how to obtain robust results under parameter uncertainty. For ease of explanation, we focus first on the nondynamic setting where the algorithm picks a single seed set in the first stage instead of a multistage policy. We also defer consideration of execution uncertainty and assume $q_v = 1 \forall v \in V$. When there is no parameter uncertainty, Kempe, Kleinberg and Tardos [7] showed that a greedy algorithm obtains a $(1 - 1/e)$ approximation to the optimal value because the objective function is submodular. We now show how to extend this approximation to a robust optimization over uncertain parameters.

Recall that we assumed interval uncertainty over the key parameters θ , that is, $\theta \in [a_\theta, b_\theta] \forall \theta \in \Theta$. The influencer’s goal is to find a seed set $S \in \mathcal{A}$ which leads to near-optimal influence spread no matter where each entry in θ lies within its interval. Our full model also includes lower-order uncertainty concerning the realized value of \mathbf{p} given θ . Since we focus for now on a single stage, the influencer does not receive observations about the realized \mathbf{p} . Thus (as noted in Section 3), the lower-order uncertainty is irrelevant because distributional uncertainty over \mathbf{p} can be removed via setting \mathbf{p} to its mean. This lower order uncertainty becomes relevant when we extend our algorithm to the dynamic setting in Section 6.

First, we formally define our objective. For each $S \in \mathcal{A}$, let $g_S(\theta)$ give the expected influence spread of seed set S with priors θ . Let the set of possible θ be $\mathcal{P} = \{\theta \mid \theta_e \in [a_{\theta_e}, b_{\theta_e}] \forall e \in E\}$. \mathcal{P} is the m -dimensional box where each θ_e may lie anywhere within the interval for that type. Let the optimal influence spread for each θ be $g_{OPT}(\theta) = \max_S g_S(\theta)$. The influencer’s utility for choosing seed set S when the parameters are θ is

$$R(S, \theta) = \frac{g_S(\theta)}{g_{OPT}(\theta)},$$

which quantifies the value that the influencer received compared to the optimum. This objective was also discussed by He and Kempe [6] and Chen et al. [3]. Since g_{OPT} is NP-hard to compute exactly, R cannot be efficiently evaluated. However, we can compute a greedy approximation g_G to the optimal value. Hence, we follow [3, 6] and instead optimize

$$R_G(S, \theta) = \frac{g_S(\theta)}{g_G(\theta)}.$$

This compares the influence spread of S to the value which would have been obtained by greedy. Note that $R(S, \theta) \leq R_G(S, \theta) \leq \frac{e}{e-1}R(S, \theta)$, so R_G is always within a constant factor of R . A randomized algorithm for the influencer is a vector $x \in \Delta^{|\mathcal{A}|}$ (where Δ^n denotes the n -dimensional probability simplex) giving the probability of selecting each seed set. As explained in Section 3, we take a worst-case approach to maximizing R_G . Hence, we seek the strategy x^* which maximizes the minimum value:

$$x^* = \arg \max_{x \in \Delta^{|\mathcal{A}|}} \min_{\theta \in \mathcal{P}} \sum_S x_S R_G(S, \theta). \quad (1)$$

This problem can be viewed as a zero-sum game between an influencer, who chooses the seed set $S \in \mathcal{A}$, and an adversary (nature), who chooses the parameters $\theta \in \mathcal{P}$. The influencer’s payoff is given by $R_G(S, \theta)$. Since we solve for a mixed strategy for the influencer, our algorithm is randomized: the influencer samples a pure strategy each time she wishes to choose a seed set. Randomization is necessary to guard against weakness to a particular set of parameters (a player in a game may be exploited if they commit to a deterministic strategy). However, it is possible that an influencer might prefer a pure strategy with a lower but “guaranteed” payoff to a probabilistic strategy which has a small chance of yielding very bad utility. In this case, it suffices to express the payoffs of our game through a risk aversion function and then solve the game normally using the algorithms described below; this will yield a desired guaranteed payoff. In practice, our proposed algorithm produces strategies with very sparse support; the few non-zero components all have comparable worst-case value so this problem does not arise.

We now introduce our techniques for solving the game. Normal methods do not apply because \mathcal{A} (the set of all sets of K nodes) is exponentially large, while \mathcal{P} (the set of all allowable parameter values) is infinite. However, we demonstrate that both sources of intractability can be overcome and present a constant-factor approximation to the minimax robust solution in polynomial time. Recall that a mixed strategy x is a vector giving the probability of choosing each seed set. Our approximation notion is as follows:

Definition 1. *An algorithm produces an (α, ϵ) -minimax robust solution if it always returns a mixed strategy x satisfying*

$$\min_{\theta} \sum_S x_S R_G(S, \theta) \geq \alpha \max_{x'} \min_{\theta} \sum_S x'_S R_G(S, \theta) - \epsilon.$$

We first tackle the adversary’s infinite strategy space by showing that this space can be reduced to a discrete set \mathcal{P}^* such that the adversary loses only an arbitrarily small ϵ by choosing from \mathcal{P}^* instead of \mathcal{P} . Recall that \mathcal{P} has infinite size because it is the continuous set of all allowable parameter values. For example, if we have types corresponding to strong and weak edges, \mathcal{P} could be the set of parameters which places the θ_e for strong edges anywhere in the continuous interval $[0.4, 0.8]$ and the θ_e for the weak edges anywhere in the

interval $[0.2, 0.4]$. There are no general approaches for solving games with infinite strategy spaces, so we show how to reduce our original game to one with a finite set of adversary strategies. The challenge is to ensure that only a small amount of approximation error is incurred in the reduction. Towards this goal, the main technical step is to prove that each g_S is sufficiently smooth for such a discretization to be possible. It is not immediately clear that g_S should always be smooth. For instance, in a complete graph, g_S rapidly increases from 0 to n even at a very small propagation probability because of the combinatorial number of potential paths (this is related to the emergence of a giant connected component in Erdős-Rényi graphs). However, Lemma 1, stated below, shows that g_S is smooth for any fixed size graph. Lemma 2 will then use Lemma 1 to show that there is a discretization \mathcal{P}^* where $|\mathcal{P}^*|$ scales reasonably.

Lemma 1. *For any $g_S, S \in \mathcal{A}$, (or g_G) and any $\theta_1, \theta_2 \in \mathcal{P}$, $|g_S(\theta_1) - g_S(\theta_2)| \leq nT \|\theta_1 - \theta_2\|_1$.*

Due to space restrictions, all proofs are deferred to the supplemental material. In Lemma 2, we use Lemma 1 to construct a \mathcal{P}^* which ensures that the loss incurred by discretizing is at most ϵ . The idea is as follows. Suppose we cover the hyperrectangle of allowable type values $\times_{\theta} [a_{\theta}, b_{\theta}]$ with a regular grid. Let our discretization $\mathcal{P}^* \subset \mathcal{P}$ be the set of θ which correspond to the grid points (by setting each θ_e to its value at the point). Then we have:

Lemma 2. *Fix $\epsilon > 0$, and construct \mathcal{P}^* using a grid with $\left(\frac{2nmT}{\epsilon}\right)^{|\Theta|}$ points. Then for any seed set $S \in \mathcal{A}$ and any point $\theta \in \mathcal{P}$ there is a $\theta^* \in \mathcal{P}^*$ satisfying $|R_G(S, \theta) - R_G(S, \theta^*)| \leq \epsilon$.*

Although $|\mathcal{P}^*|$ is exponential in $|\Theta|$, we assume that $|\Theta|$ is a small constant because organizations can only provide a small number of edge types. Moreover, this exponential dependence is likely unavoidable since when $|\Theta| = m$, we recover a model for which He and Kempe [6] showed even approximating the minimax robust strategy is NP-hard. By reducing the adversary's pure strategy space to a finite set \mathcal{P}^* , Lemma 2 allows us to formulate the minimax robustness problem as the following LP:

$$\begin{array}{ll}
 \textbf{Primal:} & \textbf{Dual:} \\
 \max_{x \in \mathbb{R}^{|\mathcal{A}|}} U \text{ s.t.} & \min_{y \in \mathbb{R}^{|\mathcal{P}^*|}} W \text{ s.t.} \\
 \sum_{S \in \mathcal{A}} x_S = 1 & \sum_{\theta \in \mathcal{P}^*} y_{\theta} = 1 \\
 x_S \geq 0 \quad \forall S \in \mathcal{A} & y_{\theta} \geq 0 \quad \forall \theta \in \mathcal{P}^* \\
 U \leq \sum_{S \in \mathcal{A}} x_S R_G(S, \theta) \quad \forall \theta & W \geq \sum_{\theta \in \mathcal{P}^*} y_{\theta} R_G(S, \theta) \quad \forall S \quad (2)
 \end{array}$$

The primal is intractable due to the exponential number of variables. However, since \mathcal{P}^* has size $\text{poly}(n, m, \frac{1}{\epsilon})$, we can instead work with the dual. Although the dual has exponentially many constraints, all that we need to ensure

polynomial time solvability is a separation oracle which checks if Constraint 2 is violated for any S . This reduces to finding the "tightest" constraint; i.e., the seed set S which maximizes the expected value of R_G under the adversary mixed strategy y . We want to compute

$$\max_{S \in \mathcal{A}} \sum_{\theta \in \mathcal{P}^*} y_\theta R_G(S, \theta) = \max_{S \in \mathcal{A}} \sum_{\theta \in \mathcal{P}^*} \frac{y_\theta}{g_G(\theta)} g_S(\theta)$$

Note that the term $\frac{y_\theta}{g_G(\theta)}$ is independent of the seed set S which we choose; it is just a nonnegative coefficient. Further, computing $\max_{S \in \mathcal{A}} g_S(\theta)$ is the problem of finding the most influential seed set for the fixed parameters θ , which is a submodular maximization problem. Since a nonnegative linear combination of submodular functions is itself submodular, greedy can be used to find an approximately tightest constraint:

Lemma 3. *For any adversary mixed strategy $y \in \Delta^{|\mathcal{P}^*|}$, running greedy with the objective $\max_{S \in \mathcal{A}} \sum_{\theta \in \mathcal{P}^*} \frac{y_\theta}{g_G(\theta)} \mathbb{E}_{\mathbf{p} \sim \theta} [f(\cdot, \mathbf{p})]$ produces a $(1 - 1/e)$ -approximate best response to y .*

This is exactly what we need to check Constraint 2 since a particular setting of y corresponds to an adversary mixed strategy, so Constraint 2 is violated if and only if the best response to y has expected R_G higher than W . Although greedy may not always detect violated constraints, we show that the ellipsoid algorithm can still be used to approximately solve the dual, and that we can then recover an approximate solution to the primal.

Theorem 1. *There is an algorithm which obtains an $(1 - 1/e, \epsilon)$ -minimax robust solution to the robust influence maximization problem in time $\text{poly}(m, n, \frac{1}{\epsilon})$.*

A practical algorithm: Theorem 1 accomplishes our goal of providing a polynomial time approximation algorithm. However, this algorithm is likely impractical since the ellipsoid algorithm is known to perform poorly. Moreover, the strategy returned could be quite large (although polynomial). While we view the construction of a polynomial time approximation algorithm as an important theoretical contribution, we also present a practical algorithm using a double oracle approach [10]. DOSIM (Double Oracle for Social Influence Maximization) solves the game using best response oracles for each player. Initially, both players choose only from a small set of arbitrarily chosen pure strategies (lines 1-2 of Algorithm 1). At each iteration, we compute the minimax mixed strategies assuming the game is restricted to just the current strategy set (line 5). The payoff matrix of this restricted game is given by $R_G(S, \theta)$ for each pair of strategies $S \in \mathcal{A}_i, \theta \in \mathcal{P}_i$ contained in the current strategy sets. Then, we find the best response of each player to their opponent's mixed strategy (lines 6-7). If either of these responses is not in the corresponding strategy set, the new entry is added and the algorithm proceeds. If both strategies are present, the algorithm terminates with a provably optimal solution. While double oracle is well known [10], it

Algorithm 1 DOSIM

```
1:  $\mathcal{A}_0 =$  arbitrary seed set
2:  $\mathcal{P}_0 = \{(\frac{1}{2}, \frac{1}{2}, \dots, \frac{1}{2})\}$ 
3:  $i = 0$ 
4: while  $\mathcal{A}_i \neq \mathcal{A}_{i-1}$  or  $\mathcal{P}_i \neq \mathcal{P}_{i-1}$  do
5:    $x_i, y_i =$  SolveLP( $\mathcal{A}_i, \mathcal{P}_i$ ) //Solve restricted game
6:    $\mathcal{A}_{i+1} = \mathcal{A}_i \cup \{\text{INFLUENCERBESTRESPONSE}(y_i)\}$ 
7:    $\mathcal{P}_{i+1} = \mathcal{P}_i \cup \{\text{ADVERSARYBESTRESPONSE}(x_i, \epsilon)\}$ 
8:    $i = i + 1$ 
9: end while
10: return SolveLP( $\mathcal{A}_i, \mathcal{P}_i$ )
```

is not an out-of-the-box approach. Our contribution is constructing appropriate oracles for robust influence maximization:

Influencer oracle: Run the greedy algorithm on the objective given by the current adversary strategy y . Lemma 3 shows that greedy is a $(1 - 1/e)$ -approximate oracle.

Adversary oracle: Search \mathcal{P}^* and choose the parameters $\theta \in \mathcal{P}^*$ which minimize $\sum_{S \in \mathcal{A}_{i-1}} x_S R_G(S, \theta)$. By Lemma 2, this approximates the adversary's best response to within an additive ϵ .

Hence, we can provide an approximation guarantee for the best response oracle for both players. McMahan et al. [10] showed that when double oracle is given an *optimal* best response oracle for each player, the strategies it returns form an equilibrium of the game. We now generalize this reasoning to show that *approximate* best response oracles yield an approximate equilibrium. We say that mixed strategies x and y form a c -approximate Nash equilibrium if neither player can gain more than a multiplicative factor c by deviating.

Theorem 2. *If double oracle is given an α -approximation to the best response of Player 1 and a β -approximation to the best response of Player 2, on termination it returns strategies which form a $\max\{\frac{1}{\alpha}, \frac{1}{\beta}\}$ -approximate Nash equilibrium.*

By taking $\beta = 1$ and accounting for the additive loss of ϵ introduced by discretizing the adversary's strategy space, we get:

Corollary 1. *DOSIM obtains a $(1 - 1/e, \epsilon)$ -minimax robust solution to the robust influence maximization problem.*

DOSIM may run in exponential time (since it may need to add all pure strategies in the worst case). However, we see in Section 7 that it converges quickly in practice and so Corollary 1 provides guaranteed solution quality.

Robustness to execution uncertainty: Recall that each node v is present with probability q_v . To illustrate the implication for the optimal strategy, suppose that there are two very influential nodes u and v with $q_u = q_v = q$ whose neighborhoods entirely overlap. Clearly, we should only select one of the two if q is high. However, if q is low, we might select both to ensure that we reach

their large set of neighbors. We create a new graph G' to represent this tradeoff as follows. For each node $v \in V$, add a new “designer” node v' , which has a single edge (v', v) . This edge has a new type θ_d whose prior outputs $p_{v',v} = 1$ with probability $\theta_d = q_v$ and $p_{v',v} = 0$ otherwise. If actions are restricted to the designer nodes, i.e., $\mathcal{A} = \{\{v'_1, v'_2, \dots, v'_K\} | v_i \in V\}$, then maximizing influence in G' exactly corresponds to maximizing influence under execution errors in G . We can also incorporate our higher-level uncertainty by allowing θ_d to be adversarially chosen within an interval. Therefore, we can compute a seed set which is robust to an adversarially chosen θ_d by running DOSIM on G' .

6 DOSIM in the dynamic setting

We now extend DOSIM to the dynamic setting, where we seek a multistage policy instead of a seed set. Recall that a policy selects a set of nodes at each stage conditioned on observations from earlier stages. We will show how to find a policy which is robust to uncertainty over the prior θ . When we seek a robust policy, observations give a quantile of the prior but do not fully specify the posterior since the exact prior is unknown. E.g., if $\theta_e \in [0.2, 0.7]$ and $w = 0.2$, the prior could be either $U[0.2, 0.4]$ or $U[0.5, 0.7]$ and an observation of “top half” implies a posterior of either $U[0.3, 0.4]$ or $U[0.6, 0.7]$. Our algorithm uses this information while remaining robust to the higher-level uncertainty over θ . Recall that the lower-level uncertainty cannot be replaced by setting $p_e = \theta_e$ anymore because the algorithm receives observations about the realized value.

DOSIM translates naturally into the dynamic setting. Specifically, our scheme for discretizing the adversary’s strategy space (Lemmas 1 and 2) generalizes to the case where the influencer chooses a multistage policy instead of a single seed set (see supplemental material for proof). Thus, our adversary oracle, and its theoretical guarantees, are unchanged. Theorem 2 applies to any zero-sum game, and hence implies that if we can supply an influencer oracle for the dynamic setting, then running DOSIM with this oracle will produce robust *policies*.

Unfortunately, supplying an influencer oracle with guaranteed approximation ratio is difficult; there are no known theoretical guarantees for dynamic influence maximization because the objective function is not adaptive submodular. Yadav et al. [18] proved this for their DIME model, which our model generalizes. Our approach is to use a heuristic oracle and show experimentally that it produces good influence spread. This is sufficient to ensure that we obtain robust results on real world problems since DOSIM performs well on a particular network so long as its influencer oracle does:

Lemma 4. *If the influencer oracle achieves an α -approximation for any $\theta \in \mathcal{P}^*$ on a given G , then DOSIM provides an (α, ϵ) -minimax robust solution on G .*

Essentially, we can be sure that DOSIM’s solution is robust if the oracle performs well on our particular network since DOSIM preserves the quality of its input oracle. We now turn to finding an appropriate heuristic for the influencer oracle. One natural idea would be to use the current state of the art for

dynamic influence maximization, the POMDP based heuristic algorithm HEAL [18]. However, HEAL is unsuitable for two reasons. First, our model is more general than HEAL’s. Specifically, HEAL cannot represent the p_e as being drawn from an interval. Second, we show experimentally that HEAL is not scalable and is hence unsuitable as a subroutine which will be called many times by DOSIM. Hence, we use an alternative heuristic oracle, presented on below.

Algorithm 2 Dynamic greedy

```

1:  $S_{\text{prev}} = \emptyset$  //nodes selected in previous stages
2:  $O = \emptyset$  //set of all observations
3: for  $t \leftarrow 1$  to  $T$  do
4:    $S_t = \emptyset$  //nodes selected in this stage
5:   for  $i \leftarrow 1$  to  $K$  do
6:      $v^* = \arg \max_{\mathbf{p}} \mathbb{E}_{\mathbf{p}}[f(S_{\text{prev}} \cup S_t \cup \{v\}, \mathbf{p})$ 
7:        $- f(S_{\text{prev}} \cup S_t, \mathbf{p}) | O]$ 
8:      $S_t = S_t \cup \{v^*\}$ 
9:   end for
10:  Influence  $S_t$ ; Update  $S_{\text{prev}} = S_{\text{prev}} \cup S_t$ 
11:  Receive observation  $O_t$ ; Update  $O = O \cup O_t$ 
12: end for

```

When examining Algorithm 1, our extension of DOSIM to the dynamic setting only needs to supply an influencer oracle (line 6). The rest of DOSIM is completely unchanged because we can use the same adversary oracle. For the influencer oracle, which best responds to an adversary mixed strategy y , we will use a dynamic version of the greedy influence maximization algorithm, described in Algorithm 2. The dynamic greedy algorithm operates greedily on two layers. In each stage (line 3), it chooses the set of K nodes which maximize the immediate gain in influence spread, conditioned on the observations gathered thus far. However, since finding this set is essentially a classical influence maximization problem (and hence NP-hard), we use a further greedy selection process to approximate the optimal set of K nodes for each stage. Line 5 selects the node v^* which maximizes the expected marginal gain. The expectation is over parameters θ sampled from the adversary mixed strategy y , and propagation probabilities \mathbf{p} sampled according to θ . After identifying these K nodes, the algorithm marks them as influenced and simulates new observations (lines 10-11). We can repeatedly simulate different sets of observations (repeating line 10) to build up the best response policy.

7 Experiments

We present experimental results for our example domain, preventing HIV spread in homeless youth. We use two datasets (Network A and Network B) collected by homeless shelters using surveys and interviews [18]. Both have 140-170 nodes




Fig. 1: Influence with varying K on Network A (left) and Network B (right).




Fig. 2: Runtime with varying T on Network A (left) and WS Networks (right).

and 300-400 edges. We also show results on artificial Watts-Strogatz networks (model parameters $p = 0.1$, $k = 7$, all results averaged over 30 networks), since these mirror our datasets' small diameter.

Dynamic influence maximization with known parameters:

The most closely related work is the HEAL algorithm [18], so we begin by setting up a comparison between DOSIM and HEAL. HEAL only handles the known parameter case, where only the lower order uncertainty over the realized value of \mathbf{p} is present. Thus, we run DOSIM assuming that the adversary strategy is fixed to a single, known θ . Our results validate the even when only lower-level uncertainty is present, DOSIM still outperforms the closest competitor.

Our model is more general than HEAL's, so we align the models as follows. There is one type of "certain" edges with fixed $p_e = p$, and one type of "uncertain" edges which have $p_e = p$ with probability u and $p_e = 0$ with probability $1 - u$. We give DOSIM the same observations as HEAL by revealing whether an edge is in the top $100 \cdot u\%$ of the prior or bottom $100 \cdot (1 - u)\%$. Figure 1 shows the influence spread of both algorithms on the two real networks as we vary the intervention size K (fixing $T = 5$, $p = 0.6$, $u = 0.1$ as in [18]). We also compare to picking nodes with the highest degree centrality (DC), since this is standard in health policy [17]. The maximum intervention size is $K = 6$ since the shelter cannot accommodate more participants. We see that DOSIM and HEAL perform very similarly, while DC performs poorly (and is omitted hereafter). Table 1 shows the percentage gain of DOSIM over HEAL as the parameters u and p are varied (for $K = 2$, $T = 10$) for Network A (Network B is similar). DOSIM performs no more than 4% worse, and up to 25% better. Overall, the influence obtained by DOSIM is comparable to HEAL, and sometimes better.

However, only DOSIM is scalable. Figures 2 shows runtime on Network A and Watts-Strogatz networks. DOSIM performs better on larger graphs. For $T = 10$ on Network A, it is roughly 2.5 times faster. On Watts-Strogatz networks with 500 nodes, DOSIM finishes in about 8 minutes while we cut off HEAL's execution after 5 hours. For networks larger than 500 nodes, HEAL runs out of memory. Both algorithms were run on a 2.33 GHz Intel processor with 48 GB of RAM.

Robustness: We now investigate the performance of the full DOSIM algorithm when we have interval uncertainty over θ . Shelter officials annotated

p $\backslash u$	0.1	0.2	0.3
0.5	-0.1	15.0	6.0
0.6	6.4	13.3	24.9
0.7	8.9	21.0	-3.4

Table 1: % gain over HEAL with varying parameters




Fig. 3: Robustness comparison.

the edges as being either strong or weak ties, so we have two corresponding types. We add a third type for the designer edges reflecting execution uncertainty. A uniform prior is adversarially chosen for the first two types, using width $w = 0.4$ (truncated if the support would exceed $[0,1]$). We allow any mean $\theta_{\text{strong}} \in [0.2, 0.8]$ and $\theta_{\text{weak}} \in [0, 0.4]$. These ranges contain the parameters used in previous work [18], but are large enough to represent our considerable uncertainty. We allow any $\theta_d \in [0.2, 0.8]$. Observations specified whether outgoing edges were in the bottom or top quantile of the prior. $K = 4$ nodes were selected in each of $T = 5$ stages with $\epsilon = 0.1$. We compare DOSIM to two benchmarks. Since no previous work (including HEAL) handles both levels of uncertainty simultaneously, our benchmarks run DOSIM with a fixed adversary strategy. “DOSIM-fixed” uses fixed parameter settings reported previously ($\theta_{\text{strong}} = 0.6$, $\theta_{\text{weak}} = 0.05$) without considering execution uncertainty ($\theta_d = 1$). “DOSIM-SE” uses the same settings for θ_{strong} and θ_{weak} , but also accounts for stochastic execution uncertainty by setting $\theta_d = 0.5$ (based on shelter officials’ experience). These two algorithms assume a single allowable point value for θ . Figure 3 shows the worst case ratio R_G (over all allowable parameter values) on the y axis for DOSIM and the benchmarks. DOSIM obtains approximately 90% of the optimal value across all parameter settings while using fixed parameters leads to substantially worse performance (about 72% in the worst case). Incorporating execution uncertainty results in a value roughly halfway in between. Lastly, we note that on both datasets, DOSIM converges rapidly (within 13 iterations).

8 Conclusion

We address dynamic influence maximization under uncertainty about both the network parameters and the efficacy of interventions. First, we give algorithms with provable guarantees for minimax robustness under unknown parameters. Second, our algorithm handles execution uncertainties. Third, we extend these results to the dynamic setting using an experimentally validated greedy algorithm. Lastly, experiments on real world networks collected from homeless youth demonstrate our approach’s advantages.

References

1. J. Brown and P. Reingen. Social ties and word-of-mouth referral behavior. *Journal of Consumer Research*, 14(3):350–362, 1987.
2. CDC. HIV surveillance report. 26, 2014.
3. W. Chen, T. Lin, Z. Tan, M. Zhao, and X. Zhou. Robust influence maximization. *CoRR*, abs/1601.06551, 2016.
4. J.-P. Cointet and C. Roth. How realistic should knowledge diffusion models be? *Journal of Artificial Societies and Social Simulation*, 10(3):5, 2007.
5. M. Gomez Rodriguez, J. Leskovec, and A. Krause. Inferring networks of diffusion and influence. In *KDD*, 2010.
6. X. He and D. Kempe. Robust influence maximization. *CoRR*, abs/1602.05240, 2016.
7. D. Kempe, J. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *KDD*, 2003.
8. D. Kim et al. Social network targeting to maximise population behaviour change: a cluster randomised controlled trial. *The Lancet*, 386(9989):145–153, 2015.
9. D. McAdam and R. Paulsen. Specifying the relationship between social ties and activism. *American Journal of Sociology*, 99(3):640–667, 1993.
10. H. B. McMahan, G. J. Gordon, and A. Blum. Planning in the presence of cost functions controlled by an adversary. In *ICML*, pages 536–543, 2003.
11. P. Netrapalli and S. Sanghavi. Learning the graph of epidemic cascades. *SIGMETRICS Perform. Eval. Rev.*, 40(1):211–222, June 2012.
12. N.H. Council. HIV/AIDS among persons experiencing homelessness: Risk factors, predictors of testing, and promising testing strategies. 2012.
13. E. Rice. The positive role of social networks and social networking technology in the condom-using behaviors of homeless young people. *Public Health Reports*, pages 588–595, 2010.
14. E. Rice, A. Fulginiti, H. Winetrobe, J. Montoya, A. Plant, and T. Kordic. Sexuality and homelessness in Los Angeles public schools. *American Journal of Public Health*, 102(2):200a–3201, 2012.
15. F. Starkey, S. Audrey, J. Holliday, L. Moore, and R. Campbell. Identifying influential young people to undertake effective peer-led health promotion. *Health Education Research*, 2009.
16. Y. Tang, X. Xiao, and Y. Shi. Influence maximization: Near-optimal time complexity meets practical efficiency. In *Proceedings of the 2014 ACM SIGMOD international conference on Management of data*, pages 75–86. ACM, 2014.
17. T. W. Valente. Network interventions. *Science*, 337(6090):49–53, 2012.
18. A. Yadav, H. Chan, A. Jiang, H. Xu, E. Rice, and M. Tambe. Using social networks to aid homeless shelters: Dynamic influence maximization under uncertainty. In *AAMAS*, 2016.
19. A. Yadav, L. Marcolino, E. Rice, R. Petering, H. Winetrobe, H. Rhoades, M. Tambe, and H. Carmichael. Preventing HIV spread in homeless populations using PSINET. In *IAAI*, 2015.
20. A. Yadav, B. Wilder, E. Rice, R. Petering, J. Craddock, A. Yoshioka-Maxwell, M. Hemler, L. Onasch-Vera, M. Tambe, and D. Woo. Influence maximization in the field: The arduous journey from emerging to deployed application. In *AAMAS*, 2017.