Competitive influence maximisation with nonlinear cost of allocations*

Sukankana Chakraborty[0000−0002−2115−8531] and Sebastian Stein[0000−0003−2858−8857]

University of Southampton, Highfield Campus, SO17 1BJ, UK.
{sc8n15,ss2}@ecs.soton.ac.uk

Abstract. We explore the competitive influence maximisation problem in the voter model. We extend past work by modelling real-world settings where the strength of influence changes nonlinearly with external allocations to the network. We use this approach to identify two distinct regimes — one where optimal intervention strategies offer significant gain in outcomes, and the other where they yield no gains. The two regimes also vary in their sensitivity to budget availability, and we find that in some cases, even a tenfold increase in the budget only marginally improves the outcome of an intervention in a population.

Keywords: influence maximisation · nonlinear model · voter dynamics.

1 Introduction

Our interactions with our peers often impact our personal choices, behaviours and opinions. Social networks have therefore attracted considerable attention as a medium to control collective behaviours in populations through external interventions [17]. A key challenge in this aspect is to determine the optimal allocation of external influence on the network that can maximise the outcome of an intervention in a population. The influence maximisation approach addresses this problem by exploiting interpersonal ties in a social network, to maximise the adoption of an innovation or a behaviour in a population [11]. It is typically framed as an optimisation problem that identifies the most influential individuals in a network who can maximise the spread of a desired behaviour in the

* 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. The authors would like to thank Dr. Markus Brede for the insightful discussions. The authors acknowledge the use of the IRIDIS High Performance Computing Facility at the University of Southampton for the completion of this work.

rest of the population. In the past, influence maximisation has been categorically studied using diffusion models [11]. While these models aptly describe how decisions and behaviours virally spread in social systems, their representation of individual states reflect long-term commitments such as buying a car. The one-off, immutable nature of individual states in these models make them unsuitable for studying settings where individual choices (or opinions) are transient and free of abiding commitments [13]. In contrast, dynamical models are used to study instances where individuals frequently change their states as they interact with their social neighbourhood, and thus is a more appropriate model when studying behaviour and opinion dynamics in social networks[1].

In this paper, we consider a dynamical model known as the voter model to capture influence flow in a population [5,9]. This paradigmatic model is characterised by its simple but effective approach to study reality-based social dynamics [16]. The voter model has attracted considerable attention within influence maximisation research [12,23]. Most of this work however, mimics the traditional setting where a limited budget is used to convert ("seed" or activate) a small number of individuals who subsequently influence the rest of the population. Such an approach focuses on identifying the most influential individuals in the network and abstracts all other information about how the intervention budget (e.g. marketing budget) should be used. In the real world, influence maximisation efforts are typically led with resources such as time and money, and a strategy detailing the optimal distribution of these resources would be a useful result. With this in mind, we assume continuous allocations of resources (e.g. time and money) on the network, where individuals are targeted with varying intensities based on their importance in the influence spread process. We focus on studying the influence maximisation problem in competitive settings, as dynamics in the voter model either converge to an ordered consensus, or reach a fragmented state at equilibrium [2], and consensus is rarely ever achieved in the real world.

It is important to note that when considering traditional methods, the intervention budget constrains the number of nodes "seeded" at the start of the dynamics. In the continuous approach however, resources are spread heterogeneously over the network and thus an explicit relationship between the amount of resources allocated to a node and the strength of influence experienced by them needs to be defined. Past work in this area has strictly assumed this relationship to be linear [4,18], i.e. the amount of allocated resources is directly proportional to the strength of influence experienced by the node. However, this assumption may not consistently apply to all real-world settings. For instance, some interventions take longer to be understood or adopted (e.g. adoption of green technologies), and hence require more resources [6]. Similarly, studies from marketing research show that more resources is not always better, and that the duration and complexity of advertisements often have a *diminishing returns* effect on customer engagement and interest [7]. With this in mind, here we present a novel model that considers nonlinear relationships between the strength of influence and external allocations. So far nonlinearity has only been considered in

[1] For a comprehensive review see [2].

the spread dynamics by introducing noise such as in the q-voter model [3] or by adding contrarians to the network [20], and to the best of our knowledge, it has never been studied in terms of allocations to the network.

2 Model

We consider a population of N individuals, connected via a social network. The structure of the network is given by a graph $G(V, E)$, where vertices V represent individuals and edges E the relationships between them. Any vertex $i \in V = \{1, 2, \ldots, N\}$ is connected to other individuals in the network $\{j \in V; j \neq i\}$ through a subset of E. Edges between individuals are indicated using weights w_{ij}. Here we consider unweighted and undirected graphs, where binary weights are used to capture the structure of the network, such that $w_{ij} = 1$, if i and j have an edge between them, else $w_{ij} = 0$. Additionally, W here is symmetric as we consider undirected graphs.

We explore a setting where two controllers (A and B) compete to maximise their influence (or opinions) in the population. At any given point in time, individuals in the network strictly adhere to one of two opinions (A or B), corresponding to each controller. Opinions are characterised using binary state variables $\sigma_{A,i}(t) \in \{0, 1\}$, where $\sigma_{A,i}(t) = 1$ implies node i is in state A, or in state B ($\sigma_{A,i}(t) = 0$) at time t. From here, it follows that $\sigma_{B,i}(t) = 1 - \sigma_{A,i}(t)$. Controllers maximise their opinion shares in the population by influencing the network externally. Here we assume a nonlinear relation between allocations and the strength of influence experienced by the node. For any node i, external influence from controllers A and B are $p_{A,i}$ and $p_{B,i}$, when $p^{\gamma}_{A,i}$ and $p^{\gamma}_{B,i}$ amount of resources are allocated to it. Influence over the entire network is described using non-negative vectors, $p_A \in R^N_+$ and $p_B \in R^N_+$ which are nonlinearly constrained by controller budgets B_A and B_B, as $\sum_i p^{\gamma}_{A,i} = B_A$ and $\sum_i p^{\gamma}_{B,i} = B_B$.

Nodes update their opinions using voter dynamics [9], at every time step, a node is selected uniformly at random to update their opinion state where they copy the state of a neighbouring node j with the probability $w_{ji}/(\sum_{j\in\mathcal{N}_i} w_{ji} + p_{A,i} + p_{B,i})$ where $\mathcal{N}_i$ is the immediate social neighbourhood of i, or copy the state of an external controller (say A) with the probability $p_{A,i}/(\sum_{j\in\mathcal{N}} w_{ji} + p_{A,i} + p_{B,i})$. As opinions are stochastic, we approximate the global behaviour in the system by assuming that $x_{A,i}$ is the probability a node i is in state $\sigma_{A,i} = 1$, which gives us the rate at which it chooses to remain in opinion state A as,

$$\frac{dx_{A,i}}{dt} = (1 - x_{A,i})\frac{\sum_j w_{ji}x_{A,j} + p_{A,i}}{\sum_j w_{ji} + p_{A,i} + p_{B,i}} - x_{A,i}\frac{\sum_j w_{ji}(1 - x_{A,j}) + p_{B,i}}{\sum_j w_{ji} + p_{A,i} + p_{B,i}}. \quad (1)$$

Here the terms $\frac{\sum_j w_{ji}x_{A,j}+p_{A,i}}{\sum_j w_{ji}+p_{A,i}+p_{B,i}}$ and $\frac{\sum_j w_{ji}(1-x_{A,j})+p_{B,i}}{\sum_j w_{ji}+p_{A,i}+p_{B,i}}$ quantify the total influence a node i experiences from their immediate neighbourhood and from external controllers in favour of opinions A and B respectively. We estimate the global behaviour of the population by estimating the total share of opinions obtained by each controller at equilibrium. We determine steady-state conditions by setting $\frac{du_{A,i}}{dt} = 0$ in Eq. (1), which for an arbitrary network of size N yields $[L + diag(p_A + p_B)]x_A = p_A$, where L is the Laplacian of the network given by a $N \times N$ matrix with diagonal elements representing the total strength of all edges on a node ($L_{ii} = \sum_j w_{ji}$) and off-diagonal elements are $L_{ij} = -w_{ij}$. The total vote-share obtained by controller A at equilibrium is then given by $X_A = \frac{1}{N}\vec{1}^T x_A = \frac{1}{N}\vec{1}^T[L + diag(p_A + p_B)]^{-1} p_A$, which can be used to present the formal optimisation problem as,

$$p_A^* = \underset{p_A \in \mathcal{P}}{\arg\max}\, X_A^*(L, p_B), \tag{2}$$

where $\mathcal{P}$ is a set of all possible allocations p_A such that $\sum_{i=1}^{N} p_{A,i}^{\gamma} = B_A$. [2]

3 Methods

We now propose approaches to solve the optimisation problem given in Eq. (2). Depending on the value of γ, the constraint space changes yielding the following *two* settings: (i) when $\gamma > 1$, we observe *diminishing returns* in terms of the strength of influence experienced by nodes as allocations are increased, and (ii) when $0 < \gamma < 1$, we find that individuals take longer (or more resources) to respond to external influence and hence illustrates a *delayed influence* effect. Below we present numerical approaches for each case.

3.1 Numerical methods

Case $\gamma > 1$: We first consider the *diminishing returns* instance where $\gamma > 1$. The strength of influence experienced by a node is given by the γ-th root of allocated resources $p_{A,i}^{\gamma}$ and thus the norm of the allocation vector is constrained by $B_A^{1/\gamma}$, implying that the constraint set here is convex in nature. Given that the vote-share X_A is a concave function of allocations p_A [18], the problem we have at hand is therefore a convex optimisation problem. A common approach taken to solve this type of constrained optimisation problem uses the Lagrange method where the objective function is maximised by iteratively stepping in the optimal direction, within the constraint space. To employ this method in our work we follow the approach discussed in [21] which yields the optimal direction as,

$$p_{A,i}(t+1) \leftarrow p_{A,i}(t) + \eta \frac{\nabla_{p_A(t)} X_{A,i}(t)}{p_{A,i}^{\gamma-2}(t)}, \tag{3}$$

[2] The total vote-share is a function of the network structure L, competitor allocations p_B and controller allocations p_A. The maximum vote-share X_A^* is achieved by the optimal allocation vector p_A^*.

where η is the step-length and $\nabla_{p_A(t)} X_{A,i}(t)$ is the gradient of the vote-share function wrt to allocations[3]. To solve the optimisation problem we first initialise a random, feasible allocation vector which then uses Eq. (3) to iteratively update allocations until the total vote-share can no longer be improved, or alternatively we obtain a $\mu-$approximation of the optimal allocation configuration. Note that the allocation vector is also normalised by scaling the entries in p_A at every time step to satisfy the budget constraint $\sum_i p_{A,i}^{\gamma} = B_A$. The step-length is adjusted using back-tracking[4] to ensure convergence, and given that the problem at hand is a convex optimisation problem, the algorithm is guaranteed to converge to the global maximum.

Case $0 < \gamma < 1$: Next, we consider the *delayed influence* setting where $0 < \gamma < 1$. The constraint set in this case is clearly nonconvex, and hence the Lagrange method cannot be employed here. In general, nonconvex optimisation problems are difficult to solve. In some cases however, the structure of the objective function and the constraint function can be exploited to design polynomial-time algorithms that yield near-optimal solutions [10]. Given that our objective function here is similar to the one shown in [18] and the constraint space is in the shape of an ℓ_p-norm ball where $0 < p < 1$, we employ the projected gradient ascent algorithm proposed in [18], and modify the projection step to meet the nonlinear budget constraint considered in our work. For the projection method, we use the IRBP algorithm[5] which is an instance of a majorisation-minimisation algorithm, where the algorithm iteratively alternates between a majorisation step and a minimisation step[6] until it converges [22]. As this is a nonconvex problem, we do not have any theoretical guarantees of reaching the optimal solution, and therefore we run the algorithm for multiple initialisations of the allocation vector and consider the mean result obtained over all simulations.

4 Results

We now use the above approaches to study the problem in synthetic and real world networks. We also present analytical approaches in synthetic networks to provide benchmarks for our numerical results.

4.1 Analytical benchmark

First we propose an analytical method to determine optimal allocations. Analytical methods typically apply to simplified network structures, and here we

[3] Given by $\nabla_{p_A} X_A = \frac{1}{N}\vec{1}^T[L + diag(p_A + p_B)]^{-1}(I - diag(x_A))$.

[4] When $X_A(t+1) < X_A(t)$, the solution $p_A(t+1)$ at time step $(t+1)$ is rejected and the allocation vector is optimised again with an updated step-length $\eta(t+1) = \frac{\eta(t)}{2}$.

[5] See `https://github.com/Optimizater/Lp-ball-Projection` for more details.

[6] The majorisation step relaxes and linearises the ℓ_p ball to obtain a weighted ℓ_1 ball, and the minimisation step obtains the projection of the point on the ℓ_1 ball.

consider the core-periphery network which has a core of highly connected nodes and other sparsely connected peripheral nodes. The bimodal degree-distribution resembles many real-world leader-follower type networks, and the simplified network structure limits the degrees of freedom which allows us to apply analytical methods to the problem at hand. We employ a degree-based mean-field approximation to determine optimal allocations analytically. The approximation method assumes that nodes with the same degrees have similar behaviours. Nodes are grouped based on their degrees, and the behaviour of the population is approximated by averaging over the behaviours of each group of nodes. Although such an approximation is not always effective, it has been observed to work well in networks where there are no degree correlations [15]. We obtain a degree-based mean-field approximation by following the approach taken in [18], and modifying it to reflect our budget constraint as follows,

$$X_{MF} = \frac{\left(\sum_k \frac{P_k k}{k + a_k^{\frac{1}{\gamma}} + b_k^{\frac{1}{\gamma}}}\right)\left(\sum_k \frac{P_k k a_k^{\frac{1}{\gamma}}}{k + a_k^{\frac{1}{\gamma}} + b_k^{\frac{1}{\gamma}}}\right)}{\sum_k \frac{P_k k (a_k^{\frac{1}{\gamma}} + b_k^{\frac{1}{\gamma}})}{k + a_k^{\frac{1}{\gamma}} + b_k^{\frac{1}{\gamma}}}} + \sum_k \frac{P_k a_k^{\frac{1}{\gamma}}}{k + a_k^{\frac{1}{\gamma}} + b_k^{\frac{1}{\gamma}}}. \tag{4}$$

Here P_k is the degree-distribution of the network and k is the degree of nodes in the network. Additionally, b_k is the competitor's allocation to the group of nodes with degree k and a_k is the controller's allocation to the same group of nodes. The influence experienced by nodes with degree k from both controllers are $a_k^{1/\gamma}$ and $b_k^{1/\gamma}$, and are uniform across all nodes in a given group. We can now use Eq. (4) to determine the optimal allocation patterns in any large, arbitrary core-periphery network structure.

4.2 Core-periphery network

Core-periphery networks of size $N = 1000$ are considered, with a core formed by $P_1 = p_r = 0.25$ (or 25% of the total nodes in the network). Nodes in the highly clustered core have a degree of $k_1 = 30$ and the sparsely connected peripheral nodes have degree $k_2 = 3$. We examine two settings, one where the competitor targets the core and another where they target the peripheral nodes.

Competitor allocations to the core: We first consider the instance where B targets the hub nodes. The competitor budget is distributed uniformly across all nodes in the core of the network. Assuming ϵ_A is the fraction of the total budget allocated to hub nodes by controller A, we obtain an expression for total vote-share using Eq. (4), and use semi-analytical methods to determine the optimal allocation ϵ_A^* that maximise the total vote-shares. We also use the computational methods described in Section 3.1 to determine optimal allocations numerically. For numerical results, we consider 10 instances of core-periphery networks, each of size $N = 1000$, $p_r = 0.25$, $k_1 = 30$ and $k_2 = 3$. Networks are generated using the configuration model [14]. We consider three budget scenarios: (i) insufficient budget $B_A/B_B = 0.1$, (ii) equal budget $B_A/B_B = 1$ and finally (iii) excess budget $B_A/B_B = 10$. In each case, $B_B = N$.

Fig. 1a illustrates the fraction of total resources allocated to the hub nodes in the network, as γ is varied. Fig. 1b shows the corresponding vote-shares obtained by the controller. We find a stark contrast in optimal allocations and vote-shares between the two regimes: (i) $0 < \gamma < 1$ and (ii) $\gamma > 1$. We find that allocation strategies are highly sensitive to controller budgets when allocations have a *delayed effect* on influence i.e. $0 < \gamma < 1$. Whereas when allocations yield *diminishing returns* on influence ($\gamma > 1$), optimal strategies do not change significantly even when budgets differ considerably. A similar phenomenon is also reflected in Fig. 1b, where controller budgets significantly affect vote-shares when $0 < \gamma < 1$, and have significantly less effect on vote-shares when $\gamma \gg 1$. Observe that the linear case, $\gamma = 1$, clearly acts the transition point between the two regimes. Taking a closer look at the region where $0 < \gamma < 1$, we find that optimal allocations oscillate between discrete and continuous configurations (i.e. hub nodes are exclusively targeted or entirely avoided) particularly for limited budget. These fluctuations are more abrupt in analytical results, which maybe due to the *all-or-none* strategy adopted in the mean-field approximation, i.e. all nodes with the same degree are uniformly targeted in the analytical approach. This assumption also results in discrepancies between the analytical and the numerical results. For instance, when $B_A/B_B = 0.1$ we find that while the analytical solution opts for a discrete strategy (targets only the hub nodes) when $0.3 \leq \gamma \leq 0.5$, the numerical solution does not allocate any resources to the hub nodes in these settings and only targets a fraction of the peripheral nodes. Whereas from Fig. 1b we find that the numerical method yields higher vote-shares than the analytical results in this region, thus highlighting the limitations of the mean-field approximation in such instances. Finally, we also observe inconsistencies in Fig. 1a for very high values of γ, i.e. $\gamma > 4$ likely caused by numerical instabilities [1].

Competitor allocations to the periphery: We now consider the instance where competitor B targets the periphery, and we illustrate our results in Fig. 1. Once again, we find that controller budgets heavily impact optimal allocations and vote-shares in the region where $0 < \gamma < 1$. The variation in optimal allocations decreases as γ crosses over into the region where $\gamma \geq 1$. We also find that, for the most part numerical results closely replicate analytical results, with some exceptions. As before, numerical inconsistencies are observed when $\gamma > 4$. We also observe discrepancies between the analytical and the numerical results in the region $0.4 \leq \gamma \leq 0.6$, and argue that such disparities exist as allocations are artificially constrained in the analytical method, whereas numerical approaches are more flexible, thus yielding higher vote-shares.

4.3 Real-world collaboration network

Next, to explore how these results apply to the real world, we explore a collaboration network among network scientists consisting of $N = 379$ scientists

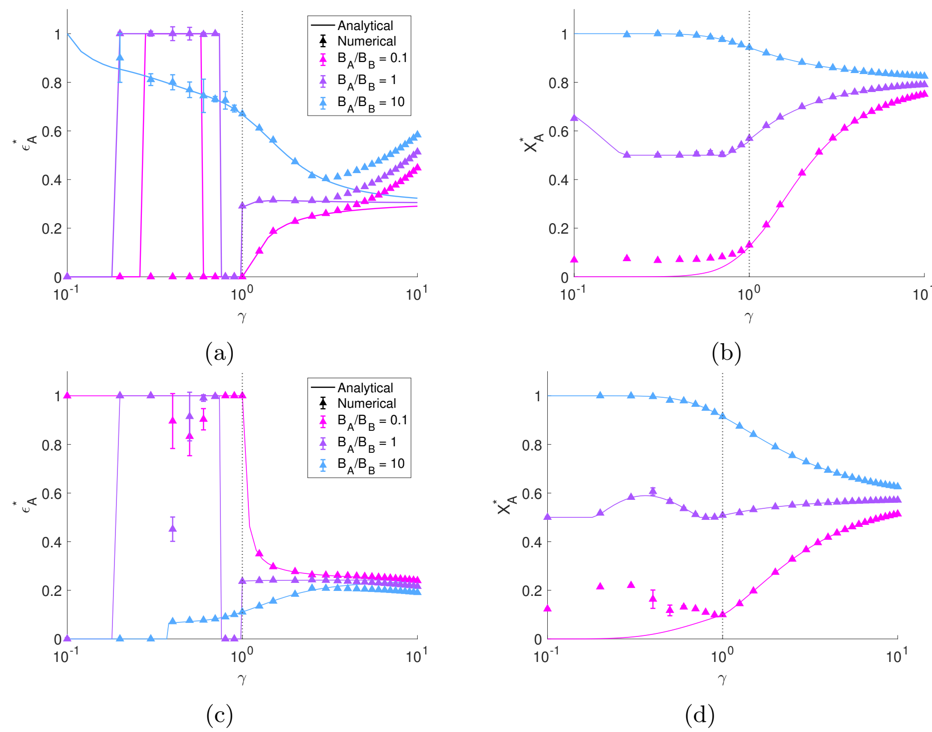

Fig. 1: Figure showing analytical and numerical results for optimal configurations of allocations in a core-periphery network of size $N = 1000$ and average degree $\langle k \rangle = 9.75$. Hub nodes here have degree $k_1 = 30$ while peripheral nodes have degree $k_2 = 3$. The top panel ((a) and (b)) shows results against competitor allocations to the hub nodes. Figure (a) shows the optimal fraction ϵ^*_A of the total budget allocated to hub nodes as γ is varied. Figure (b) shows the corresponding optimal vote-shares X^*_A obtained by the controller. The bottom panel ((c) and (d)) shows results against competitor allocations to the periphery. Figure (c) shows optimal allocations ϵ^*_A to the hub nodes while (d) shows the corresponding vote-shares. Numerical results obtained are a $\mu-$approximation of the optimal solutions where $\mu = 10^{-10}$. Results are averaged over 100 simulations, where algorithms for both the convex and nonconvex optimisation problem are run 10 times with random initialisations of p_A (and $\eta(0) = 1$) on 10 realisations of core-periphery networks. Errorbars show the 95% confidence intervals. The result for $\gamma = 0.1$ and $B_A = 10N$ is missing due to runtime errors caused by numerical overflows.

connected through coauthorship of papers [19,8]. Given that both competitor allocation settings yield similar results in synthetic networks, here for the sake of brevity we focus on the instance where B targets hub nodes. Identifying hub nodes in a heterogeneous network is not a straight-forward process. Here for simplicity, we use the degree centrality measure to distinguish between hubs and peripheral nodes. Given that the average degree of the network is $\langle k \rangle = 4.8$, we assume nodes with degrees above $k > 5$ are hubs, and those with $k \leq 5$ form the periphery of the network. We find that this method classifies nearly 30% of the network as hubs and the rest as the periphery. Since it is infeasible to apply analytical methods to highly heterogeneous network structures, here we rely on numerical approaches to determine optimal allocations. Optimal allocations are determined as the nonlinearity constraint of the allocation vector changes between $0.1 \leq \gamma \leq 10$ and the controller budget is varied between $0.1 \leq B_A/B_B \leq 10$, where $B_B = N$. Results are averaged over 10 simulations, and shown in Fig. 2.

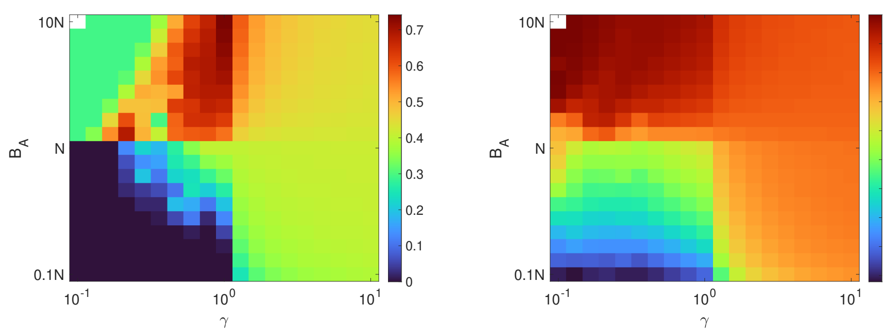

(a) (b)

Fig. 2: Figure showing optimal allocations and vote-shares in a real-world network depicting collaborations among network scientists ($N = 379$ and average degree $\langle k \rangle = 4.8$). Here the competitor B targets hubs nodes. Figure (a) illustrates the optimal fraction of total budget that is allocated to the hub nodes, and Figure (b) shows the corresponding optimal vote-shares for varying γ and B_A. Numerical results obtained are a $\mu-$approximation of the optimal solutions where $\mu = 10^{-10}$. Results are averaged over 10 simulations, where algorithms for both the convex and nonconvex optimisation problems are run 10 times with random initialisations of p_A and $\eta(0) = 1$. The missing values for $\gamma = 0.1$ and $B_A = 10N$ across the heatmaps are due to numerical overflows.

Fig. 2a illustrates the optimal fraction of the total budget allocated to the hub nodes when the competitor targets the hubs. Here too we observe that optimal allocation patterns are highly sensitive to change in budget conditions in the region $0 < \gamma < 1$, whereas optimal allocation patterns are more uniform when $\gamma > 1$. Furthermore, we observe that under low budget conditions, optimal allocation configurations allocate less resources to nodes targeted by the competitor and focus more resources on nodes avoided by the competitor. As the budget increases $B_A > N$, we find that more resources are allocated to the nodes targeted by the competitor. Fig. 2b illustrates the optimal vote-shares obtained in this instance. Similar to results in Section 4.2, we find that vote-shares vary significantly with budgets when $0 < \gamma < 1$, but not as much when $\gamma \geq 1$.

So far we have observed patterns of optimal allocations in settings with nonlinear budget constraints. We now examine how much vote-share a controller gains from optimally targeting a network as opposed to employing a naïve strategy. For comparison, we consider two simple heuristics: (i) the degree-based approach and (ii) uniform allocation. We consider the collaboration network among network scientists for our simulations. We vary γ and B_A, and in each instance determine the optimal allocation and vote-shares numerically. For comparison, we define the allocation vector for the degree-based approach as $p_{A,i} \propto k_i$, where k_i is the degree of a node i, normalised to meet the budget constraint. Additionally, the allocation vector for uniform allocation is given by $p_{A,i} = (B_A/N)^{1/\gamma}$, $1 \leq i \leq N$. The corresponding vote-shares are determined for each instance as X_A^{deg} and X_A^{uni} respectively. Gain in vote-shares is then measured as $[X_A^*/X_A^{deg} - 1]$ and $[X_A^*/X_A^{uni} - 1]$, and results are shown in Fig. 3.

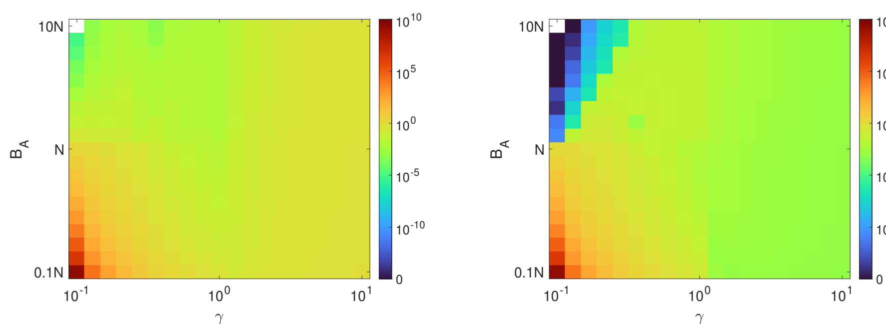

(a) (b)

Fig. 3: Figures showing the gain in vote-shares obtained when employing optimal allocation strategies in a real-world collaboration network ($N = 379$ and $\langle k \rangle = 4.8$), as opposed to simple heuristics such as (i) degree-based targeting or (ii) uniform allocations. The competitor targets the hub nodes. Figure (a) shows the gain in vote-shares $[X_A^*/X_A^{deg} - 1]$ when the optimal strategy is compared to degree-based targeting, whereas Figure (b) shows the gain in vote-shares $[X_A^*/X_A^{uni} - 1]$ when the optimal strategy is compared to uniform allocations. Numerical results obtained are a $\mu-$approximation of the optimal solutions where $\mu = 10^{-10}$. Results shown are mean values obtained over 10 simulations, where algorithms for both the convex and nonconvex optimisation problems are run 10 times with random initialisations of p_A and $\eta(0) = 1$. The missing values for $\gamma = 0.1$ and $B_A = 10N$ across the heatmaps are due to numerical overflows.

Figs. 3a and 3b illustrate gain in vote-shares against competitor allocations to the hub. We find that the sensitivity of gain in vote-shares to controller budget is more in the *delayed influence* setting, as compared to the *diminishing returns* setting. In particular, we observe that the controller can gain significant vote-shares ($\approx 10^{10}$ times) by targeting the network optimally for low values of γ and a low budget B_A. This implies that the effectiveness of the optimal strategy in comparison to other heuristics is higher when individuals take longer (or more resources) to experience or respond to external influence under low budget conditions. We further observe that for larger budgets and low γ, the strategy of targeting the network uniformly yields near-optimal results.

5 Conclusions

Here we explore the competitive influence maximisation problem with continuous allocations in the voter model. Contrary to traditional methods, where nodes are typically targeted in a binary fashion, here we consider continuous allocation of influence to the network where an array of nodes are targeted with varying intensities. We assume that two controllers compete to maximise their vote-shares in the network. Traditionally, the influence maximisation problem has been studied in a linear setting where the cost of influence (or allocations) is analogous to the strength of influence experienced by individuals, and thus the relationship between the cost and effect of influence was expressed using a linear function. However, assuming a linear cost function may be an over-simplification that may not apply to all real-world settings. Thus, here we study the competitive influence maximisation problem for nonlinear cost functions, and we consider settings where the effect of influence varies nonlinearly with the cost of influence. Specifically, we consider two scenarios: (i) where increasing allocations diminishes the marginal strength of influence experienced by nodes, and (ii) where nodes take longer or more allocations to start experiencing the effect of influence (observed as the delayed effect of influence). We find that optimal allocations and vote-shares are highly sensitive to budget conditions where allocations have a delayed effect on influence experienced by nodes. On the contrary, when allocations have a diminishing effect on influence, optimal allocations and vote-shares show limited sensitivity to budget conditions. We further show that targeting the network optimally under low budget conditions and under nonlinear budget constraints can result in significant gain in vote-shares, when compared to more naïve approaches. Our results consider optimal allocations only for known competitor allocations, and an interesting future direction for this work would be to study the problem in a game-theoretic framework that assumes incomplete knowledge of competitor allocations.

References

1. Barbero, A., Sra, S.: Modular proximal optimization for multidimensional total-variation regularization. arXiv preprint arXiv:1411.0589 (2014)

2. Castellano, C., Fortunato, S., Loreto, V.: Statistical physics of social dynamics. Reviews of modern physics **81**(2), 591 (2009)
3. Castellano, C., Muñoz, M.A., Pastor-Satorras, R.: Nonlinear q-voter model. Physical Review E **80**(4), 041129 (2009)
4. Chakraborty, S., Stein, S., Brede, M., Swami, A., de Mel, G., Restocchi, V.: Competitive influence maximisation using voting dynamics. In: Proceedings of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining. pp. 978–985 (2019)
5. Clifford, P., Sudbury, A.: A model for spatial conflict. Biometrika **60**(3), 581–588 (1973)
6. Darko, A., Chan, A.P.: Review of barriers to green building adoption. Sustainable Development **25**(3), 167–179 (2017)
7. Goldstein, D.G., McAfee, R.P., Suri, S.: The effects of exposure time on memory of display advertisements. In: Proceedings of the 12th ACM conference on Electronic commerce. pp. 49–58. ACM (2011)
8. Guimera, R., Danon, L., Diaz-Guilera, A., Giralt, F., Arenas, A.: Self-similar community structure in a network of human interactions. Physical review E **68**(6), 065103 (2003)
9. Holley, R.A., Liggett, T.M.: Ergodic theorems for weakly interacting infinite systems and the voter model. The annals of probability pp. 643–663 (1975)
10. Jain, P., Kar, P., et al.: Non-convex optimization for machine learning. Foundations and Trends in Machine Learning **10**(3-4), 142–363 (2017)
11. Kempe, D., Kleinberg, J., Tardos, É.: 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 (2003)
12. Masuda, N.: Opinion control in complex networks. New Journal of Physics **17**(3), 033031 (2015)
13. Moon, S.: Analysis of twitter unfollow: how often do people unfollow in twitter and why? In: Proceedings of the Third international conference on Social informatics. pp. 7–7. Springer-Verlag (2011)
14. Newman, M.: Networks. Oxford university press (2018)
15. Pastor-Satorras, R., Vespignani, A.: Epidemic spreading in scale-free networks. Physical review letters **86**(14), 3200 (2001)
16. Redner, S.: Reality-inspired voter models: A mini-review. Comptes Rendus Physique **20**(4), 275–292 (2019)
17. Rogers, E.M.: Diffusion of innovations. Simon and Schuster (2010)
18. Romero Moreno, G., Chakraborty, S., Brede, M.: Shadowing and shielding: Effective heuristics for continuous influence maximisation in the voting dynamics. Plos one **16**(6), e0252515 (2021)
19. Rossi, R.A., Ahmed, N.K.: The network data repository with interactive graph analytics and visualization. In: AAAI (2015), `http://networkrepository.com`
20. Tanabe, S., Masuda, N.: Complex dynamics of a nonlinear voter model with contrarian agents. Chaos: An Interdisciplinary Journal of Nonlinear Science **23**(4), 043136 (2013)
21. Tanaka, T., Aoyagi, T.: Optimal weighted networks of phase oscillators for synchronization. Physical Review E **78**(4), 046210 (2008)
22. Yang, X., Wang, J., Wang, H.: Towards an efficient approach for the nonconvex lp ball projection: algorithm and analysis. arXiv preprint arXiv:2101.01350 (2021)
23. Yildiz, E., Ozdaglar, A., Acemoglu, D., Saberi, A., Scaglione, A.: Binary opinion dynamics with stubborn agents. ACM Transactions on Economics and Computation (TEAC) **1**(4), 1–30 (2013)