

Preventing HIV Spread in Homeless Populations Using PSINET

A. Yadav, L. Marcolino, E. Rice, R. Petering, H. Winetrobe, H. Rhoades, M. Tambe

University of Southern California, Los Angeles, CA, 90089
 {amulyaya, sorianom, ericr, petering, hwinetro, hrhoades, tambe}@usc.edu

H. Carmichael

LCSW, Executive Director, My Friend's Place, Los Angeles, CA 90028
 hcarmichael@myfriendsplace.org

Abstract

Homeless youth are prone to HIV due to their engagement in high risk behavior. Many agencies conduct interventions to educate/train a select group of homeless youth about HIV prevention practices and rely on word-of-mouth spread of information through their social network. Previous work in strategic selection of intervention participants does not handle uncertainties in the social network's structure and in the evolving network state, potentially causing significant shortcomings in spread of information. Thus, we developed PSINET, a decision support system to aid the agencies in this task. PSINET includes the following key novelties: (i) it handles uncertainties in network structure and evolving network state; (ii) it addresses these uncertainties by using POMDPs in influence maximization; (iii) it provides algorithmic advances to allow high quality approximate solutions for such POMDPs. Simulations show that PSINET achieves $\sim 60\%$ more information spread over the current state-of-the-art. PSINET was developed in collaboration with My Friend's Place (a drop-in agency serving homeless youth in Los Angeles) and is currently being reviewed by their officials.

1 Introduction

Homelessness affects ~ 2 million youths in USA annually, 11% of whom are HIV positive (10 times the infection rate in the general population) (NAHC 2011). Peer-led HIV prevention programs such as POL (Kelly et al. 1997) try to spread HIV prevention information through network ties and recommend selecting intervention participants based on Degree Centrality (i.e., highest degree nodes first). Such peer-led programs are highly desirable to agencies working with homeless youth as these youth have minimal health care access and are distrustful of adults (Rice and Rhoades 2013).

Agencies working with homeless youth prefer a series of small size interventions deployed sequentially as they have limited manpower to direct towards these programs. This fact and emotional/behavioral problems of youth makes managing groups of more than 5-6 youth at a time very difficult (Rice et al. 2012). Strategically choosing intervention participants is important so that information percolates through their social network in the most efficient way.

Copyright © 2015, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

This paper introduces PSINET (POMDP based Social Interventions in Networks for Enhanced HIV Testing), a novel Partially Observable Markov Decision Process (POMDP) based system which chooses the participants of successive interventions in a social network. The key novelty of our work is a unique combination of POMDPs and influence maximization to handle uncertainties about (i) friendships between people in the social network; and (ii) evolution of the network state in between two successive interventions. Traditionally, influence maximization has not dealt with these uncertainties, which greatly complicates the process of choosing intervention participants. Moreover, this problem is a very good fit for POMDPs as (i) we conduct several interventions sequentially, similar to sequential actions taken in a POMDP; and (ii) we must handle uncertainty over network structure and evolving state.

Unfortunately, our POMDP's state (2^{300} states) and action spaces ($\binom{150}{10}$ actions) are beyond the reach of current state-of-the-art POMDP solvers/algorithms. To address this scale-up challenge, PSINET provides a novel on-line algorithm, that relies on the following key ideas: (i) compact representation of transition probabilities to manage the intractable state and action spaces; (ii) combination of the QMDP heuristic with Monte-Carlo simulations to avoid exhaustive search of the entire belief space; and (iii) voting on multiple POMDP solutions, each of which efficiently searches a portion of the solution space to improve accuracy.

Our work is done in collaboration with My Friend's Place¹, a non-profit agency assisting Los Angeles's homeless youth to build self-sufficient lives by providing education/support to reduce high-risk behavior. Thus, we evaluate PSINET on real social networks of youth frequenting this agency. This work is being reviewed by officials at My Friend's Place towards final deployment.

2 Related work

There are two primary areas of related work that we discuss in this section. First, we discuss work in influence maximization. Kempe, Kleinberg, and Tardos (2003) provided a constant-ratio approximation algorithm to find 'seed' sets of nodes to optimally spread influence in a graph. This was followed by many speed up techniques (Kimura and Saito

¹See <http://myfriendsplace.org/>

2006; Chen, Wang, and Wang 2010). All these algorithms assume no uncertainty in the network structure and select a single seed set. In contrast, we select several seed sets sequentially in our work to select intervention participants. Also, our problem takes into account uncertainty about the network structure and evolving network state. Golovin and Krause (2010) introduced adaptive submodularity and discussed adaptive sequential selection (similar to our work) in viral marketing. However, unlike our work, they assume no uncertainty in network structure and state evolution.

The second field of related work is online planning in POMDPs, since offline planning approaches are unable to scale up to problems of interest in our work (Smith 2013). We focus on the literature on Monte-Carlo (MC) sampling based online POMDP solvers since that sub-field is most related to our work. Silver and Veness (2010) proposed POMCP algorithm that uses Monte-Carlo tree search in online planning. Somani et al. (2013) improved the worst case performance of POMCP in DESPOT algorithm. These two algorithms maintain a search tree for all sampled histories to find best actions, which may lead to better solution qualities, but it makes the algorithm less scalable (as we show in our experiments). Therefore, our algorithm does not maintain a search tree and uses the QMDP heuristic to find best actions.

3 Our Approach

A POMDP is a tuple $\varphi = \langle \mathbf{S}, \mathbf{A}, \mathbf{O}, \mathbf{T}, \Omega, \mathbf{R} \rangle$, where \mathbf{S} , \mathbf{A} and \mathbf{O} are sets of possible world states, actions and observations respectively; $\mathbf{T}(s, a, s')$ is the transition probability of reaching s' by taking action a in s , $\Omega(o, a, s')$ is the observation probability of observing o , by taking action a to reach s' ; and $\mathbf{R}(s, a)$ is the reward of taking action a in s .

A POMDP policy Π maps every possible belief state b to an action $a = \Pi(b)$. Our aim is to find an optimal policy $\Pi^* = \arg \max_{\pi} P^{\pi}$ (given an initial belief b_0), which maximizes the expected long term reward $P^{\Pi} = \sum_{t=1..H} E[R(s^t, a^t)]$ where H is the horizon. Computing optimal policies offline for finite horizon POMDPs is PSPACE-Complete. Thus, focus has recently turned towards online algorithms, which only find the best action for the current belief state. Upon reaching a new belief state, online planning again plans for this new belief. Thus, online planning interleaves planning and execution at every time step.

POMDP Model of our Domain

In describing our model, we first outline the homeless youth social network and then map it onto our POMDP. The social network of homeless youth is a digraph $G = (V, E)$ with $|V| = n$. Every $v \in V$ represents a homeless youth, and every $\{e = (a, b) | a, b \in V\} \in E$ represents that youth a has nominated youth b in their social circle. Further, $E = E_c \cup E_u$, where $E_c (|E_c| = l)$ is the set of certain edges, i.e., friendships which we are certain about. Conversely, $E_u (|E_u| = m)$ is the set of uncertain edges, i.e., friendships which we are uncertain about. For example, youth may describe their friends “vaguely”, which is not enough for accurate identification (Rice et al. 2012). In this case, there would be uncertain edges from the youth to each of his “suspected” friends.

Each uncertain edge ($e \in E_u$) exists with an *existence probability* $u(e)$, the exact value of which is determined from domain experts. For example, if it is uncertain whether node B is node A’s friend, then $u(A, B) = 0.5$ (say) implies that B is A’s friend with a 0.5 chance. Accounting for these uncertain edges is important as our node selection might depend heavily on whether these edges exist with certainty or not. We call this graph G an “uncertain graph” henceforth. Figure 1 shows an uncertain graph on 6 nodes (A to F) and 7 edges. The dashed/solid edges represent uncertain/certain edges respectively.




Figure 1: Graph G

In our work, we use the independent cascade model, a well studied influence propagation model (Kimura and Saito 2006). In this model, every node $v \in V$ has an h-value, where $h : V \rightarrow \{0, 1\}$. $h(v) = 1$ and 0 determines whether a node is influenced or not, respectively. Nodes only change their h-value (from 0 to 1) when they first get influenced. If node $v \in V$ gets influenced at time step t , it influences each of its 1-hop un-influenced neighbors with a *propagation probability* $p(e) \forall e \in E$ for all future time steps. Moreover, every edge $e \in E_u$ has an f-value (which represents a sampled instance of $u(e)$) and is unknown apriori, where $f : E_u \rightarrow \{0, 1\}$. $f(e) = 1$ and 0 determines whether the uncertain edge exists with certainty in the real graph or not, respectively. For $e \in E_u$, the influence probability (given by $p(e) * u(e)$) is contingent on the edge’s actual existence.

Recall that we need a policy for selecting nodes for successive interventions in order to maximize the influence spread in the network. Nodes selected for interventions are assumed to be influenced ($h(v) = 1$) post-intervention with certainty. However, there is uncertainty in how the h-value of the unselected nodes changes in between successive interventions. For example, in Figure 1, if we choose nodes B and D for the 1st intervention, we are uncertain whether nodes C and E (adjacent to nodes B and D) are influenced before nodes for the 2nd intervention are chosen. We now provide a POMDP mapping onto our problem.

States Consider strict total orders $<_v$ and $<_u$ on the sets V and E_u respectively. A state $S = \langle H, F \rangle$ is a 2-tuple. $H = \langle h(v_1), h(v_2), \dots, h(v_i), \dots \rangle \forall i \in 1..n$ is a binary tuple representing the h-values of all nodes (ordered by $<_v$). Also, $F = \langle f(e_1), f(e_2), \dots, f(e_i), \dots \rangle \forall i \in 1..m$ is a binary tuple representing the f-values of all uncertain edges (ordered by $<_u$) in the graph. Our POMDP has 2^{n+m} states.

Actions Every $\alpha \subset V$ s.t. $|\alpha| = k$ (k is the number of nodes selected per intervention) is a POMDP action. For example, in Figure 1, one possible action is $\{A, B\}$ (assuming $k = 2$). Our POMDP has $\binom{n}{k}$ actions.

Observations We assume that we can “observe” the f-values of uncertain edges outgoing from the nodes chosen in an action. This translates to asking intervention participants about their 1-hop social circles, which is within the agency’s capacity (Rice et al. 2012). For example, by taking action $\{B, C\}$ in Figure 1, the f-values of edge 4 and 5 (i.e., uncertain edges in the 1-hop social circle of nodes B and C)

would be observed. Consider $\Theta(\alpha) = \{e \mid e = (a,b) \text{ s.t. } a \in \alpha \wedge e \in E_u\} \forall \alpha \in A$, which represents the ordered tuple of uncertain edges that are observed when the agency takes action α . Then, our POMDP observation upon taking action α is defined as $o(\alpha) = \langle f(e_1), f(e_2), \dots, f(e_i) \rangle \forall e_i \in \Theta(\alpha)$, i.e., the f-values of the observed uncertain edges.

Transition Probabilities Consider states $s = \langle H, F \rangle$ and $s' = \langle H', F' \rangle$ and action $\alpha \in A$. In order for $T(s, \alpha, s')$ to be non zero, first we require that $F'[i] = F[i] \forall i \text{ s.t. } e_i \notin \Theta(\alpha)$ i.e., uncertain edges which were not observed will not change their f-values. Second, $H'[i] = H[i] \forall i \text{ s.t. } H[i] = 1$. Finally, $H'[i] = 1 \forall i \text{ s.t. } v_i \in \alpha$. The last two conditions imply that all nodes with $h(v) = 1$ in state s , and all nodes selected by action α will remain influenced in final state s' . If any of these three conditions is not true, then $T(s, \alpha, s') = 0$. Otherwise, a heuristic method to calculate $T(s, \alpha, s')$ is described next (as accurate calculation needs to consider all possible paths in a graph through which influence could spread, which is $O(n!)$ in the worst case).

Transition Probability Heuristic Consider a weighted adjacency matrix for graph G_σ (created from graph G) s.t.

$$G_\sigma(i, j) = \begin{cases} 1 & \text{if } (i, j) \in E_c \wedge (H[i] = 1 \vee \alpha[i] = 1) \\ u(i, j) & \text{if } (i, j) \in E_u \wedge (H[i] = 1 \vee \alpha[i] = 1) \\ 0 & \text{if otherwise.} \end{cases} \quad (1)$$

G_σ is a *pruned* graph which contains only edges outgoing from influenced nodes. We prune the graph because influence can only spread through edges which are outgoing from influenced nodes. Note that G_σ does not consider influence spreading along a path consisting of more than one uninfluenced node, as this event is highly unlikely in the limited time in between successive interventions. However, nodes connected to a chain (of arbitrary length) of influenced nodes get influenced more easily due to reinforced efforts of all influenced nodes in the chain. We use G_σ to construct a diffusion vector \mathbf{D} , the i^{th} element of which gives us a measure of the probability of the i^{th} node to get influenced. This diffusion vector \mathbf{D} is then used to estimate $T(s, \alpha, s')$.

A known result states that if G is a graph's adjacency matrix, then $G^r(i, j)$ ($G^r = G$ multiplied r times) gives the number of paths of length r between nodes i and j (Diestel 2005). Additionally, note that if all edges e_i in a path of length r have different propagation probabilities $p(e_i) \forall i \in [1, r]$, the probability of influence spreading between two nodes connected through this path of length r is $\prod_{i=1}^r p(e_i)$. For simplicity, we assume the same $p(e) \forall e \in E$; hence, the probability of influence spreading becomes p^r . Using these results, we construct diffusion vector \mathbf{D} :

$$\mathbf{D}(\mathbf{p}, \mathbf{T})_{n \times 1} = \sum_{t \in [1, \mathbf{T}]} \left((\mathbf{p} \overline{\mathbf{G}}_\sigma)^t * \mathbf{1}_{n \times 1} \right) \quad (2)$$

Here, $\mathbf{D}(\mathbf{p}, \mathbf{T})$ is a column vector of size $n \times 1$, \mathbf{p} is the constant propagation probability on the edges, \mathbf{T} is a variable parameter that measures number of hops considered for influence spread (higher values of \mathbf{T} yields more accurate $\mathbf{D}(\mathbf{p}, \mathbf{T})$ but increases the runtime²), $\mathbf{1}_{n \times 1}$ is a $n \times 1$ column

²<https://www.dropbox.com/s/ar4l8eavihq6mwm/appendix.pdf> provides details/proofs.

vector of 1's and $\overline{\mathbf{G}}_\sigma$ is the transpose of \mathbf{G}_σ . This formulation is similar to diffusion centrality (Banerjee et al. 2013) where they calculate influencing power of nodes. However, we calculate power of nodes to get influenced (by using $\overline{\mathbf{G}}_\sigma$).

Proposition 1. \mathbf{D}_i , the i^{th} element of $\mathbf{D}(\mathbf{p}, \mathbf{T})_{n \times 1}$, upon normalization, gives an approximate probability of the i^{th} graph node to get influenced in the next round.²

Consider the set $\Delta = \{i \mid H'[i] = 1 \wedge H[i] = 0 \wedge \alpha[i] = 0\}$, which represents nodes which were uninfluenced in the initial state s ($H[i] = 0$) and which were not selected in the action ($\alpha[i] = 0$), but got influenced by other nodes in the final state s' ($H'[i] = 1$). Similarly, consider the set $\Phi = \{j \mid H'[j] = 0 \wedge H[j] = 0 \wedge \alpha[j] = 0\}$, which represents nodes which were not influenced even in the final state s' ($H'[j] = 0$). Using \mathbf{D}_i values, we can now calculate $T(s, \alpha, s') = \prod_{i \in \Delta} \mathbf{D}_i \prod_{j \in \Phi} (1 - \mathbf{D}_j)$, i.e., we multiply influence probabilities \mathbf{D}_i for nodes which are influenced in state s' , along with probabilities of not getting influenced ($1 - \mathbf{D}_j$) for nodes which are not influenced in state s' .

Observation Probabilities Given action $\alpha \in A$ and final state $s' = \langle H', F' \rangle$, there exists an observation $o(\alpha, s')$, which is uniquely determined by both α and s' . More formally, $o(\alpha, s') = \{F'[i] \forall e_i \in \Theta(\alpha)\}$. Therefore, $\Omega(o, \alpha, s') = 1$ if $o = o(\alpha, s')$ and 0 otherwise.

Rewards The reward of taking action $\alpha \in A$ in state $s = \langle H, F \rangle$ is given as $R(s, \alpha) = \sum_{s' \in S} T(s, \alpha, s') (\|s'\| - \|s\|)$, where $\|s'\|$ is the number of influenced nodes in s' . This gives the expected number of new influenced nodes.

PSINET

Initial experiments with the ZMDP solver (Smith 2013) showed that state-of-the-art offline POMDP planners ran out of memory on 10 node graphs. Thus, we focused on online planning algorithms and tried using POMCP (Silver and Veness 2010), a state-of-the-art online POMDP solver which relies on Monte-Carlo (MC) tree search and rollout strategies to come up with solutions quickly. However, it keeps the entire search tree over sampled histories in memory, disabling scale-up to the problems of interest in this paper. Hence, we propose a MC based online planner that utilizes the QMDP heuristic and eliminates this search tree.

POMDP black box simulator MC sampling based planners approximate the value function for a belief by the average value of n (say) MC simulations starting from states sampled from the current belief state. Such approaches depend on a POMDP black box simulator $\Gamma(s_t, \alpha_t) \sim (s_{t+1}, o_{t+1}, r_{t+1})$ which generates the state, observation and reward at time $t + 1$, given the state and action at time t , in accordance with the POMDP dynamics. In Γ , each edge $e \in \Theta(\alpha_t)$ is sampled according to $u(e)$ to generate \mathbf{o}_{t+1} . Similarly, $\mathbf{r}_{t+1} = \|\mathbf{s}_{t+1}\| - \|\mathbf{s}_t\|$, i.e., the number of new influenced nodes in s_{t+1} . To generate s_{t+1} , consider $s_{t+1} = \langle H', F' \rangle$ and $s_t = \langle H, F \rangle$. First, $\mathbf{D}(\mathbf{p}, \mathbf{T})_{n \times 1}$ is normalized to get probabilities of nodes getting influenced. Let $K = \{H[i] = 1 \vee \alpha_t[i] = 1\}$ represent the set of nodes which are certainly influenced. Then, $H'[i] = 1 \forall i \in K$ and for all other i , $H'[i]$ is sampled according to $\mathbf{D}(\mathbf{p}, \mathbf{T})_{n \times 1}[i]$. Also, $F'[i] = F[i] \forall i \notin \Theta(\alpha_t)$ and $F'[i] = o_{t+1}[i] \forall i \in$

Algorithm 1: PSINET

Input: Belief state β , Uncertain graph G

Output: Best Action κ

- 1 Sample graph to get Δ different instances;
 - 2 **for** $\delta \in \Delta$ **do**
 - 3 \lfloor $FindBestAction(\delta, \alpha_\delta, \beta)$;
 - 4 $\kappa = VoteForBestAction(\Delta, \alpha)$
 - 5 $UpdateBeliefState(\kappa, \beta)$;
 - 6 **return** κ ;
-

$\Theta(\alpha_t)$. Note that s_{t+1} calculated this way represents a state sampled according to $T(s_t, \alpha_t, s_{t+1})$.

QMDP It is a well known approximate offline planner, and it relies on $Q(s, a)$ values, which represents the value of taking action a in state s . It precomputes these $Q(s, a)$ values for every (s, a) pair by approximating them by the future expected reward obtainable if the environment is fully observable (Littman, Cassandra, and Kaelbling 1995). Finally, QMDP’s approximate policy Π is given by $\Pi(b) = \arg \max_a \sum_s Q(s, a)b(s)$ for belief b . Our intractable POMDP state/action spaces make it infeasible to calculate $Q(s, a) \forall (s, a)$. Thus, we propose to use a MC sampling based online variant of QMDP in PSINET.

Algorithm Flow Algorithm 1 shows the flow of PSINET. In Step 1, we randomly sample all $e \in E_u$ in G (according to $u(e)$) to get Δ different graph instances. Each of these instances is a different POMDP as the h-values of nodes are still partially observable. Since each of these instances fixes $f(e) \forall e \in E_u$, the belief β is represented as an un-weighted particle filter where each particle is a tuple of h-values of all nodes. This belief is shared across all instantiated POMDPs. For every graph instance $\delta \in \Delta$, we find the best action α_δ in graph δ , for the current belief β in step 3. In step 4, we find the best action κ for belief β , over all $\delta \in \Delta$ by voting amongst all the actions chosen by $\delta \in \Delta$. Then, in step 5, we update the belief state based on the chosen action κ and the current belief β . PSINET can again be used to find the best action for this or any future updated belief states. We now detail the steps in Algorithm 1.

Sampling Graphs In Step 1, we randomly keep or remove uncertain edges to create one graph instance. As a single instance might not represent the real network well, we instantiate the graph Δ times and use each of these instances to vote for the best action to be taken.

FindBestAction Step 3 uses Algorithm 2, which finds the best action for a single network instance, and works similarly for all instances. For each instance, we find the action which maximizes long term rewards averaged across n (we use $n = 2^8$) MC simulations starting from states (particles) sampled from the current belief β . Each MC simulation samples a particle from β and chooses an action to take (choice of action is explained later). Then, upon taking this action, we follow a uniform random rollout policy (until either termination, i.e., all nodes get influenced, or the horizon is breached) to find the long term reward, which we get by taking the “selected” action. This reward from each MC simulation is analogous to a $Q(s, a)$ estimate. Finally, we pick the action with the maximum average reward.

Multi-Armed Bandit We can only calculate $Q(s, a)$ for a select set of actions (due to our intractable action space). To choose these actions, we use a UCT implementation of a multi-armed bandit to select actions, with each bandit arm being one possible action. Every time we sample a new state from the belief, we run UCT, which returns the action which maximizes this quantity: $\Upsilon(s, a) = Q_{MC}(s, a) + c_0 \sqrt{\frac{\log N(s)}{N(s, a)}}$. Here, $Q_{MC}(s, a)$ is the running average of $Q(s, a)$ values across all MC simulations run so far. $N(s)$ is number of times state s has been sampled from the belief. $N(s, a)$ is number of times action a has been chosen in state s and c_0 is a constant which determines the exploration-exploitation tradeoff for UCT. High c_0 values make UCT choose rarely tried actions more frequently, and low c_0 values make UCT select actions having high $Q_{MC}(s, a)$ to get an even better $Q(s, a)$ estimate. Thus, in every MC simulation, UCT strategically chooses which action to take, after which we run the rollout policy to get the long term reward.

Voting Mechanisms In Step 4, each network instance votes for the best action (found using Step 3) for the uncertain graph and the action with the highest votes is chosen. We propose three different voting schemes:

(a)**PSINET-S** Each instance’s vote gets equal weight.

(b)**PSINET-W** Every instance’s vote gets weighted differently. The instance which removes x uncertain edges has a vote weight of $W(x) = x \forall x \leq m/2$ and $W(x) = m - x \forall x > m/2$. This weighting scheme approximates the probabilities of occurrences of real world events by giving low weights to instances which removes either too few or too many uncertain edges, since those events are less likely to occur. Instances which remove $m/2$ uncertain edges get the highest weight, since that event is most likely.

(c)**PSINET-C** Given a ranking over actions from each instance, the Copeland rule makes pairwise comparisons among all actions, and picks the one preferred by a majority of instances over the highest number of other actions (Pomerol and Barba-Romero 2000). Algorithm 2 is run D times for each instance to generate a partial ranking.

Belief State Update Recall that every MC simulation samples a particle from the belief, after which UCT chooses an action. Upon taking this action, some random state (particle) is reached using the transition probability heuristic. This particle is stored, indexed by the action taken to reach it. Finally, when all simulations are done, corresponding to every action α that was tried during the simulations, there will be a set of particles that were encountered when we took action α in that belief. The particle set corresponding to the action that we finally choose, forms our next belief state.

4 Experimental Evaluation

We provide two sets of results. First, we show results on artificial networks to understand our algorithms’ properties on abstract settings, and to gain insights on a range of networks. Next, we show results on the two real world homeless youth networks that we had access to. In all experiments, we select 2 nodes per round and average over 20 runs, unless otherwise stated. PSINET-S and W use 20 network instances and PSINET-C uses 5 network instances (each in-

Algorithm 2: FindBestAction

Input: Graph instance δ , belief β , N simulations**Output:** Best Action α_δ

- 1 Initialize *counter* = 0;
 - 2 **while** *counter* ++ < N **do**
 - 3 $s = \text{SampleStartStateFromBelief}(\beta)$;
 - 4 $a = \text{UCT_MultiArmedBandit}(s)$;
 - 5 $\{s', r\} = \text{SimulateRolloutPolicy}(s, a)$;
 - 6 $\alpha_\delta =$ action with max average reward;
 - 7 **return** α_δ ;
-

stance finds its best action 5 times) in all experiments, unless otherwise stated. The propagation and existence probability values were set to 0.5 in all experiments (based on findings by (Kelly et al. 1997)), although we relax this assumption later in the section. In this section, a $\langle X, Y, Z \rangle$ network refers to a network with X nodes, Y certain and Z uncertain edges. We use a metric of “indirect influence spread” (IIS) throughout this section, which is number of nodes “indirectly” influenced by intervention participants. For example, on a 30 node network, by selecting 2 nodes each for 10 interventions (horizon), 20 nodes (a lower bound for any strategy) are influenced with certainty. However, the total number of influenced nodes might be 26 (say) and thus, the IIS is 6. *All comparison results are statistically significant under bootstrap- t ($\alpha = 0.05$).*




Figure 2: Comparison on BTER graphs

Artificial networks First, we compare all algorithms on Block Two-Level Erdos-Renyi (BTER) networks (having degree distribution $X_d \propto d^{-1.2}$, where X_d is number of nodes of degree d) of several sizes, as they accurately capture observable properties of real-world social networks (Seshadhri, Kolda, and Pinar 2012). Figures 2(a) and 2(b) show solution quality and runtimes (respectively) of Degree Centrality (DC) (which selects nodes based on their out-degrees, and $e \in E_u$ add $u(e)$ to node degrees), POMCP and PSINET-(S,W and C). We choose DC as our baseline as it is the current modus operandi of agencies working with homeless youth. X-axis is number of network nodes and Y-axis shows IIS across varying horizons (number of interventions) in Figure 2(a) and log of runtime (in seconds) (Figure 2(b)).

Figure 2(a) shows that all POMDP based algorithms beat DC by $\sim 60\%$, which shows the value of our POMDP model. Further, it shows that PSINET-W beats PSINET-(S and C). Also, *POMCP runs out of memory on 30 node graphs*. Figure 2(b) shows that DC runs quickest (as expected) and all PSINET variants run in almost the same time. Thus, Figures 2(a) and 2(b) tell us that while DC runs quickest, it provides

the worst solutions. Amongst the POMDP based algorithms, PSINET-W is the best algorithm that can provide good solutions and can scale up as well. Surprisingly, PSINET-C performs worse than PSINET-(W and S) in terms of solution quality. Thus, we now focus on PSINET-W.




Figure 3: Increasing number of graph instances

Having shown the impact of POMDPs, we analyze the impact of increasing network instances (which implies increasing number of votes in our algorithm) on PSINET-W. Figures 3(a) and 3(b) show solution quality and runtime respectively of PSINET-W with increasing network instances, for a $\langle 40, 71, 41 \rangle$ BTER network with a horizon of 10. X-axis is number of network instances and Y-axis shows IIS (Figure 3(a)) and runtime (in seconds) (Figure 3(b)). These figures show that increasing the number of instances increases IIS as well as runtime. Thus, a solution quality-runtime tradeoff exists, which depends on the number of network instances. Greater number of instances results in better solutions and slower runtimes and vice versa. However, for 30 vs 70 instances, the gain in solution quality is $< 5\%$ whereas the runtime is $\sim 2X$, which shows that increasing instances beyond 30 yields marginal returns.




Figure 4: Comparison of DC with PSINET-W

Next, we relax our assumptions about propagation ($p(e)$) probabilities, which were set to 0.5 so far. Figure 4(a) shows the solution quality, when PSINET-W and DC are solved with different $p(e)$ values respectively, for a $\langle 40, 71, 41 \rangle$ BTER network with a horizon of 10. X-axis shows $p(e)$ and Y-axis shows IIS. This figure shows that varying $p(e)$ minimally impacts PSINET-W’s improvement over DC, which shows our algorithms’ robustness to these probability values (We get similar results upon changing $u(e)$). In Figure 4(b), we show solution qualities of PSINET-W and DC on a $\langle 30, 31, 27 \rangle$ BTER network (horizon=3) and vary number of nodes selected per round (K). X-axis shows increasing K , and Y-axis shows IIS. This figure shows that even for a small horizon of length 3, which does not give many chances for influence to spread, PSINET-W significantly beats DC with

increasing K .




Figure 5: Real world networks

Real world networks Figure 5(a) compares PSINET variants and DC (horizon = 30) on two real-world social networks (created by our collaborators through surveys and interviews of homeless youth frequenting My Friend’s Place’s) of homeless youth (each of size $\sim \langle 155, 120, 190 \rangle$). X-axis shows the two networks and Y-axis shows IIS. This figure clearly shows that all PSINET variants beat DC on both real world networks by $\sim 60\%$, which shows that PSINET works equally well on real-world networks. Also, PSINET-W beats PSINET-S, in accordance with previous results. Above all, this signifies that we could improve the quality and efficiency of HIV based interventions over the current modus operandi of agencies by $\sim 60\%$.

We now differentiate between the kinds of nodes selected by DC and PSINET-W for the sample BTER network in Figure 5(b), which contains nodes segregated into four clusters (C1 to C4), and node degrees in a cluster are almost equal. C1 is biggest, with slightly higher node degrees than other clusters, followed by C2, C3 and C4. DC would first select all nodes in cluster C1, then all nodes in C2 and so on. Selecting all nodes in a cluster is not “smart”, since selecting just a few cluster nodes influences all other nodes. PSINET-W realizes this by looking ahead and spreads more influence by picking nodes in different clusters each time. For example, assuming $K=2$, PSINET-W picks one node in both C1 and C2, then one node in both C1 and C4, etc.

5 Implementation Challenges & Conclusion

A few implementation challenges need to be solved when PSINET gets deployed by agencies working with homeless youth. A computer-based selection procedure for intervention attendees may raise suspicions about invasion of privacy for these youth. Public awareness campaigns in the agencies working with this program would help overcome such fears and to encourage participation. There is also a secondary issue about protection of privacy for the involved youth. Agencies working with homeless youth collect information on their clients, most of which is not to be shared with third parties, such as researchers etc. We propose deploying PSINET as a software package that the agencies could use without providing identifying information to our team.

This paper presents PSINET, a POMDP based decision support system to select homeless youth for HIV based interventions. Previous work in strategic selection of intervention participants does not handle uncertainties in the social network’s structure and evolving network state, potentially

causing significant shortcomings in spread of information. PSINET has the following key novelties: (i) it handles uncertainties in network structure and evolving network state; (ii) it addresses these uncertainties by using POMDPs in influence maximization; and (iii) it provides algorithmic advances to allow high quality approximate solutions for such POMDPs. Simulations show that PSINET achieves $\sim 60\%$ improvement over the current state-of-the-art. PSINET was developed in collaboration with My Friend’s Place and is currently being reviewed by their officials.

6 Acknowledgements

This research was supported by MURI Grant W911NF-11-1-0332. Rice, Rhoades, Winetrobe and Petering were supported by NIMH Grant number R01-MH093336.

References

- Banerjee, A.; Chandrasekhar, A. G.; Duflo, E.; and Jackson, M. O. 2013. The diffusion of microfinance. *Science* 341(6144).
- Chen, W.; Wang, C.; and Wang, Y. 2010. Scalable influence maximization for prevalent viral marketing in large-scale social networks. SIGKDD.
- Diestel, R. 2005. Graph theory. 2005. *Grad. Texts in Math.*
- Golovin, D., and Krause, A. 2010. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *NIPS*.
- Kelly, J. A.; Murphy, D. A.; Sikkema, K. J.; McAuliffe, T. L.; Roffman, R. A.; Solomon, L. J.; Winett, R. A.; and Kalichman, S. C. 1997. Randomised, controlled, community-level hiv-prevention intervention for sexual-risk behaviour among homosexual men in us cities. *The Lancet* 350(9090):1500.
- Kempe, D.; Kleinberg, J.; and Tardos, É. 2003. Maximizing the spread of influence through a social network. In *PKDD*.
- Kimura, M., and Saito, K. 2006. Tractable models for information diffusion in social networks. In *PKDD*.
- Littman, M. L.; Cassandra, A. R.; and Kaelbling, L. P. 1995. Learning policies for partially observable environments: Scaling up. Citeseer.
- NAHC. 2011. The link between Homelessness and HIV.
- Pomerol, J.-C., and Barba-Romero, S. 2000. *Multicriterion decision in management: principles and practice*. Springer.
- Rice, E., and Rhoades, H. 2013. How should network-based prevention for homeless youth be implemented? *Addiction* 108(9):1625.
- Rice, E.; Tulbert, E.; Cederbaum, J.; Adhikari, A. B.; and Milburn, N. G. 2012. Mobilizing homeless youth for HIV prevention: a social network analysis of the acceptability of a face-to-face and online social networking intervention. *Health education research* 27(2):226.
- Seshadhri, C.; Kolda, T. G.; and Pinar, A. 2012. Community structure and scale-free collections of erdős-rényi graphs. *Physical Review E* 85(5):056109.
- Silver, D., and Veness, J. 2010. Monte-carlo planning in large pomdps. *NIPS*.
- Smith, T. 2013. Zmdp software for pomdp/mdp planning.
- Somani, A.; Ye, N.; Hsu, D.; and Lee, W. S. 2013. Despot: Online pomdp planning with regularization. *NIPS*.