The Influence Maximisation Game

Demonstration Track

Sukankana Chakraborty
The Alan Turing Institute
London, United Kingdom
schakraborty@turing.ac.uk

Sebastian Stein
University of Southampton
Southampton, United Kingdom
ss2@ecs.soton.ac.uk

Ananthram Swami
DEVCOM U.S. Army Reserach Lab
Adelphi, MD, USA
ananthram.swami.civ@army.mil

Matthew Jones
University of Southampton
Southampton, United Kingdom
matthew.jones1@hotmail.co.uk

Lewis Hill
University of Southampton
Southampton, United Kingdom
lewis.j.hill.95@gmail.com

ABSTRACT

The problem of influence maximisation investigates efficient ways in which external influence (typically limited by resources) can be applied to a social network to maximise control over the global behaviours of a population. It is an effective tool that finds its application in many real-world scenarios, for instance it can be used to gather intelligence in crowdsourcing activities and to incentivise people to adopt desirable public policies. While the problem has been studied extensively in theoretical settings, many of these approaches can be expensive and inefficient to apply in the real world, particularly when considering an unknown or irrational competitor. The influence maximisation game was designed to bridge this gap between theory and the practical application of this knowledge. In this experiment, human subjects are presented with networks where they can employ their own tactics to maintain maximum influence against a competitor (which in this case is an AI agent). We aim to determine how people strategise to spread influence in the real world. In particular, we determine if people always act rationally in these settings or if their strategies are inherently biased —in which case we aim to identify inexpensive, yet effective strategies that can outperform these biased strategies. Observing how people strategise in the real world can help us modify our theoretical results for more efficient practical applications.

KEYWORDS

multi-agent systems; influence maximisation; network-based interventions

ACM Reference Format:
Sukankana Chakraborty, Sebastian Stein, Ananthram Swami, Matthew Jones, and Lewis Hill. 2023. The Influence Maximisation Game: Demonstration Track. 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

Social outcomes, such as election results and aggregate behaviours in populations (e.g. vaccination uptake) are often driven by the way influence propagates in social networks. The challenge for an external agent trying to control global behaviours in a population lies in determining the optimal distribution of limited resources over a network that can maximise the spread of influence in the population. This is a computationally expensive problem to solve [6] and although there are several analytical and computational methods that can approximate near-optimal solutions to the problem, in practice, it is significantly difficult to implement in the real world [11]. Moreover, influence maximisation efforts do not typically occur in isolation and are often met with competing influence (and resistance) in networks which further adds to the challenge of the problem. Existing work in this area typically determine optimal strategies while assuming that all actors involved are rational [2, 3], which in practice may not be true, as human beings often exhibit irrational behaviour and more importantly, they may not have the knowledge or tools to compute (and employ) optimal strategies in real-time. Thus, rendering theoretical solutions inefficient or expensive in real-world settings. With this in mind, we propose the influence maximisation game in which human players compete with an AI to maximise influence spread in synthetic social networks under resource constraints. While the study of collective behaviours and social contagion has received significant attention within the context of human-subject experiments [1, 4, 7], the influence maximisation problem has only been explored within the real-world context as field-experiments [8, 11], which come with their own set of challenges and often yield results specific to the setting being explored. Here with this effort we create an opportunity to extensively study the influence maximisation problem as it occurs in real world scenarios, but under controlled settings. This will help us determine the extent to which people strategise rationally within real world influence maximisation settings and if there are any biases in their strategies that can be exploited to further design effective and inexpensive counter-strategies. This is an initial effort to bridge the gap between theory and practice in this field and the results we obtain will have real-world applications, for instance they can be used to recruit members of the public in crowd-sourcing activities or to spread desirable health practices in a population [9, 10, 12].

2 EXPERIMENT DESIGN

In this experiment, human players participate in an influence maximisation task where they tactically employ limited budget to maximise influence in a synthetic social network against a competitor,

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

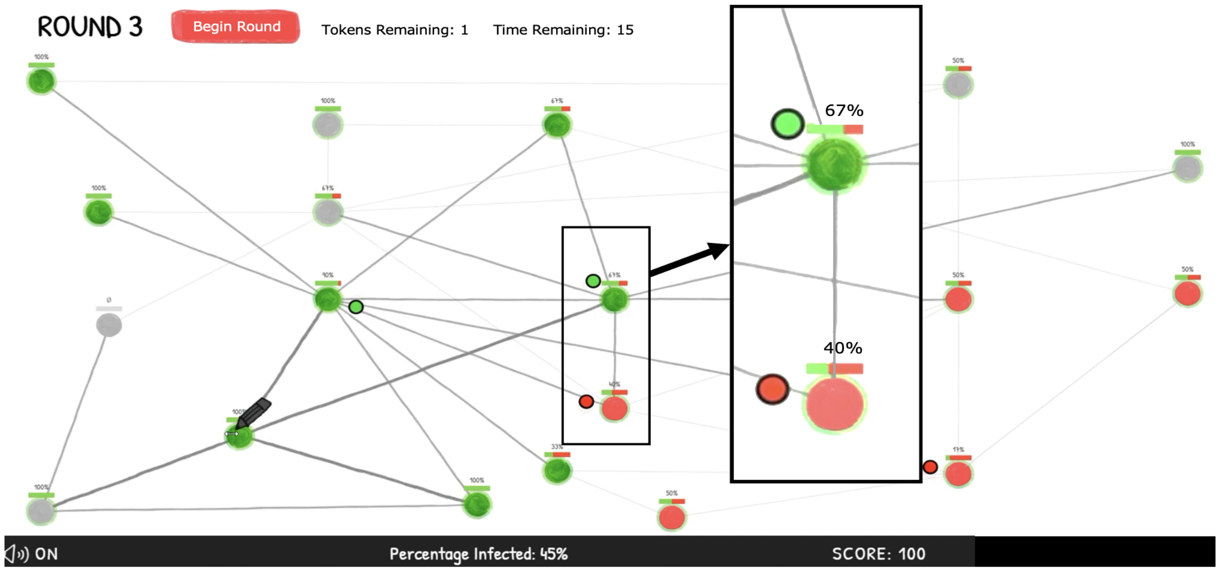

Figure 1: Figure showing a snapshot of the game. Influenced nodes are shown in green, nodes controlled by the adversary are shown in red and uninfluenced nodes are shown in grey. For more details and a tutorial of the game click this link.

which in this instance is an AI. Our ultimate aim is to study strategies adopted by players against both human and AI opponents, however as a first step we explore the problem with the AI opponent as using an AI opponent helps us control the conditions of the experiment with ease. In each game, a social network of $|V| = N = 20$ nodes and E edges is displayed to the players. The topologies of the networks are pre-determined and randomised across different instances of the game. The layout of the nodes are also randomised across games to remove any topological bias. Both the AI and the human player add one token per round to the network to maximise their spread of influence and once a token has been placed it cannot be moved in subsequent rounds. This helps us determine nodes that the players deem important at various stages of the game —and further contributes to the temporal analysis of the distribution of resources in the network.

Each game consists of ten rounds. At the start of the game, all nodes in the network are in an uninfluenced state (shown in grey in fig. 1). As the game progresses, nodes owned by the human player are shown in green and those controlled by the AI (or the opponent) are shown in red (see fig. 1). Note that, the overall strategy of the AI is unknown to the player. The action of the AI for any given round however, is revealed to the human player only after they have finalised their move. The incomplete knowledge of the opponent's strategy is reflective of real-world settings where competitors are not inclined to reveal their strategies in advance [5]. The AI is chosen at random at the start of the experiment and is fixed for each participant. The AI adopts any one of the following strategies:

(1) Degree-based strategy: The AI picks nodes based on their degrees. A degree-dependent probability function ($p \propto e^d$) determines the sequence in which nodes are to be targeted, where d is the degree of a node.
(2) Myopic-greedy strategy: The AI uses a greedy strategy to pick the node that contributes maximally to the increase in network share at the end of the (t+2)-*th* round. The look-ahead is *two* rounds for computational purposes.
(3) Random strategy: AI picks nodes uniformly at random.

Once the players commit their strategies, nodes in the network synchronously update their states based on the majority influence in their neighbourhood. For example, if a node has 1 red neighbour, 2 green neighbours and 1 green token, the fraction of green influence on them is $f = 0.75$, and thus the node adopts the green state with a probability $f^2/(f^2 + (1 - f)^2) = 0.9$. At the end of ten rounds, the player (human or AI) with the higher score wins the game. Performance indicators in the game help players tune and improve their strategies through the course of the experiment. The percentage infection (at the bottom of the screen) indicates the fraction of the network under the control of the human player (and it is updated at the end of each round). A cumulative score is also displayed to help players monitor their performance in the game. A graph of the cumulative scores (of both players) is displayed at the end of the game to aid players in reviewing their performance relative to the AI. To prevent players from running intricate calculations on the network (with the purpose of determining optimal strategies), we keep the network sufficiently complex and give players only 15 seconds to finalise their move. This further encourages players to use their perception of the network topology to make their moves. Finally, the design phase of this experiment is now complete and it is ready to be deployed on Amazon Mechanical Turk. Moves adopted by the players throughout the game will be recorded in a secure database that will later be used to analyse patterns in strategies that can be exploited to predict and outperform competitor strategies in real-time. In terms of future applications, an instance of this game can also be developed to be used as a training tool or a simulator.

ACKNOWLEDGMENTS

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement Number W911NF-16-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon. Part of this work was also funded via a UKRI Turing AI Fellowship (EP/V022067/1). The authors would also like to thank Dr. Markus Brede for the insightful discussions.

REFERENCES

[1] Joshua Becker, Devon Brackbill, and Damon Centola. 2017. Network dynamics of social influence in the wisdom of crowds. *Proceedings of the national academy of sciences* 114, 26 (2017), E5070–E5076.

[2] Shishir Bharathi, David Kempe, and Mahyar Salek. 2007. Competitive influence maximization in social networks. In *Internet and Network Economics: Third International Workshop, WINE 2007, San Diego, CA, USA, December 12-14, 2007. Proceedings 3*. Springer, 306–311.

[3] Tim Carnes, Chandrashekhar Nagarajan, Stefan M Wild, and Anke van Zuylen. 2007. Maximizing influence in a competitive social network: a follower's perspective. In *Proceedings of the ninth international conference on Electronic commerce.* 351–360.

[4] Damon Centola. 2010. The spread of behavior in an online social network experiment. *science* 329, 5996 (2010), 1194–1197.

[5] Sukankana Chakraborty, Sebastian Stein, Markus Brede, Ananthram Swami, Geeth de Mel, and Valerio Restocchi. 2019. Competitive influence maximisation using voting dynamics. In *Proceedings of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining.* 978–985.

[6] David Kempe, Jon Kleinberg, and Éva Tardos. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining.* 137–146.

[7] Milena Tsvetkova and Michael W Macy. 2014. The social contagion of generosity. *PloS one* 9, 2 (2014), e87275.

[8] Bryan Wilder, Laura Onasch-Vera, Juliana Hudson, Jose Luna, Nicole Wilson, Robin Petering, Darlene Woo, Milind Tambe, and Eric Rice. 2018. End-to-End Influence Maximization in the Field.. In *AAMAS*, Vol. 18. 1414–1422.

[9] Bryan Wilder, Sze-Chuan Suen, and Milind Tambe. 2018. Preventing infectious disease in dynamic populations under uncertainty. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.

[10] Amulya Yadav, Hau Chan, Albert Xin Jiang, Haifeng Xu, Eric Rice, and Milind Tambe. 2016. Using Social Networks to Aid Homeless Shelters: Dynamic Influence Maximization under Uncertainty.. In *AAMAS*, Vol. 16. 740–748.

[11] Amulya Yadav, Bryan Wilder, Eric Rice, Robin Petering, Jaih Craddock, Amanda Yoshioka-Maxwell, Mary Hemler, Laura Onasch-Vera, Milind Tambe, and Darlene Woo. 2017. Influence maximization in the field: The arduous journey from emerging to deployed application. In *Proceedings of the 16th conference on autonomous agents and multiagent systems.* 150–158.

[12] Amulya Yadav, Bryan Wilder, Eric Rice, Robin Petering, Jaih Craddock, Amanda Yoshioka-Maxwell, Mary Hemler, Laura Onasch-Vera, Milind Tambe, and Darlene Woo. 2018. Bridging the gap between theory and practice in influence maximization: Raising awareness about HIV among homeless youth.. In *IJCAI.* 5399–5403.