Proceedings of the Thirtieth International Conference on Automated Planning and Scheduling (ICAPS 2020)

We Mind Your Well-Being: Preventing Depression in Uncertain Social Networks by Sequential Interventions

Aye Phyu Phyu Aung,[1,2] **Xinrun Wang,**[1] **Bo An,**[1] **Xiaoli Li**[2]
1. School of Computer Science and Engineering, Nanyang Technological University, Singapore
2. Institute for Infocomm Research, Agency for Science, Technology and Research (A*STAR), Singapore
{ayep0001, xwang033, boan}@ntu.edu.sg, xlli@i2r.a-star.edu.sg

Abstract

Mental health has become a major concern according to WHO who estimates that more than 350 million people worldwide are affected by depression. Studies have shown that interventions and social support can reduce stress and depression. However, counselling centers do not have enough resources to provide counselling and social support to all the participants in their interest. This paper helps social support organizations (e.g., university counselling centers) sequentially select the participants for interventions. Unfortunately, previous works do not consider emotion propagation from other neighbours of the influencees and initial uncertainties of mental states and influence. Moreover, they fail to scale up to solve problems with a large number of participants due to the huge state space. Our contributions in this paper are fourfold. Firstly, we propose a new model that addresses the sequential intervention of participants while considering the propagation of emotions and formulate it as a Partially Observable Markov Decision Process (POMDP) to handle uncertainties about their mental states and the influence between them. Secondly, we apply reasoning to refine belief to improve solution quality for the lack of initial information on mental state values. Thirdly, we improve the scalability by the abstraction of states to reduce the number of states by representing the mental states with an abstracted discrete set. We further improve the scalability by multi-level partitioning to get smaller POMDPs. Finally, we conduct extensive experiments on both synthetic and real networks to show that our algorithm significantly improves scalability with comparable solution quality compared to the state-of-the-art algorithms.

1 Introduction

Nowadays, depression and other mood disorders have become major public health concerns worldwide. The World Health Organization (WHO) estimates that 350 million people are affected by depression throughout the world, reducing their ability to work and socialize, as well as increasing the rate of mortality from suicides (Eyre et al. 2017). Since counselling and social support help mitigate this problem by reducing people's stress (Rafferty and Griffin 2006), counselling services are emerging to provide interventions where a counsellor conducts dialogue sessions with several participants, finds out about their mental states and provides therapy. In this work, we use University Counselling Center (UCC) as an example. UCC is set up to monitor the stress level and well-being of students and decrease their stress level by interventions.

However, there are several challenges faced. Firstly, according to 2016 AUCCCD Survey, only $6-7\%$ of the students seek for counselling. Although we need to invite all the students for counselling to provide more effective social support, there are some limitations in the capacity of intervention such as the availability of counsellors in a university. Secondly, UCC cannot obtain the complete information about the students' mental states and the relationship between them at the beginning of the intervention. Thirdly, the large number of students makes it difficult to scale up to the real-world networks. Hence, it is not easy for UCC to efficiently decide the intervention plan.

There are existing works to deal with such kind of problems using sequential planning algorithms. PSINET (Yadav et al. 2015), HEAL (Yadav et al. 2016) and CAIMS (Yadav et al. 2018) maximize the HIV information spread in uncertain networks by formulating as POMDPs addressing the uncertainty of the influence between the participants. DOSIM (Wilder et al. 2017) formulates the problem as a zero-sum game between the influencer and the adversary (uncertainty). However, these algorithms cannot be directly extended to solve our problem due to two main issues. Firstly, they do not consider the initial uncertainty of mental state values which leads to poor solution quality. Secondly, the extensions of these algorithms fail to scale up to realistic networks due to the huge state space of our problem.

This paper makes four key contributions. Firstly, we propose a novel model for intervention to prevent depression in an uncertain network. The model considers the different values of nodes and the propagation which is not only affected by the influencer but also by influencee's neighbours. We formulate the problem as POMDP to address the uncertainties of mental state values and influence which we observe along the interventions. Secondly, we propose a reasoning algorithm on the students' mental states upon observation to refine the belief of the POMDP. Refining the belief with reasoning reduces the loss of solution quality by the initial uncertainty of students' mental states. Thirdly, we propose MLPRAP (Multi-Level Partition algorithm with Reasoning and Abstracted Planning) with the following novelties: (i) we abstract the POMDP states to reduce the state space; (ii) we provide theoretical bounds on uniform networks for solution quality loss due to abstraction; (iii) we partition the graph into smaller partitions by two folds: (1) balanced partitioning (MLP-B); (2) cluster partitioning (MLP-C) so that the algorithm scales up to plan interventions for a network of at least 1000 students. Finally, we provide extensive experimental analysis on scalability and solution quality of ML-PRAP compared to the state-of-the-art algorithms.

2 Motivation Scenario

Though our model can be applied to many scenarios such as mental health care for a public sector and counselling services for employees, we motivate our model by a specific case where UCC conducts a series of interventions to students. At each round of intervention, UCC invites a group of students for the counselling session. Through the intervention, UCC knows the mental states of intervened students and students within one degree of them (Rice et al. 2012), i.e., the influence between the counselled students and their friends. Since the counsellors have student registry and the collected information in previous interventions, UCC has initial students' mental states and influence estimation.

There has been evidence shown by the studies that emotions (happiness or stress) can spread from person to person via emotion propagation (Fowler and Christakis 2008; Eyre et al. 2017). Therefore, we construct the emotion propagation model where a person's stress (mental state) is reduced after the intervention, after which she spreads her happiness through emotion propagation in the network and reduces the stress levels of her neighbours. This propagation is one-degree from the seed node since the influence propagation does not normally go beyond that in real-world networks (Goel, Watts, and Goldstein 2012). Since the neighbours' mental states affect a person's mental state both positively and negatively (Rafferty and Griffin 2006), we considered the happy/ stressed emotions of each neighbour in the propagation model. We assume that the mental states of intervened students are reduced with certainty considering that the sudden external factors would not arise while being monitored during intervention (UCLA 2018).

3 Related Works

Students' Stress and Risk of Depression. Many different factors can lead to depression such as genetics, medication, physical or substance abuse and stress (Helmers et al. 1997). Among them, stress (feeling of frustration, anger, nervousness) is a significant factor for a high risk of depression and anxiety, esp. for university students (Blackmore et al. 2007; Lucassen et al. 2006; Khan and Khan 2017). Hence, we aim to reduce stress levels to reduce the risk of depression.

Influence Maximization. Influence maximization on social networks has been modelled as Independent Cascade (IC) or Linear Threshold (LT) models (Kempe, Kleinberg, and Tardos 2003). There are uncertainties in social networks such as uncertainty in the edge influence probability and uncertainty in the initial values of the nodes. There are two lines of works that address such uncertainty: (1) IC is extended to choose the seed set to optimally spread influence in a graph (Chen et al. 2016); (2) the influencer selects several seed sets sequentially to intervene the participants and the influencer receives observations about the participants' immediate social circles. To sequentially select the intervention participants, PSINET and HEAL formulate the problem as POMDP to handle the uncertainties of connection existence that are observed in each intervention (Yadav et al. 2015; 2016). As another approach, DOSIM formulates the problem as a zero-sum game between the algorithm which picks the optimal policy and the adversary (nature) which selects the connection probabilities with uncertainty (Wilder et al. 2017). However, all adopt IC model and only consider the spread from the influencer but not the effect of influencee's neighbours. Moreover, they only consider the node values as binary and set initial values as 0, i.e., they do not consider uncertainty on the initial values of the nodes. However, HEAL cannot scale up to realistic networks due to huge state and action sets in our model. Hence, we propose abstraction and multi-level partitioning to improve the scalability.

Online POMDP Solvers. We focus on online algorithms for solving POMDPs that are more scalable and suitable for our problem with huge state and action sets than offline algorithms. Monte Carlo Sampling-based online solvers, POMCP (Silver and Veness 2010) and DESPOT (Ye et al. 2017), use Monte Carlo tree search and maintain a search tree for all sampled histories. Thus, they have better solution quality but reduce scalability. HEAL does not maintain a search tree and uses QMDP heuristics (Littman, Cassandra, and Kaelbling 1995) to find the best action.

4 Model

We consider Q rounds of UCC's interventions of a group of students and at each intervention, UCC selects K students. For each intervention, UCC obtains the observations about mental states of the selected students and influence between the selected students and their neighbours. Belief for the next intervention is updated according to the newly obtained observations. The objective of UCC is to decrease the global stress level of all students to prevent depression.

4.1 Network and Dynamics

The connection network of N students is represented by a directed graph $G = \langle V, E \rangle$ with the node set $V(|V| = N)$ and the edge set E. Every $i \in V$ represents a student in the connection network. Moreover, every $e = \{(i,j) | i,j \in V\} \in E$ represents that student i is a friend of student j and is associated with real value w_{ij}, which terms the *influence* that i induces to j. Since the friendship between a pair of students is mutual (Seshadhri, Kolda, and Pinar 2012), if $(i,j) \in E$ then $(j,i) \in E$. However, different w_{ij} and w_{ji} values indicate the different influence i and j have on each other and we set $w_{ii} = 0$. For the sake of description, let $\mathcal{N}^{in}(i)$ be student i's in-neighbours where $(j,i) \in E$ with $0 < w_{ji} \leq 1$ for $j \in \mathcal{N}^{in}(i)$ and $\mathcal{N}^{out}(i)$ as out-neighbours where $(i,j) \in E, 0 < w_{ij} \leq 1$ for $j \in \mathcal{N}^{out}(i)$, respectively. The mental state of a student is one of the values in the discrete set $\mathcal{M} = \{0, 1, 2, \cdots, \mu\}$ [1] in which 0 represents the least stressed mental state and μ represents the most stressed mental state. Therefore, the students' mental states are represented by $\mathbf{v} = \langle v_1, ..., v_N \rangle$ where $v_i \in \mathcal{M}, i \in V$ is the mental state of student i.

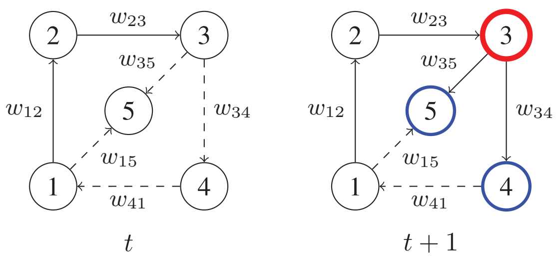

Figure 1: An illustrative example with 5 students. At round t, UCC knows the influence between the students w_{12} and w_{23}, represented by the solid lines. Influence unknown by UCC are represented by the dashed lines. Student 3 is picked by UCC to be intervened at t. w_{34} and w_{35} are known by UCC and students 4 and 5 are also influenced according to Eq. (1).

We assume that in each intervention, UCC reduces the selected student's stress level by δ [2]. Due to her mental state change, her emotion propagates to her friends $j \in \mathcal{N}^{out}(i)$ by the propagation where the extent of influence varies by *influence* w_{ij}. The process is illustrated in Figure 1.

The extent of influence on j by i is represented by $\Delta_{i \to j}$ which is defined as:

$$\Delta_{i \to j} = \lfloor \frac{w_{ij}(\mu - v_i)}{w_{ij}(\mu - v_i) + \sum_{k \in \mathcal{N}^{in}(j) \setminus \{i\}} v_k \cdot w_{kj}} \cdot \delta \rfloor \quad (1)$$

which implies that when v_i is smaller, i.e., the influencer i is less stressed, $\Delta_{i \to j}$ is larger. When $\sum_{k \in \mathcal{N}^{in}(j) \setminus \{i\}} v_k \cdot w_{kj}$ is larger, i.e., her other neighbours have more stressed mental states, $\Delta_{i \to j}$ is smaller. This is inspired by the studies that the happiness/stress of the neighbours also affect the extent of influence (Rice et al. 2012; Rafferty and Griffin 2006).

Hence, the total mental state value decrease on j is:

$$\Delta j = a_j \cdot \delta + \sum\nolimits_{a_i = 1, i \in V \setminus \{j\}} \Delta_{i \to j} \quad (2)$$

where $\mathbf{a} = \langle a_i \rangle, \forall i \in V$ such that $a_i = 1$ if student i is selected and $a_i = 0$ otherwise. The first term $a_j \cdot \delta$ is the influence induced by UCC and the second term is the influence induced by the propagation from the intervened neighbours of j such that $\Delta_{i \to j}$ are aggregated for all neighbours i of j who are intervened.

4.2 Uncertainties

UCC does not have complete information of the students' mental states and influence initially. Hence, we model the uncertainty of students' mental states at the t^{th} intervention by defining an $N \times (\mu + 1)$ matrix $\hat{P}_{t-1}$ and each row $\hat{\mathbf{p}}_i^{t-1} = \langle \hat{p}_i^{t-1}(m) \rangle$ is the probability distribution over the discrete set $\mathcal{M}$ of student i. $\hat{p}_i^{t-1}(m)$ expresses the probability of student i being evaluated as mental state value m at t. For the uncertain influence, we also define an $N \times N$ matrix $\hat{W}_0$ which represents the estimates of influence between each pair of students. The values in $\hat{W}_0$ are estimated by the counsellors based on the information collection before the intervention process.

Initially, UCC has mental state estimates $\hat{P}_0$ and influence estimates $\hat{W}_0$. In each intervention, the mental states of the selected students and influence are observed. Hence, in t^{th} intervention, UCC derives $\hat{P}_t$ from the belief which is updated during the intervention. The rule for the belief update is described in POMDP formulation section. $\hat{W}_t$ is also updated by assigning $\hat{w}_{ij} = w_{ij}, \forall j \in \mathcal{N}^{out}(i)$ and $\hat{w}_{ji} = w_{ji}, \forall j \in \mathcal{N}^{in}(i)$ for each intervened student i.

4.3 POMDP Formulation

POMDPs are sequential decision making models under uncertainty (Puterman 2014). Formally, a POMDP is defined as $\mathcal{P} = \langle S, A, O, T, \Omega, R, b^0 \rangle$.

States and Initial Belief. S is the state set. A state is defined as $\mathbf{s} = \langle \mathbf{v}, \hat{W} \rangle$ where $\mathbf{v}$ denotes the students' mental states and $\hat{W}$ is defined as $\hat{w}_{ij} = w_{ij}$ if the influence of student i on j is known by UCC and $\hat{w}_{ij} = \hat{w}^0_{ij}$ otherwise where $\hat{w}^0_{ij}$ is the initial estimation of w_{ij} by UCC. UCC has an initial belief over states b^0 which is a distribution over S and $b^0_{\mathbf{s}}$ is the probability that the POMDP is at $\mathbf{s}$ at the beginning of the interventions.

Actions and Observations. UCC's selection of K students at each intervention is defined as action $\mathbf{a}$ where $a_i = 1$ means student i is selected and $a_i = 0$ otherwise, given the constraint $\sum_{i \in V} a_i = K$. All actions belong to the set A. UCC's observation by taking the action $\mathbf{a} \in A$ at state $\mathbf{s}$ is defined as $o(\mathbf{s}, \mathbf{a}) = \{v_i, w_{ij}, w_{ji} | \forall a_i = 1, j \in V, \mathbf{v} \in \mathbf{s}\}$, i.e., the mental states and the associated influence of the intervened students. All observations belong to the set O. Ω is the observation function of the POMDP which is uniquely defined by the action $\mathbf{a}$ and the state $\mathbf{s}$:

$$\Omega(o, \mathbf{s}, \mathbf{a}) = \begin{cases} 1, & \text{if } o = o(\mathbf{s}, \mathbf{a}); \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

Transition Probabilities Heuristic. $T(\mathbf{s}, \mathbf{a}, \mathbf{s}')$ is the transition probability of reaching $\mathbf{s}'$ from $\mathbf{s}$ by taking action $\mathbf{a}$. UCC takes action $\mathbf{a}$ and the change in students' mental states is calculated as by Eq (2). Therefore, we can define $T(\mathbf{s}, \mathbf{a}, \mathbf{s}')$ as:

$$T(\mathbf{s}, \mathbf{a}, \mathbf{s}') = \begin{cases} 1, & \text{if } \mathbf{s}' = \langle \mathbf{v}', \hat{W}' \rangle; \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

where $\mathbf{v}'$ and $\hat{W}'$ are students' mental states and influence of the new state $\mathbf{s}'$ that are updated as:

$$v'_j = v_j - \Delta j; \hat{w}'_{ij} = \begin{cases} w_{ij}, & \text{if } a_i = 1 \text{ or } a_j = 1; \\ \hat{w}_{ij}, & \text{otherwise.} \end{cases} \quad (5)$$

where $v'_j \in \mathbf{v}', \hat{w}'_{ij} \in \hat{W}', \forall i, j \in V$.

[1] In current literature, mental states can only be roughly evaluated inexplicitly as *mild*, *moderate* and *severe* (WHO 1993).

[2] $\delta = f(v_i)$. If $v_i < \delta$, we assign $v_i = 0$ after decrease. This also applies to Δj in Eq (2).

Reward and Policy. The reward $R(\mathbf{s}, \mathbf{a})$ of taking action $\mathbf{a} \in A$ in state $\mathbf{s} = \langle \mathbf{v}, \hat{W} \rangle$ is defined by:

$$R(\mathbf{s}, \mathbf{a}) = \sum_{\mathbf{s}' \in S} T(\mathbf{s}, \mathbf{a}, \mathbf{s}') \left(\sum_{i \in V} (v_i - v_i') \right) \tag{6}$$

We define the history at intervention t as a sequence of past actions and observations $H_t = \langle a_1, o_1, a_2, \ldots, a_t, o_t \rangle$. We denote $\mathcal{H}_t$ as the set of all possible histories at t. The policy is defined as $\pi : \mathcal{H}_t \rightarrow A$ which takes history H_t as input and outputs the action $\mathbf{a}$. The expected reward for π starting from b^0 is defined as $V^\pi(b^0) = \sum_{t=1}^{Q} \mathbb{E}[R(\mathbf{s}, \mathbf{a}) | b^0, \pi]$. $\mathbb{E}[\cdot]$ outputs the expected value of the input. The optimal policy π^* is the policy that maximizes $V^\pi(b^0)$. Formally, $\pi^* = \arg\max_\pi V^\pi(b^0)$.

Belief Update. Since in each state $\mathbf{s}$, we have the deterministic value of $\hat{W}$ where each element is either $\hat{w}_{ij}$ or w_{ij}, the initial belief b^0 can be defined by $\hat{P}_0$ and $\hat{W}_0$ such that for $\mathbf{s} = \langle \mathbf{v}, \hat{W} \rangle$:

$$b^0_{\mathbf{s}} = \begin{cases} \prod_{v_i \in \mathbf{v}} \hat{p}^0_i(v_i), & \text{if } \hat{W} = \hat{W}_0 \\ 0, & \text{otherwise.} \end{cases} \tag{7}$$

At intervention t, each state $\mathbf{s} = (\mathbf{v}, \hat{W})$ with belief $b^{t-1}_{\mathbf{s}}$ transits to $s' = (\mathbf{v}', \hat{W}')$ upon taking action $\mathbf{a}$. UCC observes $o \in O$ with the probability of $\Omega(o, \mathbf{s}', \mathbf{a})$. Hence we update the belief by $b^t_{\mathbf{s}'} = \gamma \cdot \Omega(o, \mathbf{s}', \mathbf{a}) \cdot \sum_{\mathbf{s} \in S} T(\mathbf{s}, \mathbf{a}, \mathbf{s}') \cdot b^{t-1}_{\mathbf{s}}$ where γ is the normalizing constant: $\gamma = 1/(\sum_{\mathbf{s}' \in S} \Omega(o, \mathbf{s}', \mathbf{a}) \cdot \sum_{\mathbf{s} \in S} T(\mathbf{s}, \mathbf{a}, \mathbf{s}') \cdot b^{t-1}_{\mathbf{s}})$. After that, we update $\hat{P}_t$ based on the belief update with $\hat{p}^t_j(m) = \sum_{\mathbf{s}' \in S, v_j' = m} b^t_{\mathbf{s}'}$.

5 MLPRAP

To solve the formulated problem, we first tried online POMDP solvers such as DESPOT (Ye et al. 2017) and POMCP (Silver and Veness 2010) which can scale up to large networks. These solvers, however, limit the size of initial belief set. This makes them not suitable for our setting since the initial belief set may be as large as the state set.

We propose MLPRAP (Multi-Level Partition algorithm with Reasoning and Abstracted Planning), extended from the algorithms for the dynamic influence maximization in social networks to improve scalability and solution quality. MLPRAP sequentially selects the students in an uncertain large-scale network with three novelties: reasoning on the estimated mental states of students to refine the belief before each intervention, abstraction of the POMDP states to solve large POMDPs and multi-level partitioning of the graph.

Algorithm 1 describes the flow of MLPRAP. We partition the graph G with multi-level partitioning by Algorithm 4 or 5 to obtain a map $\mathbb{P}$ of partition par and k (line 1). Then, we generate abstracted POMDPs for each par (line 2). Next, we refine belief b^{t-1} with reasoning by Algorithm 2 (line 4). Each POMDP is solved to find the optimal one-node $action$ with QMDP Heuristics (lines 6-8). We choose K actions from $\mathbb{A}$ according to multi-level partitioning variants to get $\mathbf{a}$ which is added to policy π and belief b^t is updated (lines 9-10).

Algorithm 1: `MLPRAP`(G)

1 Obtain $\mathbb{P}$ with multi-level partitioning (Algorithm 4 or 5) ; // $\mathbb{P}$ is a map of par and k
2 Generate abstracted POMDPs $\mathcal{P}'_k$ and assign k for $\langle par, k \rangle \in \mathbb{P}$;
3 **for** $t = 1 : Q$ **do**
4   Reason to refine b^{t-1} of each $\mathcal{P}'_k$ (Algorithm 2);
5   Initialize $\mathbb{A} = \emptyset$; // $\mathbb{A}$ is a map of $action$ and k
6   **for** $\mathcal{P}'_k$ **do**
7         $action =$ `FindBestAction`$(\mathcal{P}'_k)$;
8   add $\langle action, k \rangle$ to $\mathbb{A}$;
9   Choose K actions from $\mathbb{A}$ and assign to $\mathbf{a}$;
10   Add $\mathbf{a}$ to π, and update b^t of $\mathcal{P}'_k$ according to $(b^{t-1}, \mathbf{a})$;
11 **return** π

5.1 Reasoning

In our POMDP, the initial students' mental states and *influence* ($\hat{P}_0$ and $\hat{W}_0$) are inaccurate since they are estimated by UCC without evaluation. This leads to poor solution quality as the algorithm sequentially selects the participants from inaccurate initial beliefs. At every intervention t, as we get information about the selected students' mental states and the real influence between the students, we do reasoning on the estimated students' mental states and refine b^t so that the belief estimates are closer to the true mental states.

The change in her emotions propagates to her neighbours in the social network as emotions can spread from one person to another. Moreover, her mental state is also affected by the mental states of her neighbours. There are two main conclusions that describe the relationship about the mental states in a social network: (i) the close friends in the network have similar mental state values (Hill, Griffiths, and House 2015); (ii) the mental state of a student is not just affected by one of her friends independently but affected by the mental states of all her friends (Eyre et al. 2017). These existing findings can be represented by the State Relationship Rule.

State Relationship Rule (SRR). *A student's mental state lies within the range defined by the weighted average of all her neighbours' mental state ranges. Given each neighbour i's mental state range as $[lb_i, ub_i]$ where $lb_i \leq v_i \leq ub_i$, j's range $[lb_j, ub_j]$ is defined as:*

$$lb_j = \lfloor \frac{\sum_{i \in \mathcal{N}^{in}(j)} lb_i \cdot \hat{w}_{ij}}{\sum_{i \in \mathcal{N}^{in}(j)} \hat{w}_{ij}} \rfloor ; ub_j = \lceil \frac{\sum_{i \in \mathcal{N}^{in}(j)} ub_i \cdot \hat{w}_{ij}}{\sum_{i \in \mathcal{N}^{in}(j)} \hat{w}_{ij}} \rceil \tag{8}$$

Accordingly, the prediction of the mental state depends not only on the mental state of the neighbours but also on the influence of the neighbours. We also allow uncertainty of the range predicted by SRR. Therefore, there will be a probability to predict a student to be stressed although she is surrounded by all neighbours in happy mental states.

So far, we assume that all mental state values can occur in our POMDP. There might be some states in the belief with the mental state values that are inconsistent with SRR. During the reasoning process, we take SRR into account and set the probabilities of the states in the belief that violate SRR to 0. This process can reduce the uncertainty on mental state values and thus, improve the solution quality. A concrete example is shown in Figure 2 to describe how the belief is refined during the reasoning process according to SRR.

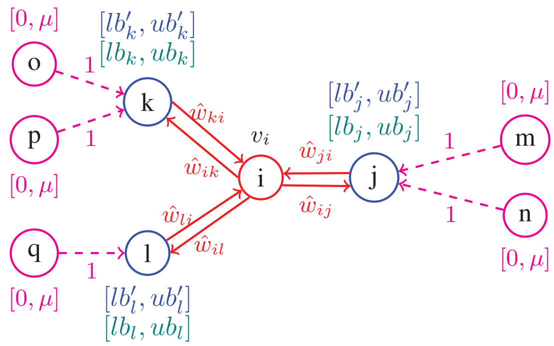

Figure 2: An illustrative example of reasoning process. When UCC intervenes student i, UCC observes the value of her mental state and influence between her neighbours shown in red. Using this observation, we predict the ranges of j, k, l in blue with observed v_i and the ranges from their other in-neighbours in magenta. The predicted ranges are in green. Using these ranges, we calculate $[lb_i, ub_i]$ for i, the j, k, l's ranges are modified if $v_i \notin [lb_i, ub_i]$. The final ranges for j, k, l are shown in blue.

Example 1. *We consider an example in Figure 2 with* $\mu = 10$ *where* $\{v_i = 5, w_{ij} = 1, w_{ji} = 1, w_{ik} = 0.5, w_{ki} = 1, w_{il} = 0.25, w_{li} = 0.25\}$ *is observed. First, with Eq. (8), we can predict the mental state ranges for* j, k, l *as* $[1, 9], [0, 10], [0, 10]$ *respectively. Using these ranges, we find that we get* $ub_i = 4$ *and* $lb_i = 6$ *after calculating* i*'s range with* $v_k = 0$ *and* $v_k = 10$*. Hence,* 0 *and* 10 *are removed from* k*'s range. Thus, we get* $[1, 9], [1, 9], [0, 10]$ *as the refined ranges for* j, k, l*. During reasoning process, we set* $b_{\mathbf{s}}$ *as zero if* $\mathbf{s}$ *contains inconsistent mental states that are not in the refined ranges and violates SRR.*

According to SRR, we do the reasoning as follows. When UCC intervenes student i, UCC gets an observation $\{v_i, w_{ij}, w_{ji} | j \in V\}$. With this observation, we predict the range for student j's mental state where $j \in \mathcal{N}^{out}(i)$. We set the mental state range for other unobserved students as $[0, \mu]$ and influence as the maximum, i.e., 1 to make the predicted range wider to compensate for uncertain values. For the robust prediction of mental states, UCC defines that the predicted range for j's mental state has the width of at least ω i.e., $ub_j - lb_j \geq \omega$. Hence, if the width of the range predicted by Eq. (8) is less than ω, we modify the range as $[\theta_j - \omega/2, \theta_j + \omega/2]$ [3] where $\theta_j = \lfloor (lb_j + ub_j)/2 \rfloor$. On the other hand, we need to make sure that the observed value v_i is in the range $[lb_i, ub_i]$ which is calculated by the predicted ranges for $v_j, \forall j \in \mathcal{N}^{in}(i)$. For each value α in range $[lb_j, ub_j]$, we use $[lb_k, ub_k]$ for $k \in \mathcal{N}^{in}(i) \setminus \{j\}$ to find $[lb_i, ub_i]$ and α is removed from the range if $ub_i < v_i$ or $lb_i > v_i$. We run the refining process until all the observed v_i values are in the range $[lb_i, ub_i]$ which is calculated by all the predicted ranges for $j \in \mathcal{N}^{in}(i)$.

[3] If $\theta_j - \omega/2 < 0$, $lb_j = 0$. If $\theta_j + \omega/2 > \mu$, $ub_j = \mu$.

We refine the belief before each intervention so that the algorithm selects action based on the belief closer to the real network. We do reasoning on the network, predict the range and check for each belief state where $b^{t-1}_{\mathbf{s}} > 0$ if all mental state values of $\mathbf{s}$ are in respective predicted ranges.

Algorithm 2: `Reasoning`$(b_{t-1}, \mathbf{d})$

1 **for** $\mathbf{s} \in S, b^{t-1}_{\mathbf{s}} > 0$ **do**
2 $\quad$ $\mathbf{lb}, \mathbf{ub}$ = `PredictMentalStateValues`$(\mathbf{s}, \mathbf{d})$;
3 $\quad$ **for** $v_j \in \mathbf{s}$ **do**
4 $\quad\quad$ **if** $v_j \notin [lb_j, ub_j]$ **then** $b^{t-1}_{\mathbf{s}} = 0$;
5 $b'^{t-1}_{\mathbf{s}} = b^{t-1}_{\mathbf{s}} / \sum_{\mathbf{s}' \in S} b^{t-1}_{\mathbf{s}'}, \forall \mathbf{s} \in S$;
6 **return** b'^{t-1};

Algorithm 2 checks each state where $b^{t-1}_{\mathbf{s}} > 0$ if the students' mental states are in the range predicted by SRR (line 1). We define a vector $\mathbf{d}$ where $d_i = 1$ if student i is observed and $d_i = 0$ otherwise. The range for each student's mental state value is predicted according to SRR in `PredictMentalStateValues(.)` which is described in Algorithm 3. After that, the algorithm checks if the mental state of all students are in the respective predicted ranges and sets $b^{t-1}_{\mathbf{s}}$ to 0 otherwise (line 4). Finally, non-zero belief values are rescaled to sum up to 1 (line 5).

Predicting Mental State Values by SRR Algorithm 3 returns the predicted ranges for the students. lb_i, ub_i are set as certain mental state of i and $\tilde{w}_{ij}$ is set as certain influence of i and j if i is observed, i.e., $d_i = 1$, and set to $lb_i = 0, ub_i = \mu, \tilde{w}_{ij} = 1$ otherwise (lines 1-7). Given the observed student i, the mental state range for her out-neighbour j is predicted by applying SRR (lines 8-15). Lines 16-22 describe that the mental state range $[lb_i, ub_i]$ of observed student i is calculated using the predicted ranges for her in-neighbours. If $\alpha \in [lb_j, ub_j], \forall j \in \mathcal{N}^{in}(i)$ predicts the range $[lb_i, ub_i]$ and $v_i \notin [lb_i, ub_i]$, it is removed from $[lb_j, ub_j]$ (line 22).

5.2 Abstraction of POMDP States

We reduce the huge state space by defining σ such that a student's mental state is one of the values in the abstracted discrete set $\mathcal{M}' = \{0, \sigma, 2\sigma, ..., \mu\}$ and we define the state set S' with $\mathcal{M}'$. Hence, the non-abstracted total number of states, i.e., $(\mu + 1)^N$ is reduced to $\lceil (\mu + 1)/\sigma \rceil^N$.

Every state in S' belongs to S, i.e., $S' \subset S$. Given action $\mathbf{a} \in A$ and state $\mathbf{s} \in S'$, the successive state $\hat{\mathbf{s}} \in S$ is obtained according to the transition function of original POMDP $\mathcal{P}$. The change in mental states is computed with observed v_i for $a_i = 1$ and abstracted v'_j for $a_j = 0$. If $\hat{\mathbf{s}} \notin S'$, we replace with $\mathbf{s}' \in S'$ such that $0 < v'_i - \hat{v}_i \leq \sigma, \forall i \in V$. The policy evaluated with S' is denoted as π'. We denote the optimal policy of the abstracted POMDP as π'^* and the expected reward as $V^{\pi'^*}$.

Algorithm 3: PredictMentalStateValues$(\mathbf{s}, \mathbf{d})$

for $i \in V$ **do**
  if $d_i == 1$ **then**
    $lb_i = ub_i = v_i$;
    for $j \in V$ **do** $\tilde{w}_{ij} = w_{ij}$;
  else
    $lb_i = 0; ub_i = \mu$;
    for $j \in V$ **do** $\tilde{w}_{ij} = 1$;
for $i \in V, d_i == 1$ **do**
  for $j \in \mathcal{N}^{out}(i), d_j == 0$ **do**
    $lb_j = \lfloor \frac{\sum_{k \in \mathcal{N}^{in}(j)} lb_k \cdot \tilde{w}_{kj}}{\sum_{k \in \mathcal{N}^{in}(j)} \tilde{w}_{kj}} \rfloor$;
    $ub_j = \lceil \frac{\sum_{k \in \mathcal{N}^{in}(j)} ub_k \cdot \tilde{w}_{kj}}{\sum_{k \in \mathcal{N}^{in}(j)} \tilde{w}_{kj}} \rceil$;
    $\theta_j = (lb_j + ub_j)/2$;
    $lb'_j = \theta_j - \omega/2, ub'_j = \theta_j + \omega/2$;
    if $lb'_j < lb_j$ **then** $lb_j = lb'_j$;
    if $ub'_j > ub_j$ **then** $ub_j = ub'_j$;
for $i \in V, d_i == 1$ **do**
  for $j \in \mathcal{N}^{in}(i), d_j == 0$ **do**
    for $\alpha \in [lb_j, ub_j]$ **do**
      $lb_i = \lfloor \frac{\alpha \cdot \tilde{w}_{ji} + \sum_{k \in \mathcal{N}^{in}(i) \setminus \{j\}} lb_k \cdot \tilde{w}_{ki}}{\sum_{k \in \mathcal{N}^{in}(i)} \tilde{w}_{ki}} \rfloor$;
      $ub_i = \lceil \frac{\alpha \cdot \tilde{w}_{ji} + \sum_{k \in \mathcal{N}^{in}(i) \setminus \{j\}} ub_k \cdot \tilde{w}_{ki}}{\sum_{k \in \mathcal{N}^{in}(i)} \tilde{w}_{ki}} \rceil$;
      if $lb_i > v_i$ *or* $ub_i < v_i$ **then**
        remove α from $[lb_j, ub_j]$;
return $\mathbf{lb}, \mathbf{ub}$;

Although the abstraction method improves scalability, there is some loss of solution quality due to the approximation, i.e., $V^{\pi^*} - V^{\pi'^*}$. Hence, in Lemma 1, we prove a theoretical bounded error, $V^{\pi^*} - V^{\pi'^*}$, for independent network at the end of round Q.

Lemma 1. *For a certain network with independent students,* $V^{\pi^*} - V^{\pi'^*} \leq Q \cdot (\sigma - 1) \cdot K$.

Proof. The maximum difference between v_i and v'_i is $(\sigma - 1)$ since we have $\mathbf{s}'$ if $0 < v'_i - \hat{v}_i \leq \sigma$. Hence, for Q rounds where K nodes are chosen at each round, the maximum difference between the total expected rewards of the optimal policies π^* and π'^* is $Q \cdot (\sigma - 1) \cdot K$.

In Lemma 2, we prove the bounds for certain networks with connections between the students. For general networks, we cannot prove the bounds since the influence propagation greatly depends on the neighbours' mental states and influence. Hence, we consider a complete graph with $w_{ij} = w, \forall i, j \in V$.

Lemma 2. *For a certain network where the students' mental state values and influence are known,*

$$V^{\pi^*} - V^{\pi'} \leq Q \cdot K \cdot \big(\frac{\delta \cdot (N \cdot \mu - 2)}{\mu + N - 2} - \delta + \sigma - 1 \big) \quad (9)$$

Proof. Let π^* be the optimal policy of the original POMDP $\mathcal{P}$ with initial state $\mathbf{s}^0$. At round t, given the state $\mathbf{s}^{t-1}$ and action $\mathbf{a}^t$, $\mathbf{s}^{t-1}$ transits to $\mathbf{s}^t$. We refer to the mental states at t as $v_i^t, \forall i \in V$ where $v_i^t \in \mathbf{v}^t, \mathbf{v}^t \in \mathbf{s}^t$. In a complete graph, the influence level on j at round t is defined using Eq. (2) as:

$$\Delta_j^t = a_j^t \cdot \delta + \sum_{a_i^t = 1, \forall i \in V \setminus \{j\}} \frac{(\mu - v_i^{t-1}) \cdot \delta}{(\mu - v_i^{t-1}) + \sum_{k \in N \setminus \{i\}} v_k^{t-1}} \quad (10)$$

Hence, we calculate R_t as:

$$R_t = K \cdot \delta + (N - 1) \cdot \sum_{a_i^t = 1, \forall i \in V} \frac{(\mu - v_i^{t-1}) \cdot \delta}{(\mu - v_i^{t-1}) + \sum_{k \in N \setminus \{i\}} v_k^{t-1}} \quad (11)$$

$v_i^{t-1} = 1, \forall i \in V$ gives the largest propagation to find the maximum total estimated reward possible for $\mathcal{P}$. Hence,

$$R_t \leq \frac{K \cdot \delta \cdot (N \cdot \mu - 2)}{\mu + N - 2} \quad (12)$$

To find the lower bound, we ignore the propagation process of the abstracted POMDP $\mathcal{P}'$. Hence,

$$R'_t \geq K \cdot (\delta - \sigma + 1) \quad (13)$$

Hence, we have the bound for $V^{\pi^*} - V^{\pi'}$ as:

$$V^{\pi^*} - V^{\pi'} \leq \hat{R} = Q \cdot K \cdot \big(\frac{\delta \cdot (N \cdot \mu - 2)}{\mu + N - 2} - \delta + \sigma - 1 \big) \quad (14)$$

Since $V^{\pi'^*} \geq V^{\pi'}$, we proved that $V^{\pi^*} - V^{\pi'^*} \leq \hat{R}$.

5.3 Multi-Level Partitions

We improve the scalability further by multi-level partitioning of the graph. The most intuitive way is to partition the graph into K partitions, i.e., the number of selected students. But the POMDPs are still very large to be solved. Hence, we divide each of K partitions into smaller subpartitions.

The first variant is MLP-B which has ηK *balanced* partitions, i.e., each partition contains a similar number of nodes. Algorithm 4 starts with initializing $\mathbb{P}$, a map of partitions and their indices, as an empty set (line 1). We use the METIS algorithm (Karypis and Kumar 1998) to partition the graph G while minimizing cross-edge influence between partitions so as not to lose the network structure of the graph (line 2). It returns a list of partitions, i.e., $pars$. Each partition is indexed (line 4) and ensembled as POMDP to compute the optimal one-node action. Finally, actions with the best K rewards are chosen to get an action of K nodes.

Algorithm 4: $\mathtt{MLP-B}(\eta)$

1 Initialize $\mathbb{P} = \emptyset$ and $k = 0$; // $\mathbb{P}$ is a map of par and k
2 $pars$ = $\mathtt{METIS}(G, \eta K)$;
3 **for** $par \in pars$ **do**
4   add $\langle par, k \rangle$ to $\mathbb{P}$; $k++$;
5 **return** $\mathbb{P}$;

As balanced partitioning may not maintain the network structure, in the second variant, MLP-C, we cluster the graph into K partitions by finding minimum cross-edge cuts without keeping the partitions balanced. Let l be the maximum

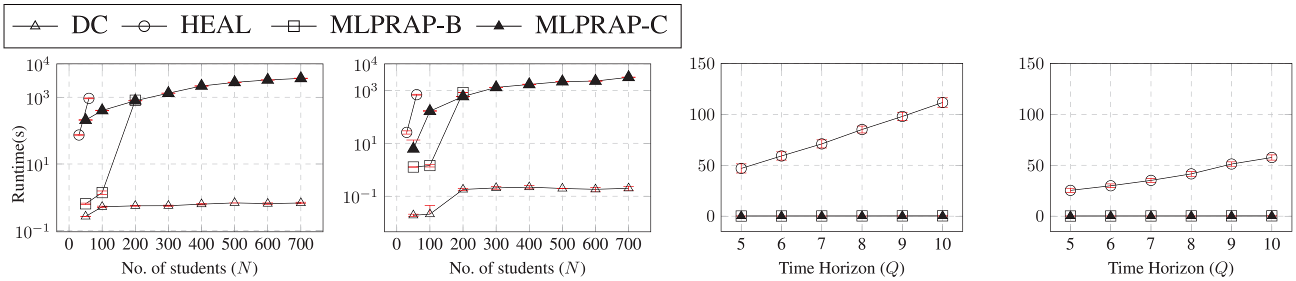

(a) BA (b) ER (c) BA (d) ER

Figure 3: Scalability comparison of MLPRAP variants with DC, HEAL (Figure 3a and 3b are plotted in log-scale)

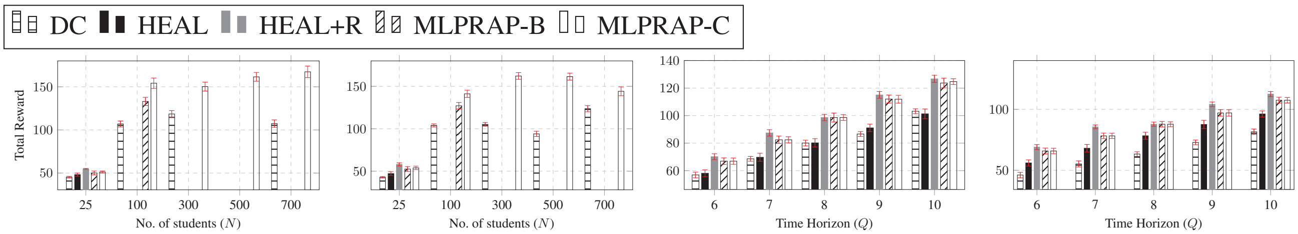

(a) BA (b) ER (c) BA (d) ER

Figure 4: Solution quality comparison of MLPRAP variants with DC, HEAL

limit of a partition that can be solved. If a partition is larger than l, we divide the partition into smaller subpartitions.

In Algorithm 5, we first cluster the graph G into K partitions by finding minimum cut and store the partitions in $pars$ (line 1). We initialize the map $\mathbb{P}$ as an empty set (line 2). The partitions are indexed and added to $\mathbb{P}$ (line 8). If a partition has more than l nodes, we divide it into $\lceil \texttt{nodecount}/l \rceil$ subpartitions using METIS and store in $subpars$. We use METIS over clustering to avoid having many levels of partitioning to get subpartitions. The subpartitions are kept with same index and added to $\mathbb{P}$ to indicate that only one node from the partitions with same index will be chosen for an action of K nodes (lines 4-7). After obtaining the optimal action from each POMDP, the action with maximum reward from the partition with same index is chosen.

Algorithm 5: $\texttt{MLP} - \texttt{C}(l)$

1 $pars = \texttt{minimumCutClustering}(G, K)$;
2 Initialize $\mathbb{P} = \emptyset$ and $k = 0$;
3 **for** $par \in pars$ **do**
4   **if** $\texttt{nodecount}(par) > l$ **then**
5         $subpars = \texttt{METIS}(par, \lceil \texttt{nodecount}(par)/l \rceil)$;
6   **for** $subpar \in subpars$ **do**
7   add $\langle subpar, k \rangle$ to $\mathbb{P}$
8   **else** add $\langle par, k \rangle$ to $\mathbb{P}$;
9     $k++$;
10 **return** $\mathbb{P}$;

6 Experimental Evaluation

In this section, we analyze the performance of MLPRAP-B and MLPRAP-C, i.e., MLPRAP with MLP-B and MLP-C with different settings. All our experiments are run on a 3.2 GHz 4-core Intel machine having 16 GB of RAM. The results are averaged over 30 trials. We use runtime and total reward as metrics to evaluate scalability and solution quality on both synthetic and real networks. The 95% confidence intervals are drawn in all figures which show that all the results are statistically significant. During experiments, we first compute UCC's policy and simulate on the networks without any uncertainties to compute the real reward obtained by the policy during the intervention process.

6.1 Experiment Setup

Problem Instance Generation. We run the lab experiments to evaluate the performance of the algorithm. We synthesize the problem instances since real-world experiments on the study of intervention process in a social network is challenging and there is no publicly available data which studies the stress level of the people in a network. However, evaluations of algorithms on synthetic data are widely accepted (Yadav et al. 2015; 2016; Wilder et al. 2017) as they serve as an important first step towards future applications of the model. We generate networks with realistic mental state values while using reasonable measures according to (Eyre et al. 2017; Hill, Griffiths, and House 2015) so that the network reflects the mental state values of each individual in the network. There are two kinds of networks considered:

- **synthetic networks**: We generate the synthetic network G by two methods. First, Barabási-Albert scale-free networks, where each new vertex is connected to ξ vertices using a preferential attachment mechanism (Barabási and Albert 1999). Since the students stay in groups of 3 or 4 which is the size of a team for a group project, we let $\xi = 3$. Second, Erdos-Ŕenyi random networks $ER(N, |E|)$, where exactly $|E|$ edges are randomly constructed between each pair of nodes (Erdos 1959). We set $|E| = 3N$ to let each node have 3 connections on average. Then, we assign $\hat{w}_{ij} \in W_0$ as randomized values from $[0, 1]$.

- **real networks**: The first network is Zachary Karate Club dataset (Karate) with 34 nodes and 78 edges (Zachary 1977) which is the friendship data of the members of a university karate club. This will closely reflect the relationship between students in the network and the effectiveness of interventions. We assign $\hat{w}_{ij} \in W_0$ as randomized values from $[0, 1]$. The second dataset is Mobile-1 dataset (Mobile) which has 107 nodes and 513 edges (Tang, Lou, and Kleinberg 2012). It consists of the logs of calls and cell tower IDsx of users for ten months. We assign communication count between users i and j as $\hat{w}_{ij}$.

After we have obtained the network, we set $\mu = 9$, pick a student $i \in V$ and assign the uniformly sampled value from $[0, \mu]$. For all other $j \in V \setminus \{i\}$, we iteratively predict the mental state ranges according to SRR and assign the sampled value from the predicted range. We repeat the process until convergence where all the nodes are assigned with the mental state values.

Baselines. We use two algorithms as baselines: (1) Degree Centrality (DC) which selects the highest degree node first; (2) K-variant HEAL as it is the most relevant algorithm which has been demonstrated to perform much better than earlier algorithms.

6.2 Experiment results

Scalability Analysis. We compare the runtimes of our algorithms and baselines by varying network sizes. We set $Q \in \{5, ..., 10\}$, $\delta = 2$, $K = 5$, $\sigma = 3$, $\eta = 2$ and $l = 10$. In Figure 3a, we compare the runtime of each algorithm along the y-axis w.r.t varying network sizes along the x-axis. For example, for a problem with a network of 30 students, it takes 0.070s for DC, 375.383 s for HEAL, 0.149 s for MLPRAP-B and 0.748 s for MLPRAP-C. While DC is the fastest, it does not result in good solution quality as we will discuss in solution quality analysis. HEAL can only solve up to 30 students. It runs out of memory for larger networks. Moreover, HEAL runs very slowly even for small networks. We also run HEAL+R (HEAL with Reasoning) and it has similar runtime with HEAL. While the running time is faster in smaller networks, MLPRAP-B runs out of memory on the network of more than 100 students with $\eta = 2$. MLPRAP-C can solve larger networks than all other algorithms. In this experiment, we show up to network with 700 students which the system can solve within the time limit of 3600s. The system can solve 1000 students network in 6141.920s. We find the same trends for ER networks shown in Figure 3b.

The second experiment runs the algorithms on BA and ER networks of 25 students ($N = 25$) with fixed $\eta = 2$ and $l = 10$. In Figures 3c and 3d, we vary time horizons along the x-axis on the same network to record the runtime along the y-axis. For a BA network of 25 students, HEAL takes 50s to plan for 5 interventions whereas DC and both variants of MLPRAP take just a fraction of a second, i.e., 0.070s, 0.195s and 0.202s respectively. While all algorithms have linearly increasing runtimes with increasing time horizon,

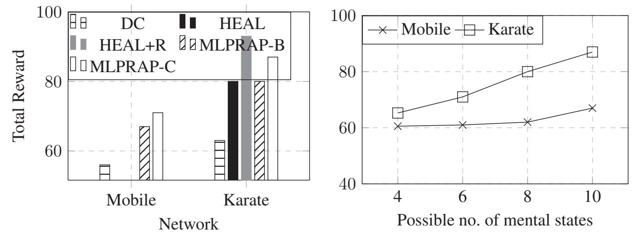

(a) Real world networks (b) Effect of estimates

Figure 5: Experiments on real networks

MLPRAP algorithms do not significantly increase and have better scalability.

Although we only scale up to a network with 700 nodes for the scalability analysis since we limit the solving time to 1hr, MLPRAP-C algorithm can solve much larger networks since the multi-level partitioning algorithm keeps partitioning the network until the partitions can be solved by the MLPRAP algorithm. With a more powerful machine and longer time, MLPRAP-C can solve much larger networks.

Solution Quality Analysis. Figures 4a and 4b show the rewards of the different algorithms along y-axis on the varying sizes of the student networks along the x-axis. We keep N=25, 100, 300, 500, 700 to highlight the limitations of baseline algorithms. For example, for BA network of 25 nodes, the total reward is 45.2, 48.17, 54.33, 50.10, 51.17 for DC, HEAL, HEAL+R, MLPRAP-B and MLPRAP-C respectively. The trends show that HEAL+R improves solution quality better than simply running HEAL and the approximated solution with abstraction does not suffer a significant loss. As networks become larger, MLPRAP variants have a larger advantage over DC. Therefore, MLPRAP-B and MLPRAP-C are more suitable to solve larger networks.

In Figure 4c and 4d, we compare the total reward (y-axis) with increasing time horizon (x-axis) for both types of random networks with 25 nodes. The trends show that MLPRAP variants result in better solution quality than DC and HEAL. While HEAL+R is the best for a network of 25 students, we have shown that it cannot scale up to larger networks. On the other hand, MLPRAP-B and MLPRAP-C do not lose much solution quality compared to HEAL+R.

Effect of uncertainty. We also analyze the effect of uncertainty of the student networks on the solution quality. The uncertainty for initial mental state values is reflected by the initial estimated mental state range. In Figure 5b, the value 4 on the x-axis represents that there are 4 possible mental state values with 0.25 probability for each student while the y-axis represents the total reward. We can conclude from the results that the homogeneous distribution, i.e., $|\mathcal{M}| = 10$, gives the best total reward. This is because homogeneous distribution assigns all the mental states with equal probability lowering the chances of incorrectly estimating the mental states and losing the solution quality in contrast to smaller ranges.

Real Networks. We also evaluate the scalability and solution quality of the algorithms on real networks (Mobile and Karate). Figure 5a describes the rewards of each network for 5 interventions with $K = 6$. HEAL and HEAL+R cannot solve Mobile and MLPRAP-C gives the best solution quality. Similar to synthetic networks, HEAL+R gives the best total reward for Karate. While MLPRAP-B only gives a similar solution quality as HEAL, MLPRAP-C gives much higher solution quality than DC and HEAL and comparable to HEAL+R.

7 Conclusion

We propose a novel model that considers emotion propagation from not only the influencer but also neighbours of the influencee while selecting the students for interventions in uncertain networks. We propose MLPRAP algorithm with reasoning, abstraction of POMDP states and multi-level partitioning of the graph into smaller POMDPs to sequentially plan to select the students for each intervention. Finally, we experiment with synthetic networks generated as BA and ER networks as well as real network datasets. We show that MLPRAP variants have significantly better scalability and solution quality comparable to the state-of-the-art algorithms.

Acknowledgement

This research is supported by Tier-1 RG13/19, NRF AISG-RP-2019-0013, and NSOE-TSS2019-01. We would like to thank Shuxin Li for giving her time in our early discussions.

References

Barabási, A.-L., and Albert, R. 1999. Emergence of scaling in random networks. *Science* 286(5439):509–512.

Blackmore, E. R.; Stansfeld, S. A.; Weller, I.; Munce, S.; Zagorski, B. M.; and Stewart, D. E. 2007. Major depressive episodes and work stress: Results from a national population survey. *American Journal of Public Health* 97(11):2088–2093.

Chen, W.; Lin, T.; Tan, Z.; Zhao, M.; and Zhou, X. 2016. Robust influence maximization. In *KDD*, 795–804.

Erdos, P. 1959. On random graphs. *Publicationes Mathematicae* 6:290–297.

Eyre, R. W.; House, T.; Hill, E. M.; and Griffiths, F. E. 2017. Spreading of components of mood in adolescent social networks. *Royal Society Open Science* 4(9):170336.

Fowler, J. H., and Christakis, N. A. 2008. Dynamic spread of happiness in a large social network: Longitudinal analysis over 20 years in the framingham heart study. *The BMJ* 337:a2338.

Goel, S.; Watts, D. J.; and Goldstein, D. G. 2012. The structure of online diffusion networks. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, 623–638.

Helmers, K. F.; Danoff, D.; Steinert, Y.; Leyton, M.; and Young, S. N. 1997. Stress and depressed mood in medical students, law students, and graduate students at mcgill university. *Academic Medicine* 72(8):708–714.

Hill, E.; Griffiths, F.; and House, T. 2015. Spreading of healthy mood in adolescent social networks. *Proc. R. Soc. B* 282(1813):20151180.

Karypis, G., and Kumar, V. 1998. Multi level k-way partitioning scheme for irregular graphs. *Journal of Parallel and Distributed Computing* 48(1):96–129.

Kempe, D.; Kleinberg, J.; and Tardos, É. 2003. Maximizing the spread of influence through a social network. In *KDD*, 137–146.

Khan, S., and Khan, R. A. 2017. Chronic stress leads to anxiety and depression. *ANN Psychiatry Ment Health* 5(1):1091.

Littman, M. L.; Cassandra, A. R.; and Kaelbling, L. P. 1995. Learning policies for partially observable environments: Scaling up. In *ICML*, 362–370.

Lucassen, P. J.; Heine, V. M.; Muller, M. B.; van der Beek, E. M.; Wiegant, V. M.; Ron De Kloet, E.; Joels, M.; Fuchs, E.; Swaab, D. F.; and Czeh, B. 2006. Stress, depression and hippocampal apoptosis. *CNS & Neurological Disorders-Drug Targets (Formerly Current Drug Targets-CNS & Neurological Disorders)* 5(5):531–546.

Puterman, M. L. 2014. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons.

Rafferty, A. E., and Griffin, M. A. 2006. Perceptions of organizational change: A stress and coping perspective. *Journal of Applied Psychology* 91(5):1154.

Rice, E.; Tulbert, E.; Cederbaum, J.; Barman Adhikari, A.; 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):226–236.

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. In *NIPS*, 2164–2172.

Tang, J.; Lou, T.; and Kleinberg, J. 2012. Inferring social ties across heterogenous networks. In *Proceedings of the 5th ACM International Conference on Web Search and Data Mining*, 743–752.

UCLA. 2018. The STAND program. https://depression.semel.ucla.edu/stand_home. Accessed: 2019-08-30.

WHO. 1993. *The ICD-10 Classification of Mental and Behavioural Disorders: Diagnostic Criteria for Research*, volume 2. World Health Organization.

Wilder, B.; Yadav, A.; Immorlica, N.; Rice, E.; and Tambe, M. 2017. Uncharted but not uninfluenced: Influence maximization with an uncertain network. In *AAMAS*, 1305–1313.

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 *AAAI*, 4006–4011.

Yadav, A.; Chan, H.; Xin Jiang, A.; Xu, H.; Rice, E.; and Tambe, M. 2016. Using social networks to aid homeless shelters: Dynamic influence maximization under uncertainty. In *AAMAS*, 740–748.

Yadav, A.; Noothigattu, R.; Rice, E.; Onasch-Vera, L.; Soriano Marcolino, L.; and Tambe, M. 2018. Please be an influencer?: Contingency-aware influence maximization. In *AAMAS*, 1423–1431.

Ye, N.; Somani, A.; Hsu, D.; and Lee, W. S. 2017. DESPOT: Online POMDP planning with regularization. *JAIR* 58:231–266.

Zachary, W. W. 1977. An information flow model for conflict and fission in small groups. *Journal of Anthropological Research* 33(4):452–473.