Competitive Influence Maximisation using Voter Dynamics

Sukankana Chakraborty (Southampton), **Sebastian Stein** (Southampton), **Markus Brede** (Southampton), **Ananthram Swami** (ARL), **Geeth De Mel** (IBM UK), **Valerio Restocchi** (University of Edinburgh).

Objectives

- We identify **optimal strategies for maximising influence**, within **social networks**, in **competitive settings**, under **budget constraints**.
- We use **voter dynamics** to characterise optimal strategies against **known and unknown adversarial strategies**. In cases where the adversarial strategy is unknown, we characterise the **Nash Equilibrium**.
- We introduce a **fixed cost** incurred to **gain access to nodes**, together with the dynamic cost proportional to the influence exerted on the nodes, constrained by the same budget.

Technical Challenges

- Controlling macro-behaviour in complex systems such as social networks, using intervention at micro-level is particularly challenging.
- People are unpredictable - it can be extremely difficult to identify the right incentives for people.

Approaches

- Assuming that influence flows in the network using voter dynamics, the rate at which an influencer A spreads its influence in the network, is given by:

$$\frac{du_{A,i}}{dt} = (1 - u_{A,i}).pr_i(A) - u_{A,i}.pr_i(B).$$

where, $pr_i(A)$ is the probability with which a node i adopts state A and $pr_i(B)$ the probability with which it adopts state B.

- At equilibrium, we get,

$$[L + diag(\boldsymbol{p_A} + \boldsymbol{p_B})]\, \boldsymbol{u_A} = \boldsymbol{p_A}.$$

where, L is the Laplacian of the given network, vectors $\mathbf{p_A}$ and $\mathbf{p_B}$ are resource allocations of competitors A and B respectively, and vector $\mathbf{u_A}$ gives the influence spread of competitor A in the network.

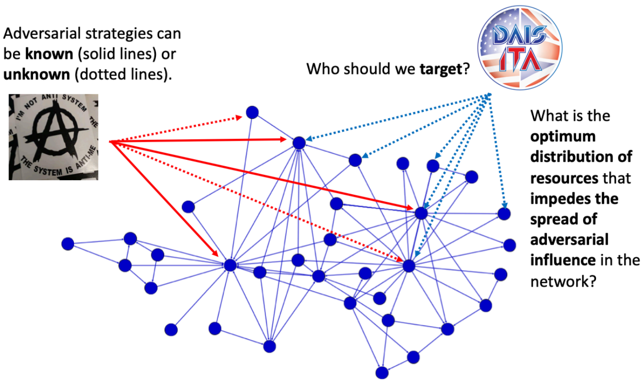

Military & Coalition Relevance

- In various security analytics scenarios, knowledge of the optimal strategy of resource allocation can help optimally influence social networks, such that more people can be helpful to the cause, even in the presence of competing influences.
- Military officials deployed in foreign lands can use similar tactics to win over the hearts of locals.

Results

- For **known adversarial strategies**, the best response strategy is to exert **more influence on nodes avoided by the adversary**, while **avoiding nodes targeted by the adversary**, when adversarial influence on the network exceeds one's own budget. As available budget increases, **more influence is expended on nodes targeted by the adversary**.

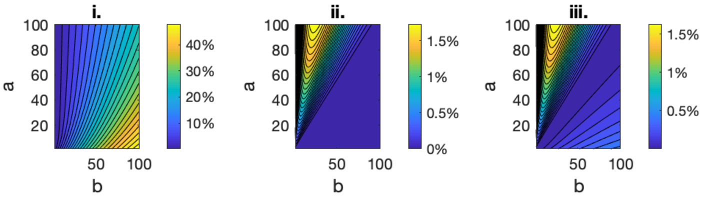

Heatmaps showing the % gain in $u_{A,avg}$, when A plays the optimal strategy in comparison to (i) targeting the hub node, (ii) targeting the periphery and (iii) targeting all nodes uniformly, against adversary B who targets the hub node.

- Where **adversarial strategy is unknown**, we show that the **pure-strategy Nash Equilibrium** is when both competitors **target all nodes equally**.

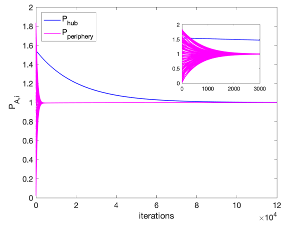

Plot showing the resource allocation vector $\mathbf{p_A}$ over several iterations for unknown adversarial strategies $\mathbf{p_B}$. The distribution converges at the Nash Equilibrium. We observe that the optimum allocation of resources is to target all nodes equally.

- We observe that, **as cost of access changes**, the system **interpolates between discrete and continuous case.**

Summary & Future Work

- We determine optimal strategies of resource allocation that maximises influence in social networks in the presence of adversarial influences.
- We recognise that our study strictly assumes rational behaviour from both players when trying to spread their influence in the network.
- This, may not always be true, or practical, as human behaviour is often irrational.
- To understand how humans strategise in real world scenarios and how it compares to rational or irrational behaviour, we devise a game, where human players compete against an AI to spread influence in a given social network.

Publication(s) & Impact

- Chakraborty S, Stein S, Brede M, Swami A, DeMel G, Restocchi V. Competitive Influence Maximisation using Voting Dynamics. In: The Social Influence Work- shop (5th edition) at the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, Vancouver, Canada.