Shielding and Shadowing: A Tale of Two Strategies for Opinion Control in the Voting Dynamics

Guillermo Romero Moreno*, Long Tran-Thanh, and Markus Brede

University of Southampton, School of Electronics and Computer Science, Southampton, UK,†
*grm1g17@soton.ac.uk

Abstract. This paper focuses on influence maximization or opinion control in the voting dynamics on social networks. We show two simple heuristics that are effective strategies to enhance vote shares: (i) avoiding the nodes controlled by your opponent when having a lower budget while focusing on them when having a larger budget (*shadowing*) and (ii) ring-fencing her influence by targeting control on adjacent nodes (*shielding*). The paper presents an empirical numerical evaluation of these strategies for various classes of complex networks which is backed up by analytical results obtained via a mean-field approach, in good agreement with numerical results. Importantly, we also show that optimal influence allocations tend to not be localized, but can include targeting nodes significant distances away from opposing influence.

Keywords: complex networks, voter dynamics, opinion control, influence maximization

1 Introduction

Issues of radicalization and polarization, or systematic application of influence on social media to disrupt elections are phenomena that have recently found much attention, both in the popular press [30, 4] and in the academic literature [35, 16, 6]. These problems can be approached by studying mathematical models of opinion dynamics on social networks (see, e.g., [10, 38] for reviews). An issue that is particularly relevant in this context is that of "opinion control", i.e. questions on how opinion dynamics can be strategically manipulated (or how such manipulation can best be prevented). This question is relevant in a wide number of contexts, ranging from political campaigning [21, 19, 17] and marketing [22] to studies of guiding technological innovation [1] and policy-making [41].

Formulations and solutions to the influence maximization problem go back to the work of [18, 22] and have been well-studied in the context of the independent cascade model, mostly in the computer science literature, but more

†This work was supported by the University of Southampton and by the Turing-sponsored pilot project *Strategic Influence in Dynamic Opinion Formation.*

recently also with elegant solutions via optimal percolation [33]. However, in the independent cascade model, opinion formation of agents is modeled as a one-off choice, making it less relevant for a number of applications of opinion formation without a strong commitment. The latter situation has found much attention in the socio-physics literature which has typically taken the approach of modeling opinions as stochastically changing due to peer-influence from network neighbors [10]. Opinions in such models have been described by either discrete or continuous variables. Prominent approaches in the first class, on which we focus in this paper, include the voter dynamics [20, 11], models directly inspired by Ising-like interactions [15], the Sznajd model [39] or majority dynamics [14]. Continuous models include approaches like the DeGroot model [13], considerations of bounded confidence introduced in [12] and, more recently, dynamics related to kinetic exchange models [25, 37]. Here, due to its mathematical tractability and prominence in the literature, we focus on the voter dynamics as a model of opinion change, which has also been proven to fit well empirical voting data [5].

Until recently, the problem of maximizing influence —or controlling opinions— has found less attention for the above class of stochastic models of opinion dynamics. Contributions to the literature mostly have focused on developing or evaluating numerical approaches, e.g. in the context of the political campaign problem for continuous opinion dynamics [19], the AB-model [2], the Glauber dynamics [26, 24, 28], and voting dynamics [23, 42, 29, 7–9]. Some very recent studies have also considered the co-evolution of control topologies and opinion dynamics at comparable time-scales, finding that active participation in influence maximization of many agents can enhance or hinder consensus formation in certain circumstances [6].

In the context of the voter model, opinion control has gained attention since the introduction of agents that change their opinions less frequently than other agents, so-called *zealots* [31, 32]. Zealots can have a substantial influence on the resulting opinion dynamics and some works have also studied the optimal placement of zealots on a social network to maximize their impact on opinions [23, 42]. Comparable to the question of finding optimal seeds in the independent cascade model, the question of finding optimal placings of zealots assumes that agents on a social network can be converted to become partisan, favoring a particular opinion. Perhaps a more interesting approach introduced by Masuda [29] takes inspiration from other approaches to network control [34, 27] and treats zealots as external, perfectly partisan agents who exert influence via strategically placed control links to the social network. To allow for mixed equilibrium states, Masuda [29] has introduced a framework in which the actions of an active (or optimizing) controller are evaluated against a passive controller who exerts her influence via random targeting. Influence maximization then translates into the problem of identifying an optimal set of targets, typically subject to a budget constraint. Addressing this problem, Masuda [29] has shown that for undirected networks targets typically follow a degree ordering, starting with nodes of the largest degrees. This is not necessarily the case for directed networks. Subsequent work has shown that noise or copying errors [9] and resistance to attempts of control [7] can shift optimal targeting to lower degree nodes in certain parameter regimes. Similar effects can also be observed if optimization is not aimed at achieving maximum vote shares in the stationary state, but rather at a finite time horizon [8]. Then, when time horizons are short, control cannot always capture hub nodes in time and targeting lower degree nodes may become optimal.

Here, we build on a numerical observation of a rule to achieve optimal control in [8] that allocates control toward nodes also influenced by the passive controller (what we term *shadowing*) or allocates control in a way to avoid nodes targeted by the passive controller. Additionally, we also introduce the notion of *shielding*, or surrounding nodes targeted by the passive controller. Below, we shall show that shadowing is generally effective when the optimizer has a larger budget than the passive agent, while avoidance is more effective otherwise. Additionally, both shadowing and avoidance strategies benefit from shielding as long as the subset of nodes targeted by the passive controller does not comprise a significant share of the network. Our results below are based on numerical experiments and backed by a mean-field approach that gives analytical support.

The paper is organized as follows. In Sect. 2, we formalize our approach to influence maximization in the voting dynamics. Section 3 first presents a summary of numerical experiments for some classes of complex networks and continues with the presentation of the mean-field approach and comparisons between numerical and analytical results. Finally, Sect. 4 gives a summary of our findings and draws conclusions.

2 Model and Methods

In the following, we model a social systems as a group of N agents who hold binary opinions $o_i \in \{A, B\}$, $i = 1, ..., N$. Agents are considered connected by a social network given by its weighted adjacency matrix w_{ij}. Below, we shall assume that the networks are undirected, contain no self-loops, and edges have positive weights. Agents update their opinions subject to the voter dynamics: at every iteration, a node is chosen at random and copies the opinion of a neighbor who is chosen with probability proportional to the weight of its in-link to the updating agent.

Following Masuda's framework [29], we assume the existence of one external controller per opinion. Controllers exert their influence via directed edges to the network, whose weights we label as $a_i, i = 1, \dots N$ for the A-controller and $b_i, i = 1, \dots N$ for the B-controller, both subject to budget constraints: $\sum_i a_i \leq a_{\max}$, $\sum_i b_i \leq b_{\max}$. Unlike [29] and other previous work on opinion control in the voter dynamics [23, 42, 9], we do not limit targeting to binary decisions (target a node with fixed weight or not target it at all) but allow for a continuous distribution of weight allocations.

To proceed, we introduce probabilities x_i that node i holds opinion $o_i = A$. The evolution of individuals' opinions can then be described through rate equations

$$\frac{\mathrm{d}x_i}{\mathrm{d}t} = (1 - x_i)\frac{\sum_{j=1}^{N} w_{ij} x_j + a_i}{k_i + a_i + b_i} - x_i \frac{\sum_{j=1}^{N} w_{ij} (1 - x_j) + b_i}{k_i + a_i + b_i} \;, \tag{1}$$

where $k_i = \sum_{j=1}^{N} w_{ij}$ is the weighted in-degree of node i.

This dynamical system converges to a single attractor $\mathbf{x}^*$ [29] that can be calculated via

$$\mathbf{x}^* = \left[L + \mathrm{diag}(\mathbf{a} + \mathbf{b})\right]^{-1} \mathbf{a} \;, \tag{2}$$

where bold symbols are vectors, L is the Laplacian matrix and $\mathrm{diag}(\mathbf{y})$ is a diagonal matrix with $\mathrm{diag}(\mathbf{y})_{ii} = y_i$. In equilibrium, the expected number of individuals with opinion A is then computed as

$$X^* = \frac{1}{N} \sum_{i=1}^{N} x_i^* = \frac{1}{N} \mathbf{1}^T \left[L + \mathrm{diag}(\mathbf{a} + \mathbf{b})\right]^{-1} \mathbf{a} \;. \tag{3}$$

In this paper, we study best-response strategies of an A-controller against a passive B-controller. We thus study the optimization problem for the A-controller, who aims to find a set of optimal targets and optimal control weights $\mathbf{a}$ to optimize its vote share

$$\mathbf{a}^* = \arg\max_{\mathbf{a}} X^* \;, \qquad \sum_{i=1}^{N} a_i \leq a_{\max} \;, \qquad a_i \geq 0 \;. \tag{4}$$

Equation (3) shows that X^* is non-linear in $\mathbf{a}$, so exact optimal solutions are very hard to obtain for non-trivial networks. Hence, we employ a numerical approach via gradient ascent [36] to find approximate solutions. Note that, as we are only focusing on equilibrium results, we will simplify notation below by dropping stars to indicate equilibrium.

3 Results

In Sect. 3.1, we first show numerical results for optimal strategies obtained via gradient ascent for various network topologies and analyze the presence of shielding and shadowing. In Sect. 3.2, we then develop an analytical understanding of shielding and shadowing via a mean-field approximation for random regular graphs.

3.1 Numerical Results

We first analyze optimal allocations obtained numerically via gradient ascent [36] on a variety of complex networks, namely 2D lattices, small-world networks [40], Barabasi-Albert (BA) networks [3] and random regular graphs. As an illustrative example, only one node is targeted by the B-controller. This node corresponds to a central position for the 2D grids, the highest degree node for BA networks and is randomly chosen for the other two network topologies. Scenarios with the passive controller targeting a bigger share of the social network are treated later on in Sect. 3.2.

By plotting the average strength of A-targeting as a function of distance from the B-targeted node, Fig. 1 gives best-response allocations of the A-controller. The A-controller is in budget disadvantage at the top row and budget advantage at the bottom row of the figure.

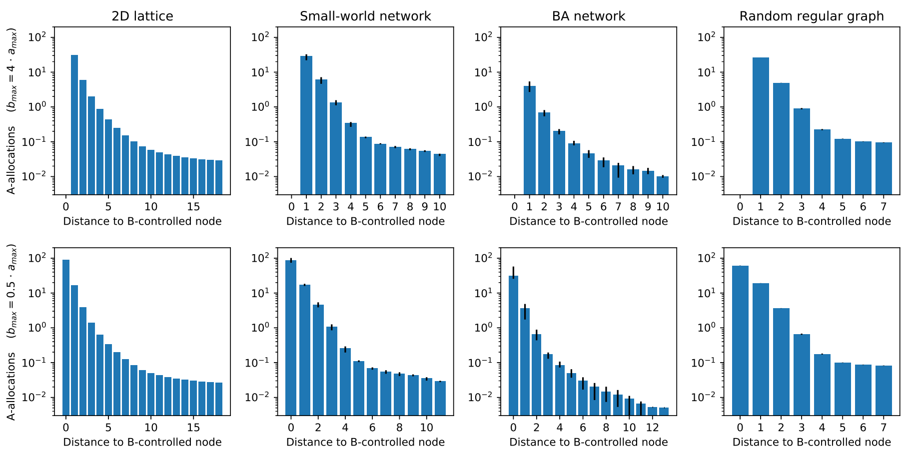

Fig. 1. Dependence of optimal influence allocation of the optimizing A-controller on the distance to the node targeted by the B-controller with budget three times larger (top) or half (bottom) than the A-controller's ($a_{\max} = 2N/3$). Bars give the median allocation for each group across 30 instances of the network class, with small bars referring to the first and third quartile. Networks are of size $N = 361$; from left to right: 2D regular lattices, small-world networks ($k=4$, $p=0.2$), Barabasi-Albert (BA) networks ($m = 1$) and random 4-regular graphs.

Inspecting the figure, we note that optimal strategies generally avoid the node targeted by the B-controller (see data points with distance zero in Fig. 1) when in resource disadvantage (top row) while focusing resources on the position of the B-controller when in resource advantage (bottom row). Furthermore, a clear pattern stands out in the allocations given to the remaining nodes. We observe a clear preference for control allocations to nodes directly surrounding the B-targeted node (distance one), with diminishing allocations the larger the distance. This shielding strategy is found independent of resource advantage or disadvantage for all shown network topologies. Moreover, we note that non-negligible control allocations tend to also be given to nodes very far away from the B-targeted nodes (note, in particular, the case of the lattices).

3.2 Analytical and Numerical Results on K-Regular Graphs

In this section, we aim to gain some analytical understanding of the effects of shadowing and shielding observed in numerical experiments above. For a basic model that can capture these effects, we neglect the effects of degree heterogeneity and focus on regular random graphs. To proceed, we consider a B-controller who targets a randomly selected fraction $0 < \rho_B < 1$ of the nodes in the network with uniform strength $b_B = b_{\max}/(N\rho_B)$, i.e. we set $b_i = b_B$, $i < \text{ceil}(\rho_B N), b_i = 0$ otherwise. Below, we develop a mean-field approximation for this situation and compare results based on this approximation to numerical results obtained via gradient ascent.

Neighbor Mean-Field Approximation for K-Regular Graphs Here, we aim to approximate the voting dynamics by grouping nodes into three disjoint classes. We shall assume that, irrespective of exact topological positions, all nodes in a class have identical states and nodes of different classes are connected at random. We define the first group B as all those nodes targeted by the B-controller: $B = \{i \mid b_i > 0\}$. The second group N is composed of nodes adjacent to nodes in B but not part of B: $N = \{i \notin B \mid (\exists j \in B : w_{ij} > 0\}$. The third group R is formed by all remaining nodes $R = \{i \notin (B \cup N)\}$, i.e. nodes which are not in direct contact to any B-targeted node. Note that this partition treats R as a uniform group; improvements of the method might be possible by distinguishing second and higher-order neighbors.

The probabilities of adopting A for the nodes of each group can be defined by x_B, x_N, and x_R, respectively, with their density in the network defined by ρ_B, ρ_N, and ρ_R, subject to $\rho_B + \rho_N + \rho_R = 1$. Whereas ρ_B is a given parameter, ρ_N and ρ_R must be derived from the network topology. For our assumption of random mixing and random targeting, this calculation is simple. Nodes belonging to R satisfy that 1) they are not targeted by the B-controller, which happens with probability $1 - \rho_B$ and 2) none of their K neighbors are targeted by the B-controller either, which happens with probability $(1 - \rho_B)^K$. Consequently, $\rho_R = (1 - \rho_B)^{K+1}$ and $\rho_N = 1 - \rho_B - \rho_R$.

To solve the influence maximization problem, the A-controller must decide how to optimally split her budget among the groups via determining a_B, a_N and a_R subject to the budget constraint $a_B\,\rho_B + a_N\,\rho_N + a_R\,\rho_R \leq a_{\max}/N$. Using (2), we obtain the mean-field vote shares for the groups:

$$
\begin{aligned}
x_B &= \frac{K\left(\gamma_{B|B}\,x_B + \gamma_{N|B}\,x_N\right) + a_B}{K + a_B + b_{\max}/\rho_B}\,, \\
x_N &= \frac{K\left(\gamma_{B|N}\,x_B + \gamma_{N|N}\,x_N + \gamma_{R|N}\,x_R\right) + a_N}{K + a_N}\,, \\
x_R &= \frac{K\left(\gamma_{N|R}\,x_N + \gamma_{R|R}\,x_R\right) + a_R}{K + a_R}\,,
\end{aligned}
\tag{5}
$$

where $\gamma_{Y|X}$ represents the probability for an edge to be attached to a node from group Y while the other edge is attached to a node from group X. The calculation of this probability is straight-forward for $\gamma_{Y|B}$:

$$\gamma_{B|B} = \gamma_B = \rho_B \ , \qquad \gamma_{N|B} = 1 - \rho_B \ . \tag{6}$$

At the other end, for $\gamma_{Y|R}$, the node at the other side of the edge will also belong to group R if none of its other $K-1$ edges is linked to a B-node:

$$\gamma_{R|R} = (1-\rho_B)^{K-1} \ , \qquad \gamma_{N|R} = 1-(1-\rho_B)^{K-1} \ . \tag{7}$$

Finally, for $\gamma_{Y|N}$, the use of Bayes' rule leads to

$$\begin{aligned}
\gamma_{B|N} &= \gamma_{N|B}\,\gamma_B/\gamma_N = (1-\rho_B)\ \rho_B/\rho_N \ , \\
\gamma_{R|N} &= \gamma_{N|R}\,\gamma_R/\gamma_N = \left[1-(1-\rho_B)^{K-1}\right]\ \rho_R/\rho_N \ , \\
\gamma_{N|N} &= 1-\gamma_{B|N}-\gamma_{R|N} \ .
\end{aligned} \tag{8}$$

Equation (5) can be made explicit via

$$\mathbf{x} = \left[\mathrm{diag}(K+\mathbf{a}+\mathbf{b}) - K\mathit{\Gamma}\ \right]^{-1}\mathbf{a}\ , \tag{9}$$

where bold symbols are vectors $\mathbf{y} = [y_B\ y_N\ y_R]^T$ and $\mathit{\Gamma}$ is the matrix containing the cross-probabilities between groups, $\mathit{\Gamma}_{ij} = \gamma_{j|i}$. Optimal allocations for the three groups can then be found by differentiating the estimated total vote share, $X^{\mathrm{mf}} = \boldsymbol{\rho}^T\mathbf{x}$, with respect to the allocation parameters, $\mathbf{a}$, and equating to zero: $\nabla_{\mathbf{a}}X^{\mathrm{mf}} = 0$, leading to a system of three non-linear, polynomial equations that can be solved numerically.

Testing the Neighbor Mean-Field Approximation To test the accuracy of the mean-field approximation, we compare predictions for stationary vote shares for the three groups of nodes to exact analytical solutions based on (2), for which we take the full network structure into account. To devise a set of test scenarios, we again assume that the B-controller distributes her resources equally among a given fraction ρ_B of nodes and the A-controller allocates control equally among all nodes in the network $a_i = a_{\max}/N,\ i = 1,\ldots N$.

Figure 2 shows the differences in vote share between exact (X) and approximate (X^{mf}) solutions for each of the three groups for varying fraction of B-targeted nodes (ρ_B) and B-controller's budget ($b_{\max}$). As expected, we generally observe a decline in vote shares with $b_{\max}$ and the qualitative dependencies are well captured by the mean-field approach. In more detail, the left panel of Fig. 2 shows the case of $\rho_B = 0.01$, implying that roughly a fraction $\rho_N = 0.03$ of nodes are adjacent to them. We note that the mean vote share of B-nodes and R-nodes is accurately estimated by the mean-field approximation ($x_B^{\mathrm{mf}} \sim x_B, x_R^{\mathrm{mf}} \sim x_R$), while the vote share of the N-nodes is slightly overestimated over the whole range of budgets ($x_N^{\mathrm{mf}} > x_N$). In the middle panel, we set $\rho_B = 0.1$, implying $\rho_N = 0.24$ and $\rho_R = 0.66$. For this scenario, mean-field results for all groups are in very good agreement with numerical results. Last, as shown in the right hand panel of Fig. 2, we also consider $\rho_B = 0.3$, $\rho_N = 0.46$, $\rho_R = 0.24$. Mean-field estimates are again in good agreement, but some slight systematic underestimation of x_R is observed.

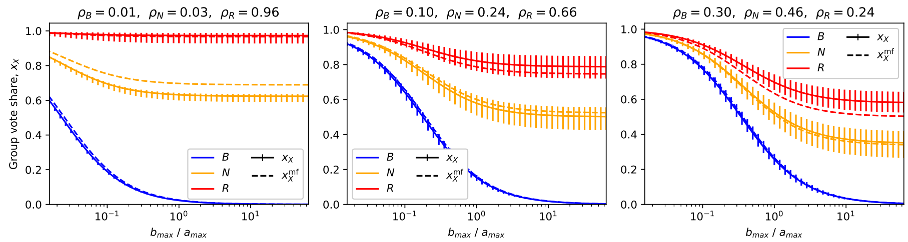

Fig. 2. Average probability to adopt A by the nodes of the three mean-approximation groups (B,N,R) for varying fraction of B-nodes, ρ_B, and B-controller's budget, $b_{\max}$. The budget for the A-controller is $a_{\max} = N/2$. Points from the exact solution, x_X, are the mean vote share in each group averaged over 15 randomly generated networks of size $N = 1000$, with vertical lines corresponding to the upper and lower mean absolute deviations of vote shares within the group. Results of the approximation are given by x_X^{mf} (*slashed lines*). The titles show the derived fractions of neighbors of each group present in the network, ρ_X.

In general, mean-field estimates for B-nodes prove highly accurate. This is expected, as B-nodes are heavily influenced by the B-controller and less susceptible to network effects. The estimation errors for the other two groups are caused by two different limitations of the approximations. The first limitation results from treating all nodes in R as a uniform group, even though they have different distances to nodes in N and, consequently, different vote shares. This limitation is most prominent for low densities of ρ_B and we see its effects in Fig. 2–left. Here, we observe an over-estimation of x_N, which is caused by an over-estimation of vote shares of second neighbors of B-controlled nodes. The second limitation of our mean-field approach is not subdividing N-nodes into separate classes depending on exact numbers of adjacent B- or R-nodes. However, as B-controls are allocated randomly, some N-nodes will have a larger number of B-neighbors than others. N-nodes with more B-neighbors than average will have smaller vote shares, but a lack of remaining connections to R-nodes implies they have a reduced impact on R. In contrast, N-nodes with fewer B-neighbors than average have higher than average vote shares and large impact on R. Hence, the mean-field approach underestimates vote shares of R-nodes. This situation is particularly prominent when ρ_B is relatively large, and we see its effects in Fig. 2–right.

Numerical and Analytical Evaluations of the Heuristics In this section, we compare optimal control allocations obtained numerically via gradient ascent and optima calculated from mean-field estimates via solving $\nabla_{\mathbf{a}} X^{\mathrm{mf}} = 0$, which allows us to quantify the roles of shielding and shadowing. As in the previous subsection, experiments are again run on K-regular networks with a varying fraction of B-targeted nodes (ρ_B) and varying B-controller's budget ($b_{\max}$).

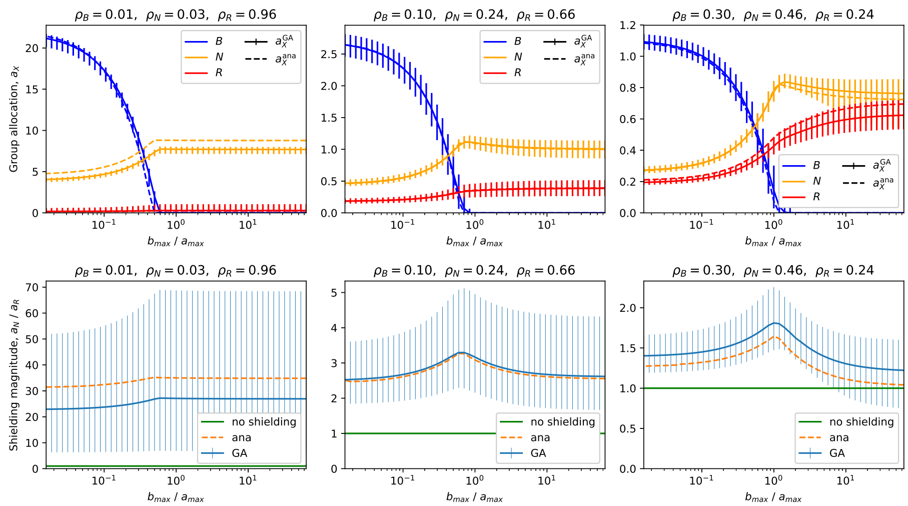

Fig. 3. Dependence of optimal control allocation for the three groups (top) and shielding strength (bottom) on B-controller's budgets ($b_{\max}$) for varying fractions of B-targeted nodes (ρ_B). We set $a_{\max} = N/2$ and use $K = 3$. Analytical solutions correspond to a_X^{ana} (*slashed lines*) and numerical solutions a_X^{GA} (*solid lines*) are averaged over 15 randomly generated networks of size $N = 1000$, with vertical lines corresponding to the upper and lower mean absolute deviations of allocations within the group. Titles of panels also show the derived fractions of N-nodes (ρ_N) and R-nodes (ρ_R).

Figure 3 presents the optimal allocations given by both numerical and analytical solutions (top row) and the intensity of shielding in the allocations (bottom row). In the top row, we note a particular pattern in optimal allocations to B-nodes: allocations are roughly inversely proportional to the budget of the B-controller and vanish at a point near budget equality between A- and B-controller. This finding confirms the numerical results related to shadowing for general classes of complex networks presented in Sect. 3.1. Considering N-nodes (*orange lines*), we note that generally $a_N > a_R$, i.e. we establish the presence of a shielding effect for all parameter settings. Shielding is even present when the A-controller has a much larger budget than the B-controller and peaks at transition points where shadowing shifts to full avoidance.

In the scenario of $\rho_B = 0.01$ (left), analytical solutions generally assign more budget to N-nodes than numerical solutions. This difference is likely caused by ignoring the role of second neighbors in the mean-field approximation, while these nodes can be used to produce a second barrier of shielding that protects further neighbors (as was seen in Fig. 1). The optimal allocations given by both methods match very well for $\rho_B = 0.1$ (middle). For $\rho_B = 0.3$ (right), the analytical allocations underestimate the benefits of shielding. This can be an effect of the mean-field approximation not differentiating between the number of

R-neighbors an N-node has; only a subgroup of N-nodes is in direct contact with R-neighbors and focusing on them makes shielding more effective. For high values of the B-controller's budget, there are N-nodes in the numerical solutions that receive even less allocation than R-nodes (panel bottom–right). These are those N-nodes that have B-nodes as their three neighbors (not shown), so resources allocated to them do not spread their impact to other nodes. The intensity of shielding is higher the lower the ρ_B (panel bottom-left).

4 Conclusions

In this paper, we have analyzed two complementary heuristics for continuous influence maximization in the voter model. These heuristics are shown to be best-response strategies of an active external controller to counter the influence of a passive opponent. The first heuristic, shielding, focuses influence on the neighbors of nodes targeted by the opponent to create a barrier that limits the spread of her opinion. Smaller barriers on further neighbors reinforce the effectiveness of the strategy. The second heuristic, shadowing, focuses on nodes directly targeted by the opponent if in budget advantage while avoiding them if in budget disadvantage. The point at which to transition from avoidance to shadowing depends on the budget of both controllers, the social network topology, and the strategy followed by the passive controller. Such best-response behavior has also been observed in previous literature on other network topologies and real networks [36], as well as for discrete allocations in the voter model [8], but was not systematically explored.

Optimal best-response behaviors obtained via numerical optimization show a consistent presence of the above heuristics for various classes of complex networks. The transition point to a shadowing behavior depends on the topological properties of the nodes targeted by the opponent. The effectiveness of the heuristics is further corroborated via analytical results on random regular graphs obtained through a mean-field approximation, which are in good agreement with the numerical results. The comparison of both approaches shows a more nuanced picture, revealing different roles played by the nodes depending on the number of adjacent opponent-controlled neighbors or their distance to the B-targeted nodes. The effectiveness of shielding decreases the more nodes the opponent targets and peaks around budget equality of the controllers. Further work should focus on refining the mean-field approach by including more distinct groups and thus overcoming the aforementioned limitations.

References

1. Alshamsi, A., Pinheiro, F.L., Hidalgo, C.A.: Optimal diversification strategies in the networks of related products and of related research areas. Nature Communications 9(1), 1328 (2018), https://doi.org/10.1038/s41467-018-03740-9
2. Arendt, D.L., Blaha, L.M.: Opinions, influence, and zealotry: a computational study on stubbornness. Comput. Math. Organ. Theory 21, 184 (2015)

3. Barabási, A.L., Albert, R.: Emergence of Scaling in Random Networks. Science 286(5439), 509–512 (oct 1999), http://www.sciencemag.org/lookup/doi/10.1126/science.286.5439.509
4. Beckett, L.: Trump digital director says facebook helped win the white house (2017), https://www.theguardian.com/technology/2017/oct/08/trump-digital-director-brad-parscale-facebook-advertising
5. Braha, D., de Aguiar, M.A.M.: Voting contagion: Modeling and analysis of a century of U.S. presidential elections. PLOS ONE 12(5), e0177970 (may 2017)
6. Brede, M.: How does active participation effect consensus: Adaptive network model of opinion dynamics and influence maximizing rewiring. Complexity 2019, 1486909 (2019)
7. Brede, M., Restocchi, V., Stein, S.: 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, 34 (2018)
8. Brede, M., Restocchi, V., Stein, S.: Effects of time horizons on influence maximization in the voter dynamics. Journal of Complex Networks 7(3), 445–468 (jun 2019), https://academic.oup.com/comnet/article/7/3/445/5149693
9. Brede, M., Restocchi, V., Stein, S.: Transmission errors and influence maximization in the voter model. Journal of Statistical Mechanics: Theory and Experiment 2019(3), 033401 (mar 2019)
10. Castellano, C., Fortunato, S., Loreto, V.: Statistical physics of social dynamics. Reviews of Modern Physics 81(2), 591–646 (2009)
11. Clifford, P., Sudbury, A.: A model for spatial conflict. Biometrika 60(3), 581–588 (1973)
12. Deffuant, G., Neau, D., Amblard, F., Weisbuch, G.: Mixing beliefs among interacting agents. Advances in Complex Systems 3, 87–89 (2000)
13. DeGroot, M.: Reaching a consensus. Journal of the American Statistical Association 69, 118–121 (1974)
14. Galam, S.: Minority opinion spreading in random geometry. Eur. Phys. J. B 25, 403 (2002)
15. Galam, S., Gefen, Y., Shapir, Y.: Sociophysics: A new approach of sociological collective behaviour. i. meanbehaviour description of a strike. J. Math. Sociol. 9, 1–13 (1982)
16. Galam, S., Javarone, M.A.: Modelling radicalization phenomena in heterogeneous populations. PLoS ONE 11, e0155407 (2016)
17. Galam, S.: Stubbornness as an unfortunate key to win a public debate: an illustration from sociophysics. Mind & Society 15(1), 117–130 (Jun 2016)
18. Goldenberg, J., Libai, B., Muller, E.: Talk of the network: A complex systems look at the underlying process of word-of-mouth. Marketing Letters 12(3), 211–223 (Aug 2001)
19. Hegselmann, R., König, S., Kurz, S., Niemann, C., Rambau, J.: Optimal opinion control: The campaign problem. Journal of Artificial Societies and Social Simulation 18 (2015)
20. Holley, R.A., Liggett, T.M.: Ergodic theorems for weakly interacting infinite systems and the voter model. The Annals of Probability 3(4), 643–663 (1975)
21. Javarone, M.A.: Network strategies in election campaigns. J. Stat. Mech. p. P08013 (2014)
22. Kempe, D., Kleinberg, J., Tardos, E.: Maximizing the spread of influence through a social network. In: Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 137–146. KDD '03, ACM, New York, NY, USA (2003), http://doi.acm.org/10.1145/956750.956769

23. Kuhlman, C.J., Kumar, V.A., Ravi, S.: Controlling opinion propagation in online networks. Computer Networks 57(10), 2121–2132 (jul 2013)
24. Laciana, C.E., Rovere, S.L.: Ising-like agent-based technology diffusion model: adoption patterns versus seeding strategies. Physica A 390, 1139 (2011)
25. Lallouache, M., Chakrabarti, A.S., Chakraborti, A., Chakrabarti, B.K.: Opinion formation in kinetic exchange models: Spontaneous symmetry-breaking transition. Physical Review E 82, 056112 (2010)
26. Liu, S., Shakkottai, S.: Influence maximization in social networks: an ising-model-based approach. Proc. of the 48th Annual Allerton Conf. p. 570 (2010)
27. Liu, Y.Y., Slotine, J.J., Barabási, A.L.: Controllability of complex networks. Nature 473(7346), 167–173 (may 2011)
28. Lynn, C.W., Lee, D.D.: Maximizing influence in an ising network: A mean-field optimal solution. In: Proceedings of the 30th International Conference on Neural Information Processing Systems. pp. 2495–2503. NIPS'16 (2016)
29. Masuda, N.: Opinion control in complex networks. New Journal of Physics 17, 1–11 (2015)
30. McFaul, M., Kass, B.: Understanding Putins Intentions and Actions in the 2016 U.S. Presidential Election. Tech. rep., Standford University (Jun 2019)
31. Mobilia, M.: Does a single zealot affect an infinite group of voters? Phys. Rev. Lett. 91, 028701 (2003)
32. Mobilia, M., Petersen, A., Redner, S.: On the role of zealotry in the voter model. Journal of Statistical Mechanics: Theory and Experiment 2007(08), P08029–P08029 (aug 2007)
33. Morone, F., Makse, H.A.: Influence maximization in complex networks through optimal percolation. Nature 524, 65 (2015)
34. Porfiri, M., di Bernardo, M.: Criteria for global pinning-controllability of complex networks. Automatica 44(12), 3100–3106 (dec 2008)
35. Ramos, M., Shao, J., Reis, S.D.S., Anteneodo, C., Andrade, J.S., Havlin, S., Makse, H.A.: How does public opinion become extreme? Sci. Rep. 5, 10032 (2015)
36. Romero Moreno, G., Tran-Thanh, L., Brede, M.: Continuous influence maximisation for the voter dynamics: Analytical solution for leader-follower networks and gradient ascent algorithm. Manuscript submitted for publication (2019)
37. Sen, P.: Phase transitions in a two-parameter model of opinion dynamics with random kinetic exchanges. Physical Review E 83, 016108 (2011)
38. Sîrbu, A., Loreto, V., Servedio, V.D.P., Tria, F.: Opinion Dynamics: Models, Extensions and External Effects, pp. 363–401. Springer International Publishing, Cham (2017), https://doi.org/10.1007/978-3-319-25658-0_17
39. Sznajd-Weron, K., Sznajd, J.: Opinion evolution in closed community. Int. J. Mod. Phys. C 11, 1157–1165 (2000)
40. Watts, D.J., Strogatz, S.H.: Collective dynamics of small-world' networks. Nature 393(6684), 440–442 (jun 1998), http://www.nature.com/articles/30918
41. Yadav, A., Wilder, B., Rice, E., Petering, R., Craddock, J., Yoshioka-Maxwell, A., Hemler, M., Onasch-Vera, L., Tambe, M., Woo, D.: Influence maximization in the field: The arduous journey from emerging to deployed application. In: Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS (2017)
42. Yildiz, E., Ozdaglar, A., Acemoglu, D., Saberi, A., Scaglione, A.: Binary Opinion Dynamics with Stubborn Agents. ACM Transactions on Economics and Computation 1(4), 1–30 (dec 2013)