Continuous Influence Maximisation for the Voter Dynamics: Is Targeting High-Degree Nodes a Good Strategy?

Extended Abstract

Guillermo Romero Moreno
University of Southampton
Southampton, United Kingdom
grm1g17@soton.ac.uk

Long Tran-Thanh
University of Southampton
Southampton, United Kingdom
L.Tran-Thanh@soton.ac.uk

Markus Brede
University of Southampton
Southampton, United Kingdom
Markus.Brede@soton.ac.uk

ABSTRACT

In this paper, we relate influence maximisation (IM) for the voting dynamics to models of network control in which external controllers interact with the intrinsic dynamics of opinion spread. In contrast to previous literature, which has mostly explored the discrete setting, our focus is on continuous allocations of control. We develop an algorithm to numerically solve our IM problem via gradient ascent. We explore optimal allocations for leader-follower type networks for different budget scenarios and observe that optimal allocations do not systematically target hub nodes, as it has been found in previous literature. Conversely, strategies are strongly opponent-depend, avoiding nodes targeted by the opponent if the opponent has a larger budget, while shadowing the opponent's allocation otherwise, i.e. targeting the same nodes as them.

KEYWORDS

influence maximization; voter model; complex networks; external control

ACM Reference Format:

Guillermo Romero Moreno, Long Tran-Thanh, and Markus Brede. 2020. Continuous Influence Maximisation for the Voter Dynamics: Is Targeting High-Degree Nodes a Good Strategy?. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), Auckland, New Zealand, May 9–13, 2020,* IFAAMAS, 3 pages.

1 INTRODUCTION

Studies of influence maximisation (IM) cover a wide range of applications, such as marketing [7, 11], policy making [17, 18, 21], political campaigning [9, 19], or guidance of technological innovation [1]. In this field, societies are typically modelled as collectives of agents who hold opinions that may be represented by binary, discrete or continuous variables. Agents are considered connected by social networks and opinion change is modelled through peer-influence between connected neighbours (see [5] for a review).

Much of the study on IM has been performed on *static* models derived from the independent cascade model (IC), where propagation occurs as a one-off process [10, 11, 13]. In these contexts, the goal of an external influencer is to choose the subset of initial active nodes (*seeds*) for which the number of converted nodes is maximised at the end of the cascading process. Models like the IC are only suitable for capturing opinions that remain unchanged once individuals commit to them (e.g. adhering to a telecommunications supplier, or buying a car). However, in many settings, opinions can be fast-changing (e.g., political stances or fashion trends) and models like IC fail to efficiently address these situations. *Dynamic* models, on the contrary, consider stochastic flips of opinions in both directions and can capture better fast changes in volatile situations.

In this work, due to its prominence in the literature and conceptual simplicity which allows for analytical approaches, we focus on the voter model (VM) [6]. In the simple VM with finite, connected networks, a population asymptotically reaches a consensus. This situation changes with the inclusion of *zealots*, i.e. nodes that don't change opinion, leading the population to a fragmented steady-state with both opinions present [16]. Opinions at the steady-state depend on the number of zealots backing each opinion, the network topology, and the positions of the zealots in it [15].

Previous work has analysed IM for the voter model in terms of either finding optimal placements of zealots on a given social network [12, 20] —similar to seeding on the IC— or considering zealots as outside controllers with unidirectional influence links that can be distributed at will [2–4, 15]. Different from the seeding approach, agents targeted by the external controller are not assumed to be automatically "convinced" and control over them can only be obtained provided the controlling link is strong enough. The works above mainly found that optimal allocation strategies tend to target high-degrees nodes, with deviations to low-degree nodes only observed when a high level of noise is injected in the dynamics [2, 4] or when IM is sought before the dynamics reach the steady-state [3]. Importantly, all of the mentioned studies assume a discrete targeting of the controller: a node is either a zealot or not [12, 20], or a link from the controller to a normal agent is built or not [2–4, 15]. However, controllers may want to split their budgets unevenly between nodes to achieve a higher impact, e.g. weakly targeting one group of nodes while focusing strong controlling power on another subset of targeted nodes. This approach provides a more suitable model to describe mixed modes of influence campaigns, which typically combine generic messages directed to a broad public (through radio, television, billboards, etc.) with intense, bespoke promotion to specific groups of the population.

To bridge the gap of continuous influence allocations in the VM, we make the following contributions: (i) opinion control for the VM is generalised to the case of continuous allocations of control; (ii) exploiting the smoothness of the continuous setting we generalise gradient ascent to find optimal allocations; (iii) we show that optimal allocation strategies are strongly opponent-dependent, avoiding nodes targeted by the opponent when in resource disadvantage and chasing them when in resource superiority.

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand.

2 FORMALISATION OF THE MODEL

Opinion dynamics models typically assume a graph where agents are identified with nodes, $i = 1, \dots, N$, and edges are influence links between them. Agents can hold one of two possible opinion states (A or B), with opinion dynamics following the classic voter model. The controlling framework employs two external controllers that hold fixed opinions A and B, respectively, each aiming to spread their opinion in the network. Both controllers have unidirectional influencing links of non-negative strength $a_i \geq 0$ (or $b_i \geq 0$), to agents in the network. A budget constraint limits the sum of all link weights for each controller to a certain total, $a_{max} \geq \sum_i^N a_i$, $b_{max} \geq \sum_i^N b_i$.

A common approach to this type of problem is observing the average behaviour of the system by considering the evolution of the probability $x_i \in [0, 1]$ of nodes to adopt opinion A via rate equations [8, 16]. As Masuda [15] has shown, this system has exactly one equilibrium point (attractor), which can be found via the linear relationship

$$[L + \mathrm{diag}(a_i + b_i)]\, \boldsymbol{x} = \boldsymbol{a}\,, \tag{1}$$

where L is the weighted Laplacian of the network and bold symbols are $N \times 1$ vectors. The total vote share of nodes holding opinion A in the system is finally obtained from $X = \frac{1}{N} \sum_{i=1}^N x_i$.

In the IM problem, the A-controller aims to find her budget allocation, $\boldsymbol{a}^*$, that maximises her vote share. We set the B-controller to be static, i.e. with a fixed allocation strategy, so the A-controller attempts to solve the optimisation problem

$$\boldsymbol{a}^* = \arg\max_{\boldsymbol{a}} X(W, \boldsymbol{a}, \boldsymbol{b}), \quad \textstyle\sum_i^N a_i \leq a_{max}, \quad a_i \geq 0\,. \tag{2}$$

This maximisation can be performed by solving the system of equations derived from $\nabla_{\boldsymbol{a}} X = \boldsymbol{0}$, with

$$\nabla_{\boldsymbol{a}} X = \frac{1}{N}\left[\boldsymbol{1}^T \cdot [L + \mathrm{diag}(a_i + b_i)]^{-1} \mathrm{diag}(1 - x_i)\right]^T. \tag{3}$$

The system of equations $\nabla_{\boldsymbol{a}} X = \boldsymbol{0}$ is generally nonlinear in $\boldsymbol{a}$ and thus becomes analytically intractable. Since gradients can be computed analytically via (3), we generalise the gradient ascent (GA) approach of [14] to the VM.

Algorithm 1: GA approximation for IM

```
input : a_max, L, b, μ
repeat
    a = a + μ∇_a X;                          Gradient step
    a = a − 1/N(1^T a − a_max)1;             Proj. to constraint plane
    for a_j < 0 do
        a = a − a_j/(N−1) 1;  a_j = 0;       Projection to N-simplex edge
    end
    if X < prevX then
        a = prev a;  μ = μ/2;                Backtracking
    end
until a ≈ prev a;
```

3 RESULTS

Figure 1–a shows a sample run of the gradient ascent algorithm on a network with bimodal degree distribution. At the first iteration, the A-controller starts distributing the budget equally among all nodes. The difference in the resulting vote share between this initial strategy and the final strategy found by the GA accounts for 0.0154.

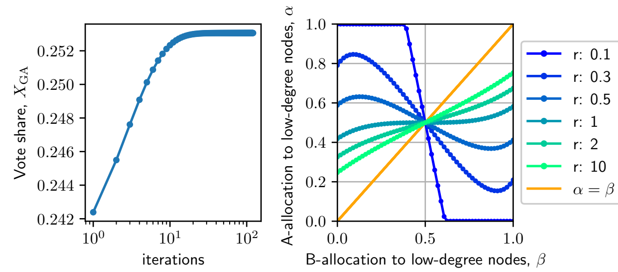

Figure 1: (a) Improvements in vote share throughout the iterations from a single sample run of the gradient ascent algorithm (log scale). (b) Optimal allocations for the A-controller for various strategies of the B-controller, β, and budget ratios, r on leader-follower type networks. Networks are of size $N = 1000$ with nodes of degrees 3 and 10 equally present. The budget of the A-controller is $a_{max} = N\langle k\rangle/3$.

We explore optimal allocation strategies for different scenarios in which the A-controller has larger or smaller budgets than the opponent. We define the ratio of budget availability of both controllers as $r = a_{max}/b_{max}$. Figure 1–b summarises the variations in optimal allocation strategies in networks with bimodal degree distribution (leader-follower type). For a particular budget ratio, r, we explore optimal A-allocations, α^*, for different B-allocations, β, where α and β represent the fraction of budget given to low-degree nodes in the network. In the figure, we observe that, when the A-controller is in large budget disadvantage ($r \ll 1$), she avoids direct confrontation with the B-controller, i.e. targets the nodes that have the least B-allocation. The avoiding behaviour damps down as the A-controller's disadvantage decreases (increasing r) and completely disappears when both budgets match ($r = 1$). In this scenario of equal budgets, the A-controller's optimal strategy is targeting nodes of low and high degree equally, largely unaffected by the strategy adopted by the B-controller. When the A-controller has an advantage of resources ($r > 1$), she devotes more resources to the group being targeted by the B-controller, but never as much as her.

We have analysed patterns of optimal influence allocation for varying budget availability in leader-follower networks. Apart from the absence of hub preferences in optimal allocations, we note that optimal strategies are strongly opponent-dependent. More specifically, the active controller always avoids the nodes targeted by the opponent when in budget disadvantage while shadowing nodes targeted by the opponent otherwise. Such a strong opponent dependency has barely been observed in previous literature, where the tendency to target high-degree nodes is prevalent.

ACKNOWLEDGMENTS

The authors acknowledge support from the Alan Turing Institute (grant R-SOU-006) and the Royal Society (grant IES\R2\192206).

REFERENCES

[1] Aamena Alshamsi, Flávio L Pinheiro, and Cesar A Hidalgo. 2018. Optimal diversification strategies in the networks of related products and of related research areas. *Nature Communications* 9, 1 (2018), 1328. https://doi.org/10.1038/s41467-018-03740-9

[2] Markus Brede, Valerio Restocchi, and Sebastian Stein. 2018. Resisting Influence: How the Strength of Predispositions to Resist Control Can Change Strategies for Optimal Opinion Control in the Voter Model. *Frontiers in Robotics and AI* 5 (2018), 34. https://doi.org/10.3389/frobt.2018.00034

[3] Markus Brede, Valerio Restocchi, and Sebastian Stein. 2019. Effects of time horizons on influence maximization in the voter dynamics. *Journal of Complex Networks* 7, 3 (jun 2019), 445–468. https://doi.org/10.1093/comnet/cny027

[4] Markus Brede, Valerio Restocchi, and Sebastian Stein. 2019. Transmission errors and influence maximization in the voter model. *Journal of Statistical Mechanics: Theory and Experiment* 2019, 3 (mar 2019), 033401. https://doi.org/10.1088/1742-5468/ab054b

[5] Claudio Castellano, Santo Fortunato, and Vittorio Loreto. 2009. Statistical physics of social dynamics. *Reviews of Modern Physics* 81, 2 (2009), 591–646. https://doi.org/10.1103/RevModPhys.81.591

[6] Peter Clifford and Aidan Sudbury. 1973. A model for spatial conflict. *Biometrika* 60, 3 (1973), 581–588. https://doi.org/10.1093/biomet/60.3.581

[7] Pedro Domingos and Matt Richardson. 2001. Mining the network value of customers. In *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data mining - KDD '01*. ACM Press, New York, New York, USA, 57–66. https://doi.org/10.1145/502512.502525

[8] Eyal Even-Dar and Asaf Shapira. 2011. A note on maximizing the spread of influence in social networks. *Inform. Process. Lett.* 111, 4 (jan 2011), 184–187. https://doi.org/10.1016/j.ipl.2010.11.015

[9] Serge Galam. 2016. Stubbornness as an unfortunate key to win a public debate: an illustration from sociophysics. *Mind & Society* 15, 1 (01 Jun 2016), 117–130. https://doi.org/10.1007/s11299-015-0175-y

[10] 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 (01 Aug 2001), 211–223. https://doi.org/10.1023/A:1011122126881

[11] 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 (KDD '03)*. 137–146. https://doi.org/10.1145/956750.956769

[12] Chris J. Kuhlman, V.S. Anil Kumar, and S.S. Ravi. 2013. Controlling opinion propagation in online networks. *Computer Networks* 57, 10 (jul 2013), 2121–2132. https://doi.org/10.1016/J.COMNET.2012.11.025

[13] Su-Chen Lin, Shou-De Lin, and Ming-Syan Chen. 2015. A Learning-based Framework to Handle Multi-round Multi-party Influence Maximization on Social Networks. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '15)*. 695–704. https://doi.org/10.1145/2783258.2783392

[14] Christopher W. Lynn and Daniel D. Lee. 2016. Maximizing Influence in an Ising Network: A Mean-field Optimal Solution. In *Proceedings of the 30th International Conference on Neural Information Processing Systems (NIPS'16)*. 2495–2503. http://dl.acm.org/citation.cfm?id=3157096.3157375

[15] Naoki Masuda. 2015. Opinion control in complex networks. *New Journal of Physics* 17 (2015), 1–11. https://doi.org/10.1088/1367-2630/17/3/033031

[16] M. Mobilia, A. Petersen, and S. Redner. 2007. On the role of zealotry in the voter model. *Journal of Statistical Mechanics: Theory and Experiment* 2007, 08 (aug 2007), P08029–P08029. https://doi.org/10.1088/1742-5468/2007/08/P08029

[17] Aida Rahmattalabi, Anamika Barman-Adhikari, Phebe Vayanos, Milind Tambe, Eric Rice, and Robin Baker. 2018. Influence Maximization for Social Network Based Substance Abuse Prevention. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, Sheila A. McIlraith and Kilian Q. Weinberger (Eds.). AAAI Press, 8139–8140. https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/17358

[18] Bryan Wilder, Sze-Chuan Suen, and Milind Tambe. 2018. Preventing Infectious Disease in Dynamic Populations Under Uncertainty. (2018). https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/16714

[19] Bryan Wilder and Yevgeniy Vorobeychik. 2018. Controlling Elections Through Social Influence. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '18)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 265–273.

[20] Ercan Yildiz, Asuman Ozdaglar, Daron Acemoglu, Amin Saberi, and Anna Scaglione. 2013. Binary Opinion Dynamics with Stubborn Agents. *ACM Transactions on Economics and Computation* 1, 4 (dec 2013), 1–30. https://doi.org/10.1145/2538508

[21] Haifeng Zhang, Yevgeniy Vorobeychik, Joshua Letchford, and Kiran Lakkaraju. 2016. Data-driven agent-based modeling, with application to rooftop solar adoption. *Autonomous Agents and Multi-Agent Systems* 30, 6 (nov 2016), 1023–1049. https://doi.org/10.1007/s10458-016-9326-8