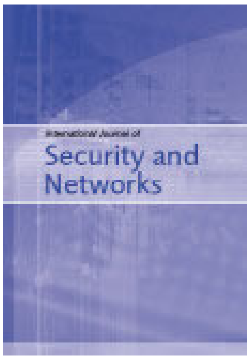

International Journal of Security and Networks

ISSN online: 1747-8413 - ISSN print: 1747-8405
https://www.inderscience.com/ijsn

MB-CIM: a multi-round budgeted competitive influence maximisation

Nadia Niknami, Jie Wu

DOI: 10.1504/IJSN.2022.10046054

Article History:

Received:	20 January 2022
Accepted:	25 January 2022
Published online:	03 April 2023

MB-CIM: a multi-round budgeted competitive influence maximisation

Nadia Niknami* and Jie Wu

Department of Computer and Information Sciences,
Temple University,
Philadelphia, PA, USA
Email: nadia.niknami@temple.edu
Email: jiewu@temple.edu
*Corresponding author

Abstract: Influence maximisation is trying to select a small set of seed users in the social network to maximise the spread of influence. An individual's decision to adopt a product or innovation will be highly dependent on the choices made by the individual's neighbours in the social network. In CIM, competitors need to decide which nodes would be an influential one and how many resources should be allocated to this member. Identifying the best algorithm for the influence maximisation under budget constraint has become a demanding task. Predicting, and controlling social influence is an exciting topic in social network analysis. We are interested in multi-round CIM where each competitor needs to decide the location and the amount of budget to invest in the most influential members simultaneously and repeatedly under a given total budget. This paper proposes a tree-approximate game-theoretical framework and introduces the new measurement as a dynamic node weight.

Keywords: budget allocation; game theory; reinforcement learning; social networks; multi-round influence maximisation.

Reference to this paper should be made as follows: Niknami, N. and Wu, J. (2023) 'MB-CIM: a multi-round budgeted competitive influence maximisation', *Int. J. Security and Networks*, Vol. 18, No. 1, pp.30–41.

Biographical notes: Nadia Niknami received her BS in Computer Science from the University of Isfahan, Iran, in 2011, and MSc from the Tarbiat Modares University, Tehran, Iran in 2015. She is currently pursuing her PhD in the Department of Computer and Information Sciences, Temple University, Philadelphia. Her current research focuses on security, privacy-preserving, attack-defence scenarios, Markov chain, social network, game theory, and software defined network.

Jie Wu is the Director of the CNC and Laura H. Carnell Professor at the Temple University. His current research interests include mobile computing and wireless networks, and network security. He regularly published in scholarly journals, conference proceedings, and books. He serves on several editorial boards, including *IEEE Transactions on Mobile Computing* and *IEEE Transactions on Service Computing*. He is/was the General Chair/Co-Chair for IEEE IPDPS'08, IEEE DOCSS'09, IEEE ICDCS'13, ACM MobiHoc'14, ICPP'16, IEEE CNS'16, and WiOps'21, as well as the Program Chair/Co-Chair for IEEE CCF CNCC'13, and ICCCN'20. He is a Fellow of the AAAS and a Fellow of the IEEE.

This paper is a revised and expanded version of a paper entitled 'A budged framework to model a multi-round competitive influence maximization problem' presented at ICC, Korea, May 2022.

1 Introduction

Viral marketing (Domingos and Richardson, 2001) is one of the most effective marketing tactics in advertising. Facebook and YouTube are two social networks that help promote products (Li et al., 2015b). Influence of social networks among individuals plays an essential role in viral marketing. The growth of online social networks has enabled them to spread quickly (Ohsaka, 2020). Influence maximisation is one of the most fundamental algorithms in social influence analysis. Over the last decade, significant effort has been put into the development of efficient algorithms for influence maximisation (Kempe et al., 2003; Nguyen and Zheng, 2013; Liontis and Pitoura, 2016; Zheng and Wu, 2017; Wang et al., 2021; Tong et al., 2018; Molinero and Riquelme, 2021). The main objectives in the IM problem are to discover which potential members, seed set, to select and how many resources to allocate to these potential members to maximise the competitors' influence. In the IM problem, all the nodes in seed set ξ are activated directly, and the remaining nodes are inactive. According to some probability distribution, when a node is activated at timestamp i, it may activate its out-neighbours at timestamp $i+1$. When no node can activate any other node, the influence propagation ends.

In the real world, there are many competitors at the same time implementing their strategies to find a considerable influence on the same social network. Actually, each rational player tries to spread its influence as maximal as possible and make its opponents as minimal as possible. That is why competitive influence maximisation (CIM) (Bharathi et al., 2007; Wu et al., 2015; Masucci and Silva, 2017; Huang et al., 2021; Shirazipourazad et al., 2012) has received a lot of attention recently. A CIM problem involves selecting the most effective seeds based on decisions made by other competitors in order to maximise their influence. The CIM model allows the influences of each player to cascade simultaneously throughout the social network, which can interfere with each other.

Considering a competitive game with two competitors, red and blue, in the given social network $\mathcal{G}(V, E, P)$, where V is the vertex set, and E is the edge set. P is a set of edge propagation probabilities, where $p(u, v)$ represents the influence probability of the edge between u and v, where $\sum_u p(u, v) < 1$. When there is no edge between u and v, $p(u, v) = 0$. The given social network is defined as a network of connections and interactions among entities. Nodes can take on one of the following states: activated by red, activated by blue, and inactive. First, competitors identify the nodes with the most influence. They compete for only these influential nodes based on the amount of budget each of them allocates to each node. After activation of a node, its influence propagates with a certain probability to their not yet activated neighbours. At each step t, each node u activated at step $t-1$ activates its neighbour v with probability $p(u, v)$. Once activated, they stay activated. Influence maximisation under both independent cascade (IC) (Kempe et al., 2003) and linear threshold (LT) (Kempe et al., 2005; Granovetter, 1978) models are NP-hard. These propagation models satisfy two important properties, submodularity, and monotonicity, in terms of their influence spread function. We will use IC in this paper. The key characteristic of this model is that diffusion events associated with every edge in the given social graph are mutually independent, and the success of the seed node u to influence one of its inactive neighbours v only depends on the propagation probability of the edge from u to v.

Consider the social network in Figure 1(a). Players red and blue compete over the nodes of this network. These players select v_1, v_3 , and v_8 as the most influential members in this network. Blue player allocates ($2, $2, $2) and red player allocates ($1, $3, $2) on members v_1, v_3, and v_8, respectively. The winning probability is proportional to the budget allocation of two parties. Red player wins v_1 with the probability of $2/(1+2) = 2/3$. Blue player wins v_3 with the probability of $3/(3+2) = 3/5$. Players have the same budget allocation on v_8. Figure 1(b) presents the result of this competition until this step. If player red adds more money, say $1 extra on his investment over v_8, his chance to win this node will be $3/2+3 = 3/5$. By doing so, he wins v_8. After finding these seed nodes, the propagation process which is based on the influence probability of relations or links between seed nodes and their friends in the given network will start. The player finding the maximum number of influenced members would be the winner of this game.

Figure 1 Budget allocation in case of tie in the CIM, (a) $S = \{v_1, v_2, v_3\}$ (b) no winner in v_3 (c) red wins v_3 (see online version for colours)

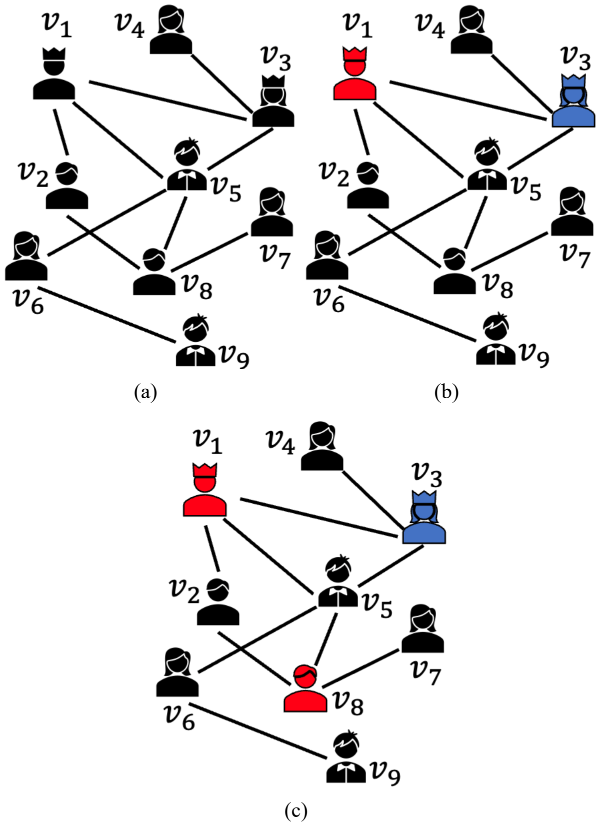

Such a scenario can be modelled by the multi-round CIM. The goal of each player in the competitive environment is to find an optimal combination of strategies to utilise their budget efficiently. The idea is to take a more realistic and practical setup, rather than selecting seeds only in the first round. In a multi-round CIM, players keep selecting seed nodes according to the current network state and the expected reactions of other players within given rounds. In addition, each player can spend a limited amount of budget in all rounds on seed nodes. Nodes with the greatest influence in a given network are selected according to different strategies. In each round, players choose a seed node ξ_t, decide the amount of budget that should be allocated to this seed node, then wait until the end of the propagation process. This assumption can be extended to multiple seed nodes in each round. Note that during each round, players take action simultaneously, but there are sequence rounds (Figure 2). As influence maximisation is NP-hard, we introduce a new notation of most reliable influence path (MRIP) as an approximation.

Table 1 Main notations

Symbol	*Meaning*
B_1/B_2	Total budget of player 1/2
$B_1(u)/B_2(u)$	Allocated budget of player 1/2 on node u
T	Total number of rounds
$N(u)$	neighbour set of node u
V	Set of nodes in the network
V^1/V^2	Set of activated nodes by player 1/2
$w(u)$	Weight of node u
$w'(u)$	Estimated total influence weight of node u
$p(u,v)$	Influence probability of edge between u and v
$R(u,v)$	Influence value of the MRIP between u and v
ξ	Seed set
s	State of network in reinforcement learning
$\pi(s)$	Policy in state s
r_t	Reward in round t in reinforcement learning
a_1/a_2	Player 1/2's action
$\mathcal{V}(s)$	Value of state s in reinforcement learning

The value of influencers varies, and competitors want to find the best value for their overall social advertising budget. It is obvious that an equal budget at each round does not sufficiently model the willingness to choose a cost-efficient seed set. Indeed, we see that the choice to use a fraction of the budget for round t is crucial: a too large budget allocation translates into a waste of budget, and a too small budget allocation translates into a waste of time (a whole round is used to influence only a few users). To circumvent this issue, instead of a budget per round, in our framework, we allow the agent to have the competition of the most influential nodes at each round under an overall budget constraint. In this paper, compared to the conference version (Niknami and Wu, 2022), we make a set of extensions in the case of explaining the approach and evaluation. We evaluate our proposed approach under different parameters such as different amount of total budget, various network structures, different densities, and different competition strategies. The contributions of this paper are summarised as follows:

- We define a new measurement called dynamic weight for nodes. Considering both fixed and dynamic weights in selecting seed nodes helps players have a more accurate selection.
- We discuss the influence spread in the social network by considering the MRIPs for each node in the process of seed selection as an approximation. MRIP is inspired by the notion of a critical path in the scheduling community.
- We consider three new features maximum weight of inactive nodes, the ratio of budget, and the weight of nodes, in case of reachability to describe the state of the network in reinforcement learning (RL).

Figure 2 A multi-stage game with two competitive players (see online version for colours)

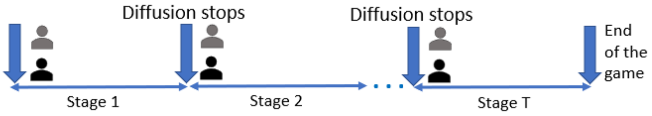

- We propose a CIM model which selects the winner of the node in case of breaking tie based on the budget proportion, rather than randomly. Players can compete on the given node by increasing their investment in this node to increase their chance.
- We evaluate the effect of our model experimentally using real datasets and some synthetic ones.

1.1 *Organisation*

The remainder of the paper is organised as follows. Section 2 briefly surveys the related works. In Section 3, we describe some preliminaries. Section 4 presents details of seed selection, budget allocation, and our proposed algorithm. Section 5 demonstrates experiment results on the proposed model in the case of different important parameters. Finally, Section 6 offers conclusions and some directions for future work.

2 Related work

In this section, we review related research efforts on the CIM problem, which analyses the implications of competing products interfering with each other. In addition, we review some reinforcement approaches in the CIM problem.

2.1 *Competitive influence maximisation*

CIM aims at finding strategies that maximise one's influence while minimising his opponents' influence in a social network (Masucci and Silva, 2014; Zhang et al., 2015). There are different extensions of the IC model and the LT model to accommodate multiple competing ideas in social networks instead of focusing on spreading a single ideas (Tong et al., 2018; Bharathi et al., 2007; Sun et al., 2018). Li et al. (2015a) consider a model for CIM. According to a graph $\mathcal{G}$ and diffusion model, the strategy space comprises all IM algorithms that players can adopt. For each player, the objective is to find a Nash equilibrium strategy that maximises his own influence. Rahaman and Hosein (2016) addressed a multi-stage version of the influence maximisation problem. They provided a new formulation and compared their approaches in terms of accuracy and computation run time.

2.2 *Reinforcement learning*

An important line of work that uses RL to solve NP-hard optimisation problems on graphs is (Zhu et al., 2019; Wei et al., 2021). Lin et al. (2015) model a multi-party CIM problem and propose a different model with the help of RL and based on the multi-round CIM method. Chung et al. (2019) propose a novel deep RL-based framework to tackle the MRCIM problem considering the network community structure under a quota-based ϵ-greedy policy. Ali et al. (2020) propose a deep RL-based model to tackle the CIM on unknown social networks. Harshavardhan et al. (2020) using automatically learned node and graph representations that encode important network structural properties, Kamarthi et al. propose a RL framework for discovering effective network sampling heuristics. Ali et al. (2018) propose a novel RL-based framework that is built on a nested Q-learning algorithm. They derive the optimal solution in both budget allocation and node selection that results in the maximum profit with time constraints.

2.3 Resource allocation against opponents

Parties in a CIM problem perform like a player in a Colonel Blotto game (CBG). CBGs are a class of two-player zero-sum games, in which both players need to allocate limited resources over several objects simultaneously. Masucci and Silva (2014) focused on competitive influence when players need to decide on resource allocation against their opponents. They proved that competition's price is unbound in such a CBG. Maehara et al. (2015) address the budget allocation scenario in maximisation influence problem. Companies can allocate different budgets to nodes in the network, and nodes will be attracted to companies whose products offer a higher value. In this case, companies compete by allocating a certain amount of budget to each node in the network. A Nash equilibrium-based model is proposed by Masucci and Silva (2017) to compete for obtaining more customers in online social networks.

Unlike most of the existing works, in this paper, we study the problem of multi-round CIM within budget constraints and while considering the remaining budget of opponents. We consider a different approach from the Blotto game for budget allocation strategy. There is a dependency between targets, and players can continue their investment in case of tie-breaking. In addition, there is propagation after any activation. In comparison with ML approaches, we consider new features to describe the state of the network.

3 Preliminaries

A social network can be modelled with a weighted and directed graph $\mathcal{G}(V, E, P, W)$, where we define W as a set of weights associated with each vertex in V. Activating a node u in $\mathcal{G}$ means accepting an idea from the player i. Once a node u accepts the idea of being occupied by a player i, it cannot change occupation to another party. If the given node does not accept any idea, it means that the state of the node u is inactive.

3.1 Competitive influence maximisation

In a multi-stage CIM problem, competitors need to select seed nodes simultaneously in each of the sequence stages. Suppose that there is a CIM game with two players, 1 and 2, and n nodes in a social network $\mathcal{G}$. Player 1 has a budget of size B_1, and player 2 has a budget size of B_2. Each node u has a value, $W(u) > 0$, which can be regarded as the reward of taking this node for players. The total value of n nodes in this social network is $W = \sum_{u \in V} W(u)$. The winner of this game would be the player who can obtain the most reward by influencing the more important nodes. Players have competition with the amount of budget they allocate in seed nodes (the most influential nodes).

In this game, three types of competition can occur. The first competition is players' competition on seed nodes by the amount of allocated budget, which can be called node-node competition. The second one is link-link, which is the competition of influence when two different links with different influences try to activate the given node in their favour. The last one is node-link. This will happen when one of the competitors allocates some budget on the given node, and the influence of another competitor reaches this node by the influence of the link.

3.1.1 Node-node influence competition

Considering a node-node competition on the node u. Suppose that $B_1(u)$ and $B_2(u)$ are the amount of budget that players 1 and 2 have allocated to node u. The winning probability of player 1 for this competition is as follows:

$$\frac{B_1(u)}{B_1(u) + B_2(u)} \tag{1}$$

3.1.2 Link-link influence competition

Link-link influence competition will happen after the budget allocation process and determining the winner of this stage in the case of taking the given seed node. During the propagation process, suppose that node u has the influence of player 1 from one of its neighbours with $p_1 = p(v, u)$. In addition, node u has influence of player 2 from another neighbour, node w, with $p_2 = p(w, u)$. The probability that node u would be activated by player 1 is as follows:

$$\frac{p_1}{(p_1 + p_2)} \times (1 - p_1 p_2), \tag{2}$$

where $(1 - p_1 p_2)$ considers the probability of activation of node u by at least one of the players. The probability that node u would be activated by player 2 is as follows:

$$\frac{p_2}{(p_1 + p_2)} \times (1 - p_1 p_2) \tag{3}$$

3.1.3 Node-link influence competition

In a multi-stage competition, competitors are able to allocate a budget at the same moment at the beginning of each stage rather than during the stage. At the beginning of each stage, competitors decide on their budget allocation, then influence propagation starts. At the end of the propagation, competitors can start the next stage and make a decision about new budget allocation. Therefore, there is node-node competition at the beginning of each stage and link-link competition during each stage. Consequently, we will avoid considering the link-node competition for the multi-stage CIM problem.

3.2 Multi-agent RL

In sequential games, players need to look forward and reason back to find the best decision. In simultaneous games, players look for the best response when they cannot see the other side's strategy. Therefore, players need to learn more about the strategies of opponents. RL is a subfield of machine learning that addresses the problem of learning optimal decisions over time. In RL, the agent keeps interacting with the environment to find the optimal policy π to maximise his expected accumulated rewards (Sutton et al., 1998). The goal of an RL is to learn a policy $\pi(s)$ to determine which action to take given a specific environment represented by state s.

The reward obtained by an agent should reinforce his behaviour. Reward reflects the success of the agent's recent activity and not all of the successes achieved by the agent so far. The agent's objective is to learn the policy that maximises the expected value of the return. The return is the measure of future cumulative reward during the rounds.

$$r_{t+1} + \gamma r_{t+2} + \gamma^2 r_{t+2} + ... = \sum\nolimits_{k=0}^{\infty} \gamma^k r_{t+k+1}. \qquad (4)$$

RL formulates the expected accumulated rewards of a state which is called the $\mathcal{V}$ function. Also, it formulates the expected accumulated rewards for each state-action pair which is called the Q function. Q function estimates how efficient the policy π is at maximising the accumulated reward r_t. The $\mathcal{V}$ function $\mathcal{V}^{\pi}(s)$ associated with a policy π tells the agent how good the policy is. The state-value function is defined as:

$$\begin{aligned} \mathcal{V}^{\pi}(s) &= E_{\pi}\{r_t | s_t = s\} \\ &= E_{\pi}\left\{\sum\nolimits_{k=0}^{\infty} \gamma^k r_{t+k+1} | s_t = s\right\}, \end{aligned} \qquad (5)$$

where γ is the discount factor. The action-value function $Q(s, a)$ is expected return starting from action a in state s, and then following policy π:

$$\begin{aligned} Q^{\pi}(s, a) &= E_{\pi}\{r_t | s_t = s, a_t = a\} \\ &= E_{\pi}\left\{\sum_{k=0}^{\infty} \gamma^k r_{t+k+1} | s_t = s, a_t = a\right\}. \end{aligned} \qquad (6)$$

The state value and action value in equations (5) and (6) can be learned through the interaction of agents with the environment. The optimal policy $\pi(s)$ can be obtained given the Q function and find the maximum value. Figure 3 displays the details of RL for a multi-round CIM. According to this diagram, at the end of each round, players can see the result of the competition in terms of reward and the current state. Then, they update their learning, compute new policy against the opponent's strategy, and select a new seed set. In Algorithm 1, we can see this process step by step. In the case of multi-agent RL, agents learn the policies through experience in the environment and interaction with each other. We assume that there are only two parties that compete with each other. We need first to define the environment, the reward, the action, and the state.

Figure 3 Approach during training process

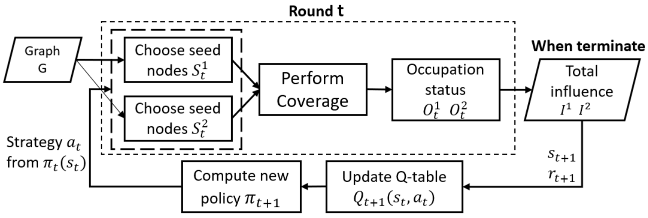
Round t
Graph G
Choose seed nodes S_t^1
Choose seed nodes S_t^2
Perform Coverage
Occupation status O_t^1 O_t^2
When terminate
Total influence I^1 I^2
Strategy a_t from $\pi_t(s_t)$
s_{t+1}
r_{t+1}
Compute new policy π_{t+1}
Update Q-table $Q_{t+1}(s_t, a_t)$

Algorithm 1 RL

```
$Q(s, a) \leftarrow$ initial value
while training is not terminal do
    $s_t \leftarrow s_0$
    while $s_t$ is not a terminal state do
        Determine $Q_t(s_t, a_t)$
        Take strategy $a_t$ based on $Q_t(s_t, a_t)$
        Simulate opponent's action
        Propagate influence to obtain reward $r_{t+1}$
        Compute next state $s_{t+1}$ based on network features
        Update $Q(s_t, a_t)$
        $s_t \leftarrow s_{t+1}$
```

Table 2 Computing $R(v)$ from source node v_2

A	*N(A)*	*R(s) ∗ p(s,v)*	*R(v)*
$\{v_2\}$	v_1	1 ∗ 0.2 = 0.2	$R(v_8)$
		1 ∗ 0.4 = 0.4	
$\{v_2, v_8\}$	v_1	1 ∗ 0.2 = 0.2	$R(v_7)$
	v_5	0.4 ∗ 0.1 = 0.04	
	v_7	**0.4 ∗ 0.7 = 0.28**	
$\{v_2, v_8, v_7\}$	v_1	**1 ∗ 0.2 = 0.2**	$R(v_1)$
	v_5	0.4 ∗ 0.1 = 0.04	
$\{v_2, v_8, v_7, v_1\}$	v_1	0.4 ∗ 0.1 = 0.04	$R(v_3)$
	v_5	**0.2 ∗ 0.5 = 1**	
		0.2 ∗ 0.4 = 0.08	
$\{v_2, v_8, v_7, v_1, v_3\}$	v_1	0.4 ∗ 0.1 = 0.04	$R(v_5)$
	v_5	0.2 ∗ 0.4 = 0.08	
		1 ∗ 0.1 = 0.1	
		1 ∗ 0.3 = 0.3	
$\{v_2, v_8, v_7, v_1, v_3, v_5\}$	v_4	1 ∗ 0.1 = 0.1	$R(v_6)$
	v_6	**0.3 ∗ 0.4 = 0.12**	
$\{v_2, v_8, v_7, v_1, v_3, v_5, v_6\}$	v_4	**1 ∗ 0.1 = 0.1**	$R(v_4)$
	v_9	0.12 ∗ 0.2 = 0.024	
$\{v_2, v_8, v_7, v_1, v_3, v_5, v_6, v_4\}$	v_9	**0.12 ∗ 0.2 = 0.024**	$R(v_9)$

In the given social network, considering V^1 and V^2 as the total number of activated nodes by players 1 and 2 respectively. The multi-round CIM can be considered as a zero-sum game for players 1 and 2 since $(V^1 -$

$V^2) + (V^2 - V^1) = 0$, where $V^1 - V^2$ and $V^2 - V^1$ are the goals of players 1 and 2, respectively. In such a game, the Nash equilibrium is guaranteed to exist with mixed strategies. The MINMAX theorem would be useful to find the equilibrium (Morgenstern and Von Neumann, 1953).

4 Methodology

Traditional RL has been successful in dealing with multi-round CIM (Lin et al., 2015). Nevertheless, this approach did not address the effect of budget on player seed selection strategy. In our approach, we integrate seed selection and budget allocation into the RL model. In the budget allocation phase, we consider convincing influential nodes to act as seeds, as well as selecting seed nodes. The player in this framework learns how to maximise the value of accumulated rewards by choosing the optimal policy π. The first step is identifying influential nodes within the network. Players then compete over only the selected nodes, rather than the entire network, depending on the budgets they allocate to each influential node. During each round t, the agent observes a set of features representing the network state $s_t \in S$, and selects one of the legal actions from the set a_t. In each round, the agent selects a seed set, $\xi_t \subset V$, based on its past observations. Note that ξ_t is the seed set selected by π at round t. The goal for the agent is to follow a learning policy π maximising the total number of activated nodes. When no budget remains or no node can be added to the seed set ξ, the algorithm terminates.

Definition 1 (budgeted multi-round CIM): Given the network $\mathcal{G}$, each player chooses seed nodes in turn, and then influence propagation is performed at round t. Players compete based upon the budgets they allocate to the most influential nodes in order to win these nodes as seed sets. The objective of each player is to maximise its overall relative influence V^i after T rounds, where V^i is the difference among activated nodes of different players.

Algorithm 2 Finding seed set by MRIP

```
$S \leftarrow \varnothing$
for all $u \in V$ do
    $w'(u) \leftarrow 0$
for $u \in V$ do
    Construct $\mathcal{T}_u$ via Algorithm 3
    for each leaf $v$ in reverse $\mathcal{T}_u$ do
        $z \leftarrow parent(v)$
        while $v \neq u$ do
            Compute $w'(z) = w'(z) + R(v) \times w(z)$
            $v \leftarrow z$
            $z \leftarrow parent(v)$
*new seed* $\leftarrow \arg\max_{u \in V/\ S} w'(u)$
$S \leftarrow S \cup \{$*new seed*$\}$
$V_A \leftarrow$ activated nodes by *new seed* node
Constructing $\mathcal{G}'$ with vertex set $V - V_A$
Recalculate $\mathcal{T}$ and $w'$ in $\mathcal{G}'$
```

Algorithm 3 Computing $\mathcal{T}_u$

```
Require: $\mathcal{G}(V, E, P)$, source node $u$
$A = \{u\}$, $R(u) = 1$
while $A \neq V$ do
    Find node $v \in N(A)$ and $v \in V - A$ such that
    $R'(v) = \max_{(s,v): s \in A, v \in V - A} R(s) \times p(s, v)$
    $R(v) = R'(v)$
    $A = A \cup \{v\}$
    Set $s$ as the parent of $v$ in spanning tree $\mathcal{T}_u$
return $\mathcal{T}_u$
```

4.1 Selecting seed nodes and propagation model

Consider a static social network $\mathcal{G}$ and B_1 and B_2 for players as their fixed budget. Each round, one seed node will be chosen. The goal of each player is to reach and activate as many nodes as possible within their total budget. Each player can decide to implement the specific strategy to maximise its overall influence in $\mathcal{G}$. The strategy refers to how the player spends their budget on selecting the seed nodes at each round. Maximal influence with a spanning tree restricts node u's influence diffusion to a local tree structure rooted at u. The influence of a node in a tree can be calculated efficiently and precisely. Note that the conflict rule is slightly different from other works. In contrast to other approaches, which prioritise one of the players or select the winner of conflicting randomly, our approach allows players to increase their investment in case of tie-breaking. The winner will be determined with the help of budget proportion.

4.2 Most reliable influence path

Since influence maximisation is NP-hard, we use the idea of the critical path in the scheduling community. Following the style of Dijkstra and Prim's greedy algorithm, an inactive node will get a chance to become active only through the shortest path from the initially active nodes. In order to find the shortest path in a maximum influence problem, we can consider the maximum influence probability of edges. The distance between node u and v can be computed as the logarithm of the inverse of the influence probability of edge (u, v). Influence propagates through the most probable paths, and the notion of the MRIP can be considered as an approximation. It is helpful to estimate the local influence of nodes for seed selection. The influence of each node when considering the most reliable paths that originate from the given node can be regarded as a new measurement for ranking nodes. This paper calls this value the weighted influence of each node u, $w'(u)$. Considering $R(v)$ as the influence value of the most reliable path on node v originated from the source u, we construct a spanning tree $\mathcal{T}$ with the most reliable paths helps us to find w' for all nodes.

In fact, Prim's algorithm allows us to determine the spanning-tree $\mathcal{T}_v$ rooted at v such that each node is reached from the source node v via MRIP. The value of $R(v)$ for any two nodes u and v in V is the value of the shortest path from node u to v, where $\mathcal{P}_{u,v} = (u = u_1, u_2, ..., u_m = v)$, and there is no duplicate nodes. The probability that node v is activated by u through the path $\mathcal{P}_{u,v}$ is calculated as $\prod_{i=1}^{m-1} p(u_i, u_{i+1})$. All nodes along the path from u to v need to be successfully activated, then node v would be activated. As an extension, to more efficiently compute the increased influence spread within the tolerance of error, we can use an influence threshold to filter out the insignificant maximal influence paths whose values are less than due to having a very small impact on the influence spread computation.

For all $v \in V$ in the $\mathcal{G}$, we need to find $\mathcal{T}$. For simplicity, we explain the process of computing w' just by considering node u as the tree's root node. Suppose that $R(u) = 1$, among all of the neighbours of node u finding the edge (u, v) with maximum $R(u) \times p(u, v)$ is the first step. This is a greedy algorithm. In each step, we consider all of the edges that the source of them is in the explored node set A, and its destination is in $V - A$. We continue this process until A includes all of the nodes in V. Algorithm 3 represents these steps in detail. After constructing the spanning-tree $\mathcal{T}_u$, we compute the influence weight of each node by traversing this tree reversely. For each node v, the parent of v is u, and the weight of node u, which is illustrated as the weighted influence, will be measured by:

$$w'(u) = \sum\nolimits_{\forall v \in V} R(v) \times w(u), \tag{7}$$

where $w(u)$ is the weight of node u and $R(v)$ presents the value of shortest originated from u to v. If we consider $w(u)$ as the *fixed weight* for node u, which can be show the importance of node in the case of degree or centrality, $w'(u)$ can be called *dynamic weight* of this node. Considering the example in Figure 4, for any pair of nodes u and v, we need to find the maximum influence path from u to v and construct a spanning tree $\mathcal{T}$. Figure 4 shows the process for node v_2. Table 2 presents some early steps of finding $R(v)$ for each nodes when v_2 is the source node. Using the calculated $R(v)$ and reverse traversing the $\mathcal{T}_{v_2}$ in Figure 4(b), the influence weights of all nodes are shown in Figure 4(c). The intuition behind the proposed algorithm comes from Dijisktra's algorithm. We can prove the proposed algorithm can find the most reliable path correctly.

Theorem 1: If $\mathcal{T}_s$ is the spanning tree selected by MRIP's algorithm for source node s in the social network $\mathcal{G} = (V, E, P, W)$, then $\mathcal{T}_s$ is a most reliable influence tree rooted in s in $\mathcal{G}$ and $R(v)$ for each node $v \in V$ shows the influence value of the most reliable path on node v.

Proof: In Figure 5, the grey area includes the explored nodes. Suppose that w is the next vertex added to $\mathcal{T}$ and P^* be the path from source s to destination w through node v. Considering any other path P from s to w, node x be the first node on path outside $\mathcal{T}$. Path P is already as long as P^* as soon as it reaches x by greedy choice. Thus, $R(w)$ is the length of the most reliable path from s to w. This completes the proof. ■

Figure 4 Computing shortest paths and influence weights, (a) original graph (b) constructing $\mathcal{T}_{v_2}$ (c) calculating w' (see online version for colours)

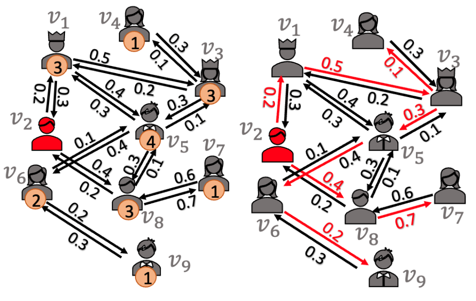

(a) (b)

$w'(v_1) = 3 * R(v_3) = 3$
$w'(v_2) = 2 * (R(v_1) + R(v_8)) = 1.2$
$w'(v_3) = 3 * (R(v_4) + R(v_5)) = 1.2$
$w'(v_4) = 0$
$w'(v_5) = 4 * R(v_6) = 0.48$
$w'(v_6) = 2 * R(v_9) = 0.048$
$w'(v_7) = 0$
$w'(v_8) = 3 * R(v_7) = 0.56$
$w'(v_9) = 0$

(c)

Figure 5 Correctness of MRIP algorithm (see online version for colours)

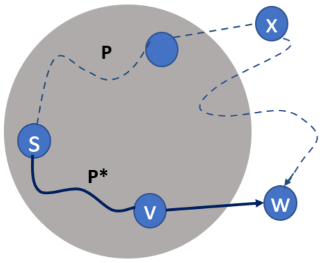

Algorithm 2 presents the processes of selecting the seed node based on the influence weight of nodes. After finding $\mathcal{T}_v$ for each node v in Algorithm 3, by considering $w(v)$ of nodes as the weight of node or ranking measurement in the case of the importance of node and $R(v)$ as the value of the most reliable path, the influence weights $w'(v)$ of all of the nodes can be calculated. The node with the highest w' would be selected as the seed node in each round. After choosing a seed node and propagating its influence, the next step is to recalculate $\mathcal{T}$ and the weighted influence of nodes in the graph $\mathcal{G}'$ with $V - V_A$ nodes, where V_A is the set of activated nodes. Therefore, after selecting any seed node and the propagation process, there are new w's for nodes. That is because we called this weight as *dynamic weight*.

The time complexity of Dijkstra's algorithm is $O(|E| \cdot log|V|)$, but here we need to find the shortest paths for all pairs of nodes. Now, the time complexity becomes $O(|E|^2 \cdot log|V|)$. After selecting a seed node, we need to remove the activated nodes, V_A, from $\mathcal{G}$ and consider a new social network $\mathcal{G}'$ including the set of nodes $V - V_A$ to recalculate $\mathcal{T}_v$ for each node v as well as new weight w'. Therefore, the total number of nodes in these paths should be considered in the algorithm's time complexity as well.

4.3 RL settings

As we are considering a multi-round scenario, the opponent's past decisions can be taken into account, but the opponent's future decisions are not known. There are several parameters we need to define in order to implement RL. The propagation of influence is treated as an environmental effect, whereby activated nodes spread their influence to their neighbours and activate new ones. The reward we receive after T steps is the number of nodes that have been influenced in the entire graph. Through Q-function updates, rewards are propagated back to previous states.

4.3.1 Action

Players can allocate different amounts of budget to nodes in $\mathcal{G}$. Competition is based on how much budget each player allocates to each node. The possible actions are allocating budget on new seed nodes or feeding an activated seed node to increase its influence on neighbours. We use the idea of meta-learning (Lin et al., 2015; Ali et al., 2018) in RL. We consider the following actions:

1 selecting a new seed node

2 feeding a node in case of tie.

Selecting seed nodes can include max-degree, max-weight, centrality, randomly, voting, and learning-based strategies. In case of investment, we consider investing \$1 or all of the remaining budget.

4.3.2 State

In order to represent the network and environment status, we must model the state. The design of features will reflect both the current status of the network and the current occupation status. Correlations with rewards, the choice of actions, and the condition of networks require certain features. Below are the features we have designed:

1 number of inactive nodes

2 summation of degrees of all inactive nodes

3 maximum degree among all inactive nodes

4 summation of the weight of the edges for which both vertices are inactive

5 summation of the inactive out-edge weight for nodes which are the neighbours of player i

6 maximum sum of the inactive out-edge weight of a node among all nodes

7 ratio of budgets

8 weight of nodes in case of reachability.

Features 1 to 5 help players find the condition of network in terms of the status of nodes as well as the weight of edges. Features 6, 7, and 8 are new ones to describe the states of the network. These features help players to learn more about the environment, as well as the opponent's strategy. As a result of the dependence between some features, not all combinations of states are possible. There is a correlation between the candidate strategies we use to choose our actions and these features. The player continually updates both Q-tables, that is, seed-selection, and budget-allocation Q-tables, during the training. Meanwhile, it updates its policy throughout the training in order to find an optimal policy for budget utilisation from budget-allocation and seed-selection Q-tables.

Table 3 Q-table for strategies

(a) Budget-allocation		
State	*Budget*	*Q-value*
33311001	Unit	0.7
33311001	All	0.2

(b) Seed-selection		
State	*Seed*	*Q-value*
333110010	Degree	0.26
333110010	Weight	0.24
333110010	MRIP	0.3
333110010	Compete	0.6
333110011	Degree	0.26
333110011	Weight	0.24
333110011	MRIP	0.3
333110011	Compete	0.6

5 Experiments

We conducted experiments to evaluate the efficiency of the proposed models in terms of influence spread to other algorithms. Also, we evaluate our algorithm for different datasets with different densities.

5.1 Experiment setup

We used the igraph Python library to represent the graphs and the shortest path calculations. The datasets consist of two real-world social networks and two synthetic ones. We used the IC as the diffusion model. The edge weights are set randomly in a range between 0 and 1. In order to check the impact of influence propagation, we consider normal distribution, with the same $\mu = 0$ and different σ^2. We train the model by doing 1,000 runs and then selecting the best result as the final result of the model. We use the random tree generation algorithm, as discussed in experimental settings.

Table 4 Social networks

Name	*Nodes*	*Edges*	*Description*
Facebook	4,039	88,234	Facebook social network
Ca-HepTh	9,877	51,971	Arxiv high energy physics
Cit-HepPh	620	827	Paper citation network
Synthetic	100	500	Randomly generated network

Table 5 Evaluation of different features

Dataset	*Reward*		*Dataset*	*Reward*	
Facebook	OPT-F6	49%	Synthetic	OPT-F6	45%
	OPT-F7	52%		OPT-F7	53%
	OPT-F8	55%		OPT-F8	50%
	OPT-F6F7	58%		OPT-F6F7	48%
	OPT-F7F8	56%		OPT-F7F8	58%
	OPT-F6F8	58%		OPT-F6F8	51%
	OPT	65%		OPT	68%

5.2 Comparison methods

To find the performance of our approach, we consider different baseline IM methods and the state-of-the-art multi-round competitive approach, which is called *STORM* (Lin et al., 2015). *OPT* is the name of the current paper's approach, which selects seed nodes based on both the fixed and dynamic weight of nodes. We consider the following approaches:

- STORM: A RL-based algorithm that finds an optimal seed selection using Q-learning.
- Max-degree: Traditional influence maximisation strategy as the algorithm selects nodes with the highest degree in the network as seed nodes.
- Centrality: This strategy select seed nodes based on the location of nodes in the network.
- Max-weight: One of the baseline methods that finds the seed node based on the maximum summation of out-edge weights.
- Random: This strategy is a baseline algorithm that randomly chooses one of the seed selection methods.
- Voting: This method lets the other three strategies vote for a node as the seed node.
- MRIP: The algorithm selects seed nodes based on both the node's fixed and dynamic weight.
- OPT: Our proposed learning-based approach.

5.3 Experiment results

We compare the influence spread of different algorithms on real-world datasets. Table 4 shows the details of these real datasets, which are accessible from http://snap.stanford.edu/data/. Each round is defined as players choosing a seed node and propagating influence. The number of active nodes after the diffusion process is used to evaluate the effectiveness of influence maximisation algorithms. We consider the evaluation of our approach in the cases of different budgets, network structures, competing strategies, and ranges for the weight of the edges. Table 5 shows the evaluation of approaches in the case of different combinations of features. *OPT-F6* is the approach that we do not consider features 7 and 8. Similarly, others show the approaches with different features. It can be seen from Table 5 with the three features 6, 7, and 8 there is the best result in real datasets. We call our approach as *OPT*.

5.3.1 Evaluation on budget setting

In the first experiment, we examine the effectiveness of the proposed models' performances in terms of reward by assuming players have a different budget. We consider a fixed budget for one of the players, then analyse the result of competition with a varied amount of budget for the opponent side. Clearly, the larger the budget, the more the increase of spread. Figures 6(a) and 6(b) shows the result of this experiment for the network with a topology that is like a tree. Figures 6(c) and 6(d) display the result in a network with a fat-tree topology. Figures 6(a) and (c) illustrate the effect of varying budget for player 2 when player 1 has a fixed budget of $10 for three algorithms of *random*, *STORM*, and *OPT*. Moreover, Figures 6(b) and 6(d) present the effect of varying budget for player 1 while player 2 has a fixed budget of $10 on the spread of player 1's influence. It should be noted that we have trained the models by assuming both parties have the same budget. It can be seen from the figures that *OPT* achieves better performance in comparison with other models.

5.3.2 Evaluation based on different topologies

We examine the effectiveness of the proposed models' performances on different networks in terms of reward by assuming players have different budgets. We consider a fixed budget for one of the players, then analyse the competition result with a varied budget for the opponent side. Clearly, the larger the budget, the more the spread increases. It should be noted that we have trained the models by assuming both parties have the same budget. It can be seen from the figures that *OPT* achieves better performance in comparison with other models. Also, we illustrate the performance of the proposed framework on networks with different structures. It can been seen from Figure 7 that in different real datasets with different topologies, *OPT* has better results than *STORM*. In addition, *OPT* can find more rewards when a player has a higher amount of budget.

Figure 6 Evaluation of player 1's reward with varying budget setting in synthetic dataset with different height, (a) fixed player 1 budget 10 (b) fixed player 2 budget 10 (c) fixed player 1 budget 10 (d) fixed player 2 budget 10 (see online version for colours)

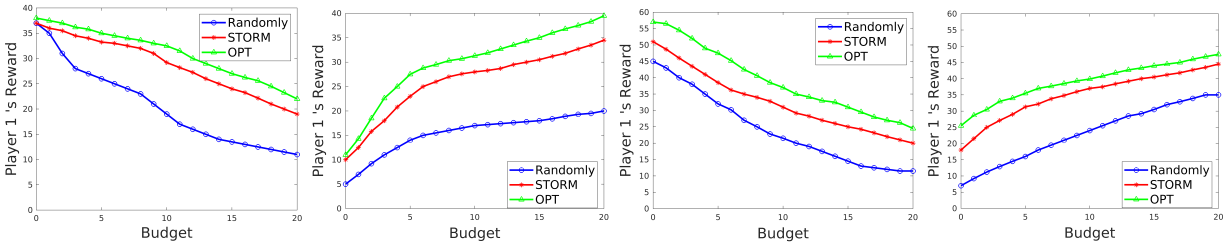

(a) (b) (c) (d)

Figure 7 Evaluation of player 1's reward with varying budget setting in real datasets, (a) Facebook (b) Ca-HepTh (c) Cit-Hepph (d) average all networks (see online version for colours)

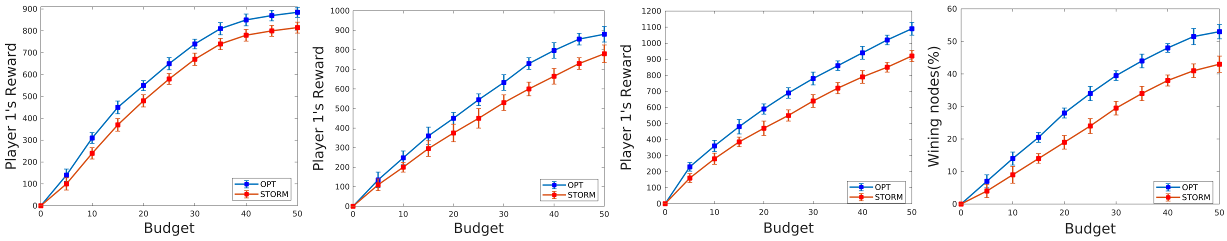

(a) (b) (c) (d)

Figure 8 Evaluation of player 1's reward with different influence distributions in average all network, (a) edge weights in (0.4, 0.7) (b) edge weights in (0.1, 0.4) (c) edge weight *U* (0, 0.2) (d) edge weight *U* (0, 1) (see online version for colours)

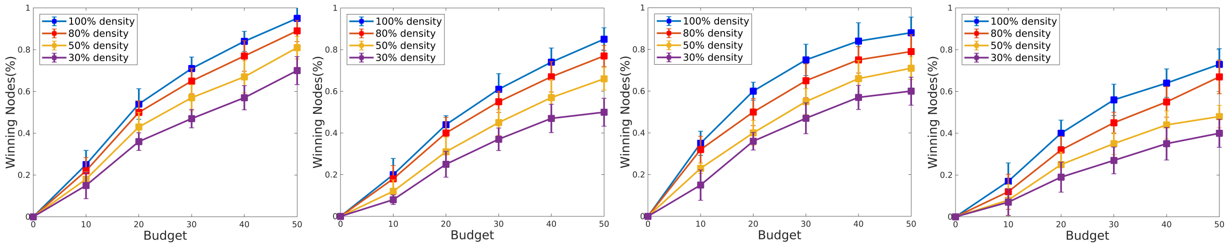

(a) (b) (c) (d)

Figure 9 Evaluation of player 1's reward with competing strategies in case of different budget, (a) Facebook (b) Ca-HepTh (c) Cit-Hepph (d) average all networks (see online version for colours)

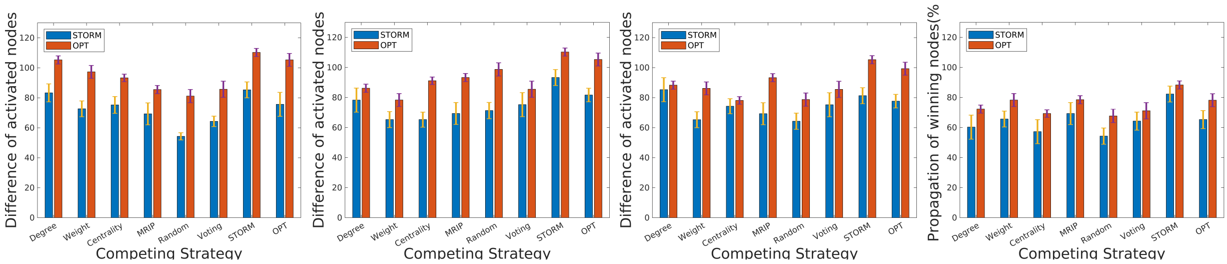

(a) (b) (c) (d)

5.3.3 *Evaluation on edge-weight setting*

We analyse the effect of different edge-weight settings on the proposed model. We consider the weight of the edges in the range of [0.1, 0.4] and [0.4, 0.7]. In addition, the weight for edges are randomly sampled from the normal distribution of *U* (0, 0.2) and *U* (0, 1). In addition, we consider different densities for the network to evaluate the performance of the approach in the case of the sparsity of the network. From Figure 8 can observe that the influence will diffuse more nodes when there are higher weights for edges. That happens because seed nodes can affect mode nodes. Also, the results show that *OPT* performs better if there is a high-density network.

5.3.4 Evaluation on different competing strategies

We evaluate our approach for player 1 against a competitor with a different strategies such as *degree*, *weight*, *MRIP*, as well as the learned-based strategy *STORM*. In this part of experiment, we consider some baseline strategies such as *degree*, *weight*, *MRIP*, as well as the learned-based strategy *STORM*. For example, in the second to last column in Figure 9, the blue one shows the result of the competition when players 1 and 2 have *STORM* approach, and the red one shows the result of the competition when player 1 uses *OPT* and player 2 uses *STORM* approaches. The blue one in the last column in Figure 9 shows the results when player 1 uses *STORM* and player 2 uses *OPT* approach. The red one shows the result when players 1 and 2 has *OPT* approaches. We can conclude from Figure 9 that *OPT* has the best performance against all the competing strategies, even against the *STORM* which is the learned-based model. According to the result of this experiment, based on the network structure, there are different results with baseline competing strategies.

In summary, according to the results of experiments considering the new extra features to describe the state of the environment when there is a budget constraint for the players would be helpful to find better final rewards. In the case of different datasets with different topologies and different numbers of nodes and edges, MRIP helps players have maximum influence against the opponent's propagated influence. Considering the total amount of budget, density of dataset, weight distribution of edges, and competitions strategies as the different parameters of simulations illustrate this learning approach is helpful in different network structures, influence probabilities distributions, and different amounts of budget for players.

6 Conclusions

In this work, we propose a RL framework to tackle the multi-round CIM problem considering budget ratio for players. A large body of related research did not focus on the impact of different budgets for players in a CIM problem. We look into identifying the set of seed nodes to maximise the spread by considering opponents' capabilities. In fact, our framework considers the combination of seed-selection and budget-allocation strategies to invest the budget efficiently to achieve better rewards considering budget constraints. To summarise, our main contribution is designing and evaluating a budgeted learned-based framework that handles the multi-round CIM. Our experimental results show that our approach successfully increases the influence on the given network compared to some known baseline approaches and a learned-based CIM approach. One possible future research is to investigate whether it is possible to accelerate the process of learning and study which parameters have significant in fact in the speed of learning. We also plan to study the partial-observed MDP (POMDP) algorithms to handle the players' partial information about the environment, opponents' strategies, and diffusion process.

Acknowledgements

This research was supported in part by NSF grants CNS 2128378, CNS 2107014, CNS 1824440, CNS 1828363, and CNS 1757533.

References

[online] http://snap.stanford.edu/data/ (accessed August 2021).

Ali, K., Wang, C-Y. and Chen, Y-S. (2018) 'A novel nested q-learning method to tackle time-constrained competitive influence maximization', *IEEE Access*, Vol. 7, pp.6337–6352.

Ali, K., Wang, C-Y., Yeh, M-Y. and Chen, Y-S. (2020) 'Addressing competitive influence maximization on unknown social network with deep reinforcement learning', in *Proc. of ASONAM*, IEEE, pp.196–203.

Bharathi, S., Kempe, D. and Salek, M. (2007) 'Competitive influence maximization in social networks', in *International Workshop on Web and Internet Economics*, Springer, pp.306–311.

Chung, T-Y., Ali, K. and Wang, C-Y. (2019) 'Deep reinforcement learning-based approach to tackle competitive influence maximization', in *Proc. of MLG Workshop*.

Domingos, P. and Richardson, M. (2001) 'Mining the network value of customers', in *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp.57–66.

Granovetter, M. (1978) 'Threshold models of collective behavior', *American Journal of Sociology*, Vol. 83, No. 6, pp.1420–1443.

Harshavardhan, K., Priyesh, V., Bryan, W., Balaraman, R. and Milind, T. (2020) 'Influence maximization in unknown social networks: learning policies for effective graph sampling', in *Proc. of Autonomous Agents and MultiAgent Systems*.

Huang, H., Meng, Z. and Shen, H. (2021) 'Competitive and complementary influence maximization in social network: a follower's perspective', *Knowledge-Based Systems*, Vol. 213, p.106600.

Kempe, D., Kleinberg, J. and Tardos, É. (2003) 'Maximizing the spread of influence through a social network', in *Proc. of ACM SIGKDD*, pp.137–146.

Kempe, D., Kleinberg, J. and Tardos, É. (2005) 'Influential nodes in a diffusion model for social networks', in *International Colloquium on Automata, Languages, and Programming*, pp.1127–1138.

Li, H., Bhowmick, S.S., Cui, J., Gao, Y. and Ma, J. (2015a) 'GetReal: towards realistic selection of influence maximization strategies in competitive networks', in *Proc. of ACM SIGMOD*, pp.1525–1537.

Li, H., Cui, J-T. and Ma, J-F. (2015b) 'Social influence study in online networks: a three-level review', *Journal of Computer Science and Technology*, Vol. 30, No. 1, pp.184–199.

Lin, S-C., Lin, S-D. and Chen, M-S. (2015) 'A learning-based framework to handle multi-round multi-party influence maximization on social networks', in *Proc. of ACM SIGKDD*.

Liontis, K. and Pitoura, E. (2016) *Boosting Nodes for Improving the Spread of Influence*, arXiv preprint arXiv:1609.03478.

Maehara, T., Yabe, A. and Kawarabayashi, K. (2015) 'Budget allocation problem with multiple advertisers: a game theoretic view', in *Proc. of Machine Learning*, PMLR, pp.428–437.

Masucci, A.M. and Silva, A. (2014) 'Strategic resource allocation for competitive influence in social networks', in *Proc. of Communication, Control, and Computing (Allerton)*, IEEE, pp.951–958.

Masucci, A.M. and Silva, A. (2017) 'Advertising competitions in social networks', in *Proc. of ACC*, IEEE, pp.4619–4624.

Molinero, X. and Riquelme, F. (2021) 'Influence decision models: from cooperative game theory to social network analysis', *Computer Science Review*, Vol. 39, p.100343.

Morgenstern, O. and Von Neumann, J. (1953) *Theory of Games and Economic Behavior*, Princeton University Press, New Jersey.

Nguyen, H. and Zheng, R. (2013) 'On budgeted influence maximization in social networks', *IEEE Journal on Selected Areas in Communications*, Vol. 31, No. 6, pp.1084–1094.

Niknami, N. and Wu, J. (2022) 'A budged framework to model a multi-round competitive influence maximization problem', in *Proc. of IEEE International Conference on Communications (ICC)*.

Ohsaka, N. (2020) 'The solution distribution of influence maximization: a high-level experimental study on three algorithmic approaches', in *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, pp.2151–2166.

Rahaman, I. and Hosein, P. (2016) 'On the multi-stage influence maximization problem', in *Proc. of LA-CCI*, IEEE, pp.1–6.

Shirazipourazad, S., Bogard, B., Vachhani, H., Sen, A. and Horn, P. (2012) 'Influence propagation in adversarial setting: how to defeat competition with least amount of investment', in *Proc. of Information and Knowledge Management*, pp.585–594.

Sun, L., Huang, W., Yu, P.S. and Chen, W. (2018) 'Multi-round influence maximization', in *Proc. of ACM SIGKDD*.

Sutton, R.S., Barto, A.G. et al. (1998) *Introduction to Reinforcement Learning*, Vol. 135, MIT Press, Cambridge.

Tong, G., Wu, W. and Du, D-Z. (2018) 'Distributed rumor blocking with multiple positive cascades', *IEEE Transactions on Computational Social Systems*, Vol. 5, No. 2, pp.468–480.

Wang, X., Tong, X., Fan, H., Wang, C., Li, J. and Wang, X. (2021) 'Multi-community influence maximization in device-to-device social networks', *Knowledge-Based Systems*, Vol. 221, p.106944.

Wei, Y., Zhang, L., Zhang, R., Si, S., Zhang, H. and Carin, L. (2021) *Reinforcement Learning for Flexibility Design Problems*, arXiv preprint arXiv:2101.00355.

Wu, H., Liu, W., Yue, K., Huang, W. and Yang, K. (2015) 'Maximizing the spread of competitive influence in a social network oriented to viral marketing', in *Proc. of Web-Age Information Management*, Springer, pp.516–519.

Zhang, H., Procaccia, A.D. and Vorobeychik, Y. (2015) 'Dynamic influence maximization under increasing returns to scale', in *Proc. of Autonomous Agents and Multiagent Systems*, pp.949–957.

Zheng, H. and Wu, J. (2017) 'Friend recommendation in online social networks: perspective of social influence maximization', in *Proc. of ICCCN*, IEEE, pp.1–9.

Zhu, S., Ng, I. and Chen, Z. (2019) *Causal Discovery with Reinforcement Learning*, arXiv preprint arXiv:1906.04477.