

A Learning Approach to Complex Contagion Influence Maximization

Extended Abstract

Haipeng Chen
William & Mary
hchen23@wm.edu

Bryan Wilder
Carnegie Mellon University
bwilder@andrew.cmu.edu

Wei Qiu
Nanyang Technological University
qiuw0008@e.ntu.edu.sg

Bo An
Nanyang Technological University
boan@ntu.edu.sg

Eric Rice
University of Southern California
ericr@usc.edu

Milind Tambe
Harvard University
milind_tambe@harvard.edu

ABSTRACT

Influence maximization (IM) aims to find a set of seed nodes in a social network that maximizes the influence spread. While most IM problems focus on classical influence cascades (e.g., Independent Cascade and Linear Threshold) which assume individual influence cascade probability is independent of the number of neighbors, recent studies by sociologists show that many influence cascades follow a pattern called complex contagion (CC), where influence cascade probability is much higher when more neighbors are influenced. Nonetheless, there are very limited studies on complex contagion influence maximization (CCIM) problems. This is partly because CC is *non-submodular*, the solution of which has been an open challenge. In this study, we propose the first reinforcement learning (RL) approach to CCIM. We find that a key obstacle in applying existing RL approaches to CCIM is the reward sparseness issue, which comes from two distinct sources. We then design a new RL algorithm that uses the CCIM problem structure to address the issue. Empirical results show that our approach achieves the state-of-the-art performance on four real-world networks.

KEYWORDS

Influence Maximization; Complex Contagion; Non-submodularity; Combinatorial Optimization; Reinforcement Learning

ACM Reference Format:

Haipeng Chen, Bryan Wilder, Wei Qiu, Bo An, Eric Rice, and Milind Tambe. 2023. A Learning Approach to Complex Contagion Influence Maximization: Extended Abstract. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 3 pages.

1 INTRODUCTION & BACKGROUND

We study the problem of influence maximization with a complex contagion model, where the cascade probability to a target node has a non-concave jump when the number of activated neighbors exceeds a threshold. Formally, a social network is represented as a graph $G = (V, E)$, where V and E are the nodes and edges, respectively. Each node is either *activated* or *inactivated*, which means the node is influenced or not. We assume all nodes are initially inactivated unless chosen as the seed node. Nodes which are linked

by edges have a probability p of influencing each other. For each node v , its neighbors are represented as $\mathcal{N}(v)$. In simple contagions such as the Independent Cascade (IC) [13] and the Linear Threshold (LT) [14], the probability p is assumed to be a constant that is independent of the number of its neighbors. In complex contagion, the assumption is relaxed, and p is represented as a dependent variable of the number of activated neighbors k . Without loss of generality, we consider the classical K -complex contagion model [5, 6, 11]:

$$p(k) = \begin{cases} p_0, & \text{if } k \geq K \\ p_1, & \text{if } k < K \end{cases} \quad (1)$$

where $0 \leq p_1 \ll p_0 \leq 1$. $K \geq 1$ is an integer threshold value. K is interpreted as the threshold to make a qualitative change to $p(k)$. Figure 1 shows the cascade probability (y-axis) p given numbers of activated neighbors (x-axis) in simple vs complex contagions ($K = 3$).

Given a set of seed nodes $S \subseteq V$, we represent the influence of complex contagion as $\sigma(G, S)$. Therefore, Complex Contagion Influence Maximization (CCIM) problem is to select the optimal set of seeds S^* given a budget $|S| \leq T$, such that the total influence $\sigma(G, S)$ is maximized:

$$S^* = \arg \max_{S \subseteq V} \sigma(G, S) \quad (2)$$

Despite evidence that many influence cascades display complex dynamics [2, 10, 20, 24], there are very limited studies on CCIM. Mainstream IM algorithms rely crucially on a property of simple contagions called *submodularity*, which reflects diminishing returns to the selection of additional seeds (the blue curve in Figure 1). Under submodularity, a simple greedy algorithm is highly effective [16]. However, the CC models' surge of cascade probability in adoption probability violates submodularity. This greatly complicates optimization by introducing complementarities: the marginal gain to selecting any single seed is small since its value is only revealed in combination with a specific set of other seeds. Formally, this means that CCIM is NP-hard to approximate even under simplified network models [23]. To our knowledge, the only practical




Figure 1: Simple (green) v.s. complex (orange) contagions.

algorithm for CCIM thus far is the Dynamic Programming for Influence Maximization (DPIM) approach of [1]. However, DPIM heavily relies on the assumption that the network has a hierarchical structure and its effectiveness may be limited when a good hierarchical decomposition cannot be found.

2 APPROACH & RESULTS

In this work, we treat CCIM as a stochastic combinatorial optimization problem (COP) with a non-submodular objective function. Inspired by recent works that combine RL and graph representation techniques to address COPs on graphs [4, 7, 9, 17, 19] and influence maximization in particular [8, 18], we design a new RL algorithm, Reinforcement Learning for IM with Complex Contagion (RL4IM-CC) to solve the CCIM problem.

The underlying idea of RL-based approaches to COPs is to decompose the original seed *set* selection into a *sequence* of seeds, and it does so greedily based on the marginal “score” of each node. The scores are estimated using deep function approximators. Despite the initial success of applying RL for IM [8, 18], we found that directly applying these methods to the CCIM problems often yields suboptimal performance, mainly because the reward signal in CCIM is much more sparse than simple contagion IM. A major challenge is that the reward signal in CCIM is much more sparse than simple contagion IM. We identify two distinct causes of reward sparseness. First, the effective solution space (i.e., solutions with non-negligible influence) is much smaller. This is because CC requires a harsher condition for influence to spread. Second, the marginal contribution of each node (action) in CCIM is more sparse than regular IM, because the reward becomes non-negligible only after multiple seeds are selected. This yields small or zero reward at the first few time steps even with the credit assignment mechanism used in previous works. We refer to the two causes of reward sparseness as *effective solution sparseness* and *credit sparseness*.

Our algorithm addresses the above issues with key innovations including (i) a solution filtering step that yields more effective policy exploration, (ii) an adapted return-based prioritized experience replay (PER) [22] component that increases the chances of sampling training transitions with higher rewards, and (iii) a new reward-shaping component which, at the end of each training episode, assigns an additional reward to nodes by looking at their marginal contribution w.r.t. the global action sequence.

We present evaluations of RL4IM-CC and baselines on 4 publicly available real-world networks, including Football [12], Polbooks [21], India [3], and Exhibition [15]. The baselines include (i) **Random+**, which selects seeds randomly from the effective candidate set $\mathcal{A}$, (ii) **Greedy** [16], which is the prominent approach for submodular IM, (iii) **DPIM** [1], which is a dynamic programming based method designed for non-submodular IM including CCIM. Note that the code of DPIM is not publicly available. Therefore we implemented our own version of DPIM, where to build a hierarchical decomposition tree of the original network, we use the Jaccard Similarity variant (see Section 4.1.3 of [1]). (iv) **RL4IM** [8], which is a recent work that uses RL to address submodular IM.

We set propagation probabilities $p_1 = 0$, $p_0 = 1$, so that the propagation is deterministic. This setting is convenient for comparison as it rules out the factor that arises from stochasticity of different runs.

The threshold K values are set to 4 for smaller networks (Football, Polbooks and India), and 6 for the larger network Exhibition. The values are set as such so that the influence spread is neither too sparse (very few nodes are influenced) nor too dense (most nodes are influenced) where comparisons become trivial. We set the seed budget $T = 8$ for all the networks except for the largest network Exhibition which is $T = 12$. In deterministic settings, there is no randomness for Greedy and DPIM. Randomness of Random+ and RL methods arises from different running seeds. For Random+, we run 50 times each and take the average. For RL methods, we run 15 times and find the best model via a separate validation process (performs every 20 training time steps), and report its performance.

Table 1: Normalized influence of different methods on 4 real-world networks. Best results are highlighted.

Method \ Network	Football	Polbooks	India	Exhibition
Greedy	0.3217	0.4571	0.0396	0.0293
DPIM	0.3913	0.6762	0.3911	0.5537
Random+	0.0926	0.1486	0.0498	0.0293
RL4IM	0.513	0.7429	0.3465	0.5439
RL4IM-CC	1.0	0.819	0.5347	0.5585

Table 1 shows the results. We have the following observations.

(i) *RL4IM-CC consistently obtains the best results among all the methods across all the networks.* It obtains the optimal or close-to-optimal solutions on small networks.

(ii) The Random+ method, though being far from optimal, seems to be sufficient to serve as a warm-start for RL4IM-CC.

(iii) Surprisingly, *Greedy can be arbitrarily bad.* For example, Greedy obtains only an influence value of 0.0396 on the India network. Considering that there are 202 nodes in this network, this means that Greedy does not find any effective solution.

(iv) *RL4IM and DPIM are unstable.* RL4IM performs close to RL4IM-CC on some networks (e.g., and Exhibition). However, on the other networks (Football, Polbooks, and India), it is significantly beaten by RL4IM-CC. Our hypothesis is that the effective solution sparseness becomes more severe when the entire solution space grows larger. The same holds for DPIM, where it works well on Exhibition, but is much worse than RL4IM-CC on other networks. A main reason, in our understanding, is that DPIM highly relies on the assumption that networks are hierarchically structured, and therefore the performance of it highly depends on how the network’s structure is aligned with the assumption. Unfortunately in practice, it is hard to measure the hierarchy of networks.

Conclusion: We propose the first learning-based approach to CCIM, with innovative components addressing the reward sparseness issues that uniquely arise from CCIM. Empirical results show that our approach achieves new state-of-the-art performance on 4 real-world social networks. Our work opens up many potential future directions for learning-based approach to CCIM. For example, it is interesting to explore: (i) Can we design more efficient RL algorithms for larger networks? (ii) Can the learned RL policies generalize to new networks? (iii) What if the network structures or the complex contagion model parameters are uncertain?

ACKNOWLEDGMENTS

Bo An is supported by the Ministry of Education, Singapore, under its Academic Research Fund Tier 1 (RG13/22).

REFERENCES

- [1] Rico Angell and Grant Schoenebeck. 2017. Don't be greedy: Leveraging community structure to find high quality seed sets for influence maximization. In *WINE*. Springer, 16–29.
- [2] Lars Backstrom, Dan Huttenlocher, Jon Kleinberg, and Xiangyang Lan. 2006. Group formation in large social networks: Membership, growth, and evolution. In *KDD*. 44–54.
- [3] Abhijit Banerjee, Arun G Chandrasekhar, Esther Duflo, and Matthew O Jackson. 2013. The diffusion of microfinance. *Science* 341, 6144 (2013).
- [4] Yoshua Bengio, Andrea Lodi, and Antoine Prouvost. 2021. Machine learning for combinatorial optimization: A methodological tour d'horizon. *European Journal of Operational Research* 290, 2 (2021), 405–421.
- [5] Damon Centola. 2010. The spread of behavior in an online social network experiment. *Science* 329, 5996 (2010), 1194–1197.
- [6] Damon Centola and Michael Macy. 2007. Complex contagions and the weakness of long ties. *Amer. J. Sociology* 113, 3 (2007), 702–734.
- [7] Haipeng Chen, Bo An, Guni Sharon, Josiah Hanna, Peter Stone, Chunyan Miao, and Yeng Soh. 2018. Dyet: Dynamic electronic toll collection for traffic congestion alleviation. In *AAAI*, Vol. 32.
- [8] Haipeng Chen, Wei Qiu, Han-Ching Ou, Bo An, and Milind Tambe. 2021. Contingency-aware influence maximization: A reinforcement learning approach. *arXiv preprint arXiv:2106.07039* (2021).
- [9] Michel Deudon, Pierre Cournut, Alexandre Lacoste, Yossiri Adulyasak, and Louis-Martin Rousseau. 2018. Learning heuristics for the TSP by policy gradient. In *CPAIOR*. 170–181.
- [10] Jie Gao, Golnaz Ghasemi, Jason J Jones, and Grant Schoenebeck. 2019. Complex contagions in charitable donations. *SocArXiv* (2019).
- [11] Golnaz Ghasemiesfeh, Roozbeh Ebrahimi, and Jie Gao. 2013. Complex contagion and the weakness of long ties in social networks: revisited. In *Proceedings of the fourteenth ACM conference on Electronic Commerce*. 507–524.
- [12] Michelle Girvan and Mark EJ Newman. 2002. Community structure in social and biological networks. *Proceedings of the national academy of sciences* 99, 12 (2002), 7821–7826.
- [13] Jacob Goldenberg, Barak Libai, and Eitan Muller. 2001. Talk of the network: A complex systems look at the underlying process of word-of-mouth. *Marketing Letters* 12, 3 (2001), 211–223.
- [14] Mark Granovetter. 1978. Threshold models of collective behavior. *Amer. J. Sociology* 83, 6 (1978), 1420–1443.
- [15] L. Isella, J. Stehlé, A. Barrat, C. Cattuto, J.-F. Pinton, and W. Van den Broeck. 2011. What's in a crowd? Analysis of face-to-face behavioral networks. *J. Theor. Biol.* (2011), 166.
- [16] David Kempe, Jon Kleinberg, and Éva Tardos. 2003. Maximizing the spread of influence through a social network. In *KDD*. 137–146.
- [17] Elias Khalil, Hanjun Dai, Yuyu Zhang, Bistra Dilkina, and Le Song. 2017. Learning combinatorial optimization algorithms over graphs. In *NeurIPS*. 6348–6358.
- [18] Sahil Manchanda, Akash Mittal, Anuj Dhawan, Sourav Medya, Sayan Ranu, and Ambuj Singh. 2020. GCOMB: Learning budget-constrained combinatorial algorithms over billion-sized graphs. *NeurIPS* 33 (2020).
- [19] Han-Ching Ou, Haipeng Chen, Shahin Jabbari, and Milind Tambe. 2021. Active screening for recurrent diseases: A reinforcement learning approach. In *AAMAS*. 992–1000.
- [20] Daniel M Romero, Brendan Meeder, and Jon Kleinberg. 2011. Differences in the mechanics of information diffusion across topics: Idioms, political hashtags, and complex contagion on twitter. In *WWW*. 695–704.
- [21] Ryan A. Rossi and Nesreen K. Ahmed. 2015. The network data repository with interactive graph analytics and visualization. In *AAAI*. <https://networkrepository.com>
- [22] Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. 2016. Prioritized experience replay. In *ICLR*.
- [23] Grant Schoenebeck and Biaoshuai Tao. 2019. Beyond worst-case (in) approximability of nonsubmodular influence maximization. *ACM Transactions on Computation Theory* 11, 3 (2019), 1–56.
- [24] Johan Ugander, Lars Backstrom, Cameron Marlow, and Jon Kleinberg. 2012. Structural diversity in social contagion. *Proceedings of the National Academy of Sciences* 109, 16 (2012), 5962–5966.