

Learning a Social Network by Influencing Opinions

Dmitry Chistikov

Centre for Discrete Mathematics and its Applications &
Department of Computer Science, University of Warwick
United Kingdom
d.chistikov@warwick.ac.uk

Mike Paterson

Centre for Discrete Mathematics and its Applications &
Department of Computer Science, University of Warwick
United Kingdom
m.s.paterson@warwick.ac.uk

Luisa Estrada

Department of Computer Science, University of Warwick
United Kingdom
luisa-fernanda.estrada-plata@warwick.ac.uk

Paolo Turrini

Department of Computer Science, University of Warwick
United Kingdom
p.turrini@warwick.ac.uk

ABSTRACT

We study a campaigner who wants to learn the structure of a social network by observing the underlying diffusion process and intervening on it. Using synchronous majoritarian updates on binary opinions as the underlying dynamics, we offer upper bounds on the campaigner’s budget for learning any network with certainty, considering both observation and intervention resources, and further improving them for the case of clique networks. Additionally, we investigate the learning progress of the campaigner when her budget falls below these upper bounds. For such cases, we design a greedy campaigning strategy aimed at optimising the campaigner’s information gain at each opinion diffusion step.

KEYWORDS

Social Networks; Graph Inference; Opinion Diffusion; Synchronous Majority; Influence Maximisation.

ACM Reference Format:

Dmitry Chistikov, Luisa Estrada, Mike Paterson, and Paolo Turrini. 2024. Learning a Social Network by Influencing Opinions. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, Auckland, New Zealand, May 6 – 10, 2024, IFAAMAS, 9 pages.

1 INTRODUCTION

The study of opinion dynamics has been central for distributed artificial intelligence, with important economic [1, 11], political [7], epidemiological [20, 24] and marketing [10, 27] applications. Opinion dynamics provide us with a toolbox to understand social network properties such as network resilience, node centrality, and community structure, as well as a useful predictor for collective decision-making [4, 6]. When the graph structure is known, understanding the dynamics of the influence relations among agents [14, 17] has attracted great attention in computer science. The seminal contribution by Kempe, Kleinberg and Tardos on influence maximisation [16], for example, features a campaigner who observes a social network where nodes hold a (typically binary)

opinion, which is synchronously updated by a function of their incoming connections. In it, the campaigner aims to identify the most influential members in order to drive the network’s collective opinion to a desired outcome. Other recent papers, e.g., [5], add a manipulation layer by studying the influence spread under forms of manipulation such as modifying opinions or links.

The computational studies of influence maximisation in social networks have been largely based on the assumption that the campaigner has access to the graph structure; an assumption that is not satisfied in many real world scenarios. For instance, a marketing firm may want to reach out to the most influential agents without knowing beforehand who they are. They may have to estimate the agent’s influence power from simply observing how the opinions change in the network as a function of what they say.

In this paper, we analyse situations where the campaigner aims at learning an unknown network structure by modifying some of the agents’ opinions and observing the resulting dynamics. We carry out our analysis on social networks with binary opinions and synchronous majority dynamics [16], which, despite its descriptive simplicity, displays a highly complex structure [9]. We start with a campaigner endowed with budgets (*Obs*) and (*Int*) for the costs she incurs in terms of the observations and interventions, respectively, carried out on the network.

Clearly, under the synchronous majority dynamics, different graphs could be behaviourally identical when starting from the same network labelling. Take the example in Figure 1, where the campaigner observes the transition from time t to time $t + 1$ in Figure 1a and notices that both networks in Figure 1b are consistent with it. The reader may verify that, in total, there are only nine feasible networks that admit this single observation, compared to the original 64 possible networks that could be generated from these three agents. Therefore, even if the campaigner cannot pinpoint the exact underlying network after one step of the opinion diffusion, it still refines her knowledge, allowing for a more educated guess.

We are interested in instructing the campaigner on how to spend her budget so she can correctly identify the underlying network. Notice that each opinion diffusion step rules out all those networks that are not consistent with the observed behaviour.

Our contribution. First, we study the campaigner’s budget bounds required to learn the network exactly. Specifically, Theorem 1 proves that this can be done for any loop-free network of n agents by spending $O(n^2)$ of the observation budget and intervening on



This work is licensed under a Creative Commons Attribution International 4.0 License.




Figure 1: The campaigner’s problem. $L := \{\circ, \bullet\}$ is the set of labellings and \otimes is an unlabelled agent. Each directed edge denotes a plausible influence relation under synchronous majority dynamics, given the observations.

$O(n^3)$ agents. Moreover, we improve these bounds if inquiring about a specific network structures. We elaborate on cliques, as agents in them are the most resilient to being manipulated and the campaigner is forced to intervene since no additional information can be retrieved from observing more than one opinion diffusion step. Table 1 summarises our technical contributions to these tasks.

Furthermore, we investigate the problem of learning the network when the campaigner has a budget that falls below the general $O(n^3)$ upper bound, and we look at optimal intervention strategies to maximise the information gain.

Table 1: Summary of the campaigner’s budget upper bounds for different learning tasks.

Learning task	Observation budget	Intervention budget
Learn any network $G \in \mathcal{H}$	$O(n^2)$	$O(n^3)$
Identify an odd clique	$O(n)$	$O(n^2)$
Identify an even clique	$O(n)$	$O(n)$

Related literature. Our approach connects to three prominent lines of research in the distributed AI and multi-agent systems literature: influence maximisation [5, 7, 16], opinion dynamics [2, 3, 9] and network inference [22, 24].

Influence Maximisation. Tracing back to the seminal contribution by Kempe, Kleinberg, and Tardos [16] for the Independent Cascade and Linear Threshold models, several efficient algorithms have been devised for optimal campaigning strategies, such as DD [8], CELF [20], TIM and TIM+ [26], under the assumption that the network parameters are known. In the context of computational social choice, Bredereck and Elkind [5] studied budgeted manipulations of networks with synchronous majority dynamics, evaluating the success of targeting a set of agents based on the optimistic and pessimistic scenarios of the converging opinion distribution. Similarly, Castiglioni et al. [7] studied how manipulators exploit the network structure to transmit both positive and negative messages to sway election results, trying to maximise the improvement in the victory margin. Once again, these studies assume knowledge of the network structure.

Online influence maximisation in graphs with unknown edge weights has been investigated in the literature [18, 19, 24] developing a campaigner’s strategy based on a learning-through-feedback

pipeline. We draw on this approach, but without making assumptions on the number of influencers each node has and, instead, introduce a temporal dimension in a similar spirit to [10].

Opinion Dynamics. Models of opinion dynamics have been widely investigated in the multiagent systems community, studying asynchronous [2, 3] and synchronous consensus-reaching majority dynamics [5, 9]. In [2] and [3], the authors identified computational conditions under which a given set of agents (seeds) can determine a consensus opinion, finding feasible complexity results for such set being half of the agents, and infeasible otherwise. Similar considerations, from a manipulation angle, are drawn in Auletta et al. [1]. In contrast, our approach focuses on minimising the campaigner’s budget spent to learn a network, for which we found that a consensus is the least informative state since any network’s opinion will be stable once reaching it.

Network Inference. This branch of research, often found in epidemiology, studies a network-learning task by using a cascade model and solving it via convex optimisation. Myers and Leskovec [22] and Netrapalli and Sanghavi [24] studied contagion traces over time, knowing when a node became infected but not who infected it. Subsequent works [12, 13] proved that finding the most likely underlying graph is NP-hard to solve exactly given only some infecting times. So, instead, they developed approximation algorithms to retrieve core-periphery structures in social media. Additionally, Narasimhan, Parkes and Singer [23] demonstrated the PAC learnability of influence relations in social networks.

Paper structure. Section 2 provides the needed technical background. Section 3 establishes key upper bounds on the campaigner’s budget, and Section 4 studies learning strategies under fixed budgets. Section 5 concludes and suggests future directions.

2 PRELIMINARIES AND NOTATION

Social Networks. Following Easley and Kleinberg [11], we define a social network as a directed graph G over the set of agents $N = \{1, \dots, n\}$, $n \in \mathbb{N}$. More specifically, we think of G as a collection of edges, such that each edge $(i, j) \in G$ represents an influence relation of agent i on agent j , where we say that agent j follows agent i or, equivalently, that agent i is an *influencer* of agent j .

For technical convenience, we will consider social networks with no self-loops and define the sets $G_i := \{j \in N \mid (i, j) \in G\}$ and $G_i^{-1} := \{j \in N \mid (j, i) \in G\}$ to denote the followers and influencers, respectively, of agent $i \in N$.

Additionally, we assign labels to the nodes in G to attribute specific traits to the agents. Let $L^N := \{\ell: N \rightarrow L\}$ be the space of all L labellings over G , where each label in L embodies an opinion and each $\ell \in L^N$ records the agents’ opinions at a given time. In principle, the labelling functions in L^N can be customised to fit any opinion diffusion mechanism. Here, we focus on the binary synchronous majority protocol, where agents can choose between two alternative opinions, and every opinion diffusion step assigns to each agent the strictly most popular opinion among its influencers or maintains its previous opinion in case of a tie.

To avoid confusion between 0 and 1 as labels, integers or bits, our binary label set will be $L := \{b, w\}$ for opinions *black* and *white*, depicted as \bullet and \circ .

Definition 1 (Labelled Social Networks). *A triplet (N, G, ℓ) is a labelled social network, where G is a directed graph, N is a set of agents and $\ell : N \rightarrow L$ is a labelling function.*

Our campaigner will start with an observation budget Obs and intervention budget Int , which she will spend to learn the network. Each time step will result in a new network configuration presented to the campaigner. So, we can think of each time step as one unit spent from Obs , where, additionally, the campaigner can choose to modify the opinion of a set of agents, in exchange for spending one unit from Int per agent intervened on. Clearly, only observing is equivalent to intervening on the empty set of agents.

Furthermore, fixing a labelled social network (N, G, ℓ) we denote by A_i the set of agents that agree with agent i and by D_i the agents that do not. Also, we use $A_i^{-1} := A_i \cap G_i^{-1}$ (respectively, $D_i^{-1} := D_i \cap G_i^{-1}$) to identify the influencers agent i agrees (respectively, disagrees) with.

We say that the *opinion imbalance* of agent i is $m_i \in \mathbb{Z}$, where $m_i := m_i(\ell) = |A_i^{-1}| - |D_i^{-1}|$. When $m_i = 0$ or $m_i = \pm 1$, we will refer to agent i 's opinion as being *balanced* or *nearly balanced*, respectively. We can now define how the network evolves.

Definition 2 (Opinion Diffusion Step). *Given a labelled social network (N, G, ℓ) , a binary synchronous majority opinion diffusion step results in the labelled social network (N, G, ℓ_G^+) , with $\ell_G^+ : N \rightarrow L$,*

$$\ell_G^+(i) := \begin{cases} \ell(i)^c & \text{if } m_i(\ell) < 0, \\ \ell(i) & \text{otherwise,} \end{cases}$$

where $\ell(i)^c$ is the complement of $\ell(i)$ and $m_i(\ell)$ is the opinion imbalance of i in (N, G, ℓ) .

This protocol establishes that an agent changes its opinion only if the influencers that disagree with it on (N, G, ℓ) outnumber those who agree. Furthermore, to represent the opinions after t diffusion steps, we will use $\ell_G^{t+} : N \rightarrow L$ for $t \leq Obs$.

Remark 1. *Since the campaigner will not know the underlying graph G in (N, G, ℓ) , we omit the subscript in ℓ_G^+ when looking at an opinion diffusion step from her perspective. That is, we will use $\ell^+ : N \rightarrow L$ to denote that there is at least one directed graph G among the possible networks that satisfies that $\ell^+(i) = \ell_G^+(i)$ for every $i \in N$.*

Exact Learning. In our framework we view the campaigner as a learner and, drawing from computational learning theory and query-based learning [15, 25], view the opinion states associated with the observed diffusion steps as answers to queries submitted to an oracle. Consequently, the campaigner's observation budget is simply a bound on the total number of queries. Additionally, the campaigner uses the intervention budget to refine such queries.

Following this parallel, our *hypothesis space* \mathcal{H} is the set of all possible influence relations between the agents in N , which the learner shrinks to reach a target concept, the target graph $G \in \mathcal{H}$.

Definition 3 (Exact Learning). *Let G be a directed graph in \mathcal{H} . We say that the campaigner can exactly learn G after t opinion diffusion steps if she can infer a nested hypothesis space sequence $\mathcal{H} \supseteq \mathcal{H}_1 \supseteq \dots \supseteq \mathcal{H}_t$ such that $\mathcal{H}_t := \{G\}$.*

The success of a campaign lies in the campaigner being able to exactly learn G without exceeding her budget. Furthermore, we

will require each element of the nested hypothesis spaces to be consistent with the observed opinion diffusion steps.

Later in the paper, we will go beyond exact learning and require a measure to establish the likelihood of the campaigner correctly “guessing” the network. To do so, we assume the graphs in \mathcal{H} to be distributed following \mathcal{D} and define $X \sim \mathcal{D}$ as the random variable of choosing any one of them. Thus, when learning the graph $G \in \mathcal{H}$, we will want to estimate the probability of the event $X = G$, as it predicts the campaigner's success rate. Given a sequence of observations, the campaigner gets to choose from a shrunken hypothesis space $\mathcal{H}_t \subseteq \mathcal{H}$, so her success probability is actually $\mathbb{P}(X = G \mid \mathcal{H}_t) = \mathbb{P}(X = G) / \mathbb{P}(X \in \mathcal{H}_t)$. In the definition of information gain that follows, we express the campaigner's learning progress in terms of Shannon's information content [21].

Definition 4. *Let (N, G, ℓ) be a labelled social network, and suppose $\mathcal{H}_t \subseteq \mathcal{H}$ is the hypothesis space of networks consistent with the opinion diffusion steps observed up to time $t < Obs$. If $X \sim \mathcal{D}$ is the random variable of choosing a graph in \mathcal{H} , for some distribution \mathcal{D} , then the information gained from an additional observation is¹*

$$\begin{aligned} I_{t+1} &:= \log_2(\mathbb{P}(X = G \mid \mathcal{H}_t)) - \log_2(\mathbb{P}(X = G \mid \mathcal{H}_{t+1})) \\ &= \log_2\left(\frac{\mathbb{P}(X \in \mathcal{H}_t)}{\mathbb{P}(X \in \mathcal{H}_{t+1})}\right), \end{aligned}$$

where $\mathcal{H}_{t+1} \subseteq \mathcal{H}_t$ is the refinement of \mathcal{H}_t given the opinion diffusion observed at time $t + 1$.

3 BOUNDS ON THE CAMPAIGNER'S BUDGET

Before trying to infer a graph $G \in \mathcal{H}$, we present some sufficient conditions for determining whether there is an influence relation between two agents with absolute certainty. We break the cost of this sub-task into the number of opinion diffusion steps (capped by Obs) the campaigner needs to observe on the network as a whole and the number of interventions with agents she needs to make (capped by Int). These will be the building blocks for establishing the general upper bounds on the campaigner's budget, and will come together in Algorithm 1 to equip the campaigner with a strategy that learns any network by spending $\mathcal{O}(n^3)$ on intervention and $\mathcal{O}(n^2)$ on observation.

Lemma 1. *Let (N, G) be a social network and suppose $\ell_1, \ell_2 \in L^N$ are two network labellings such that $\ell_1 := (\ell_1(1), \dots, \ell_1(n))$ and $\ell_2 := (\ell_2(1), \dots, \ell_2(n))$ satisfy that:*

- (1) *The Hamming distance d_H between ℓ_1 and ℓ_2 , or ℓ_1 and $\ell_2^c := (\ell_2(1)^c, \dots, \ell_2(n)^c)$, is 1, differing only on agent $j \in N$.*
- (2) *Agent $i \in N \setminus \{j\}$ is such that*
 - (a) *if $|G_i^{-1}|$ is even then $m_i(\ell_1) = 0$ and $\ell_1(i) = \ell_1(j)$, or*
 - (b) *if $|G_i^{-1}|$ is odd then $m_i(\ell_1) = -1$ and $\ell_1(i) = \ell_1(j)^c$.*

Then $j \in G_i^{-1}$ if and only if $\ell^+(i) = \ell(j)$, for $\ell \in \{\ell_1, \ell_2\}$.

We obtain (\Leftarrow) by evaluating the conditions for the opinion's transitions given. Note that agent i behaves differently coming from ℓ_1 and ℓ_2 , in the sense that, in one case it changes its opinion, while in the other it maintains it. Still, the only possible source of opinion discrepancies can be agent j , the only one with $\ell_1(j) \neq \ell_2(j)$. Hence,

¹ If $X \sim \text{Uniform}(\mathcal{H})$, the information gain is simply the logarithm of the ratio between $|\mathcal{H}_t|$ and $|\mathcal{H}_{t+1}|$.

agent i must follow agent j , as otherwise the sets of influencers who agree and disagree with it would be the same.

To obtain (\Rightarrow) , we observe the opinion imbalance of agent i assuming it follows agent j . Yet, since $\ell_1(j) \neq \ell_2(j)$ and $d_H(\ell_1, \ell_2) = 1$ as per [21, p. 206], we have that $m_i(\ell_1) = 0$ implies $m_i(\ell_2) = -2$ for case (2a), as agent j 's opinion goes from agreeing with agent i in (N, G, ℓ_1) to disagreeing in (N, G, ℓ_2) . Similarly, $m_i(\ell_1) = -1$ yields $m_i(\ell_2) = 1$ for (2b), but in this case agent j goes from disagreeing to agreeing. Notice how in both cases agent i 's subsequent opinion coincides with what agent j says.

Remark 2. *Lemma 1 provides sufficient conditions for the campaigner to identify with certainty the existence of the edge (j, i) in the underlying graph G . These conditions make agent j the only possible tiebreaker for any agent whose influencers' opinions are known to be balanced (or nearly balanced). Yet, getting to know that is usually the challenging part.*

Algorithm 1

```

1: Input: A social network  $(N, G)$ .
2:  $\triangleright$  Goal: Return a list of Influencers such that  $Influencers[i] = G_i^{-1}$ , for all  $i \in N$ .
3: Initialise  $Influencers = [\{\}] \times N$ .
4: for all  $i \in N$  do
5:   Intervene on the network such that there is a consensus.
6:   Store the network's opinion state as vector  $\ell = ones(n)$ .
7:    $\triangleright$  The consensus determines if 1 encodes opinion  $b$  or  $w$ .
8:   Set  $pivot = 0$ .
9:   repeat
10:     $\triangleright$  Agents have opinion 1 if their indices are higher
        than the pivot's and 0 otherwise.
11:    if  $pivot \neq i$  then
12:      Update  $\ell[pivot] = 0$ 
13:      Intervene on the network to match the labelling in  $\ell$ .
14:      Observe an opinion diffusion step.
15:      Store the network's opinion state as vector  $\ell^+$ .
16:      Update  $pivot = pivot + 1$ .
17:    until  $\ell^+[i] = 0$  or  $(\ell^+[i] = 1$  and  $pivot = n)$ 
18:     $\triangleright$  Move on to the next agent if  $\ell^+[i] = 1$  and  $pivot = n$ .
        Agent  $i$  has no influencers.
19:    Add  $pivot - 1$  to  $Influencers[i]$ .
20:    Store the last query as the opinion state vector  $\ell_{p^*}$ .
21:    for all  $j \in N \setminus \{i\}, j \neq pivot$  do
22:      Intervene on the network so it matches  $\ell_{p^*}$ .
23:       $\triangleright$  Induce a labelling satisfying Lemma 1's hypothesis.
24:      if  $j > pivot$  then
25:         $\triangleright$  Agent  $j$  has not been intervened on.
26:        Change the opinion of the  $pivot$  agent.
27:       $\triangleright$  Agent  $j$  agrees with the  $pivot$  agent.
28:      Change the opinion of agent  $j$ .
29:      Store the current network's opinion state as vector  $\ell$ .
30:      Observe one opinion diffusion step.
31:      Store the network's opinion state as vector  $\ell^+$ .
32:      if  $\ell^+[i] = \ell[j]$  then
33:        Add  $j$  to  $Influencers[i]$ .  $\triangleright j \in G_i^{-1}$  by Lemma 1.
```

We design Algorithm 1 such that, for each agent $i \in N$, it first finds an opinion state where its opinion imbalance is $|m_i| \leq 1$ and then it artificially queries for the diffusion step of adjacent opinion states satisfying the conditions of Lemma 1. For the first part, the algorithm starts with a consensus and has the campaigner progressively intervening on agents, one by one, making them disagree with agent i until a pivotal agent $p^* \in N$ surfaces. This agent p^* will be the one who tips the opinion balance, thereby being the first to incite agent i to change its opinion in the diffusion step. Furthermore, the labellings before and after intervening on it, namely ℓ_{p^*-1} and ℓ_{p^*} , will serve as a reference to determine the remaining influencers of agent i as, in at least one of them, m_i must be (nearly) balanced.

However, it turns out that finding an adjacent labelling in the second stage of Algorithm 1, satisfying the conditions of Lemma 1 to determine if $j \in G_i^{-1}$, depends only on whether the campaigner had previously intervened on agent j . Specifically, if $\ell_{p^*}(j) = \ell_{p^*}(i)^c$, the adjacent opinions states to observe would be

$\ell_1 = (\ell_{p^*}(1), \dots, \ell_{p^*}(n))$ and $\ell_2 = (\ell_{p^*}(1), \dots, \ell_{p^*}(j)^c, \dots, \ell_{p^*}(n))$,

whereas if $\ell_{p^*}(j) = \ell_{p^*}(i)$, these would be:

$$\begin{aligned} \ell_1 &= (\ell_{p^*-1}(1), \dots, \ell_{p^*-1}(n)) \text{ and} \\ \ell_2 &= (\ell_{p^*-1}(1), \dots, \ell_{p^*-1}(j)^c, \dots, \ell_{p^*-1}(n)). \end{aligned}$$

Recall that the labellings ℓ_{p^*-1} and ℓ_{p^*} only differ on agent p^* . Finally, computing the costs incurred by the campaigner following Algorithm 1 yields the learning upper bounds of Theorem 1.

Theorem 1. *The campaigner can learn the underlying graph G of any social network (N, G) by spending $O(n^2)$ on observation and $O(n^3)$ on intervention.*

We derive these upper bounds by evaluating the number of queries done in Algorithm 1. It queries at most $n-1$ one-step opinion diffusion to find the pivot agent p^* that makes agent $i \in N$ change opinion for the first time. From Lemma 1, this implies that agent i follows p^* and provides a reference network labelling $\ell \in \{\ell_{p^*-1}, \ell_{p^*}\}$, chosen from the last two queries, where the campaigner knows that $|m_i(\ell)| \leq 1$. After the pivot is found, Algorithm 1 observes the opinion diffusion of another $n-2$ opinion states, all chosen adjacent to ℓ , to allow the campaigner to know with certainty whether agent i follows the remaining agents of the network. Iterating over all the agents in N bounds the observation budget with $O(n^2)$, while the intervention budget is $O(n^3)$, with an additional order of magnitude coming from the campaigner having to intervene on at most $n-1$ agents to generate each one of the adjacent opinion states.

Remark 3. *Theorem 1 shows that there are no two indistinguishable graphs in our hypothesis space \mathcal{H} using binary synchronous majority opinion diffusion. Thus, the campaigner can always learn the input network given enough resources, irrespective of the initial labelling.*

3.1 Determining if a network is a clique

We enquire how much easier it would be for the campaigner to identify a specific network structure. Take cliques, for example, where all the agents in the network influence each other. These structures represent tightly-knit groups where information, opinions and behaviours have higher chances to spread. Yet they are

more resilient to manipulation, since the campaigner must intervene on at least half of the agents to alter their consensus, unless opinions are evenly split. Additionally, agents will reach the said consensus in a single opinion diffusion step, meaning that no additional information can be retrieved by merely observing, forcing the campaigner to intervene. We formalise this in Lemma 2.

Lemma 2. *Let (N, G, ℓ) be a binary-labelled social network where G is a clique and $\ell : N \rightarrow \{b, w\}$. Then, for every $i \in N$ we have that*

- (1) $\ell_G^+(i) = \ell(i)^c$ and $\ell_G^{+2}(i) = \ell(i)$ if $|\{i \in N : \ell(i) = b\}| = \frac{n}{2}$,
or
- (2) $\ell_G^+(i) = \text{argmax}_{l \in L} |\{j \in N : \ell(j) = l\}|$ otherwise.

Only in even cliques opinions can split agents exactly in half. In this case, all agents have $n/2$ influencers who disagree with them and $n/2 - 1$ influencers who agree since G does not have loops. Thus, $m_i = -1$ for every $i \in N$, and so, every agent changes its opinion, leading to an opinion state with swapped labels. Consequently, all agents will behave the same for the next opinion diffusion step, changing back to their original opinion and making a 2-cycle. Conversely, as soon as the opinions in the clique are not balanced, the agents will always succumb to the majority's opinion. Note that this will always be the case with an odd number of agents.

3.1.1 Odd cliques.

Proposition 1. *Let (N, G, ℓ) be a binary-labelled social network where n is odd. The campaigner can determine if G is a clique by using $2n + 1$ observations and making n^2 interventions.*

When G is an odd clique, opinions will reach consensus in one step, as per Lemma 2. Thus, to get the upper bound in Proposition 1, the campaigner will pick an agent $i \in N$ in a consensus and intervene on half of the remaining agents. If the network is a clique, its opinion should go back to the previous consensus. Next, if she switches again the opinions of that same half of agents together with agent i , all opinions in G should converge to the opposite consensus, which means that every agent in $N \setminus \{i\}$ follows agent i by Lemma 1. By repeating this over all agents, the campaigner learns that everyone follows everyone in the social network.

3.1.2 Even cliques.

We establish the upper bounds on the campaigner's observation and intervention budgets for even cliques by employing a contradiction-based algorithm. She will submit a set of queries, either until observing an infeasible transition for a clique or until no other network admits the history of opinion diffusion steps except a clique.

We present our learning technique in Algorithm 2, which we will refer to as *The Windmill*. Intuitively, it takes advantage of the cyclic behaviour presented in Lemma 2 by instructing the campaigner to intervene on pairs of agents without disrupting the opinion balance of the network. This way, the campaigner will still believe that the network is a clique if and only if her intervention leads to a complementary network labelling after one opinion diffusion step.

Intuitively, the algorithm is called *The Windmill* because, if the network's opinion is balanced, the campaigner can gather the agents in a circle and draw an imaginary vertical line such that all agents with opposite opinions sit on opposite sides of it. Then, to decide which agents to intervene on, she can rotate the line, just like a windmill, and flip the opinions of the agents who changed

sides. See Figure 2 for an example. Also, notice how it suffices to observe one step of the opinion diffusion since the resulting labelling is equivalent to the preceding opinion state, in the sense that all agents have their same opinion imbalance.




Figure 2: Graphical intuition of *The Windmill*.

Algorithm 2 The Windmill

Input: A social network (N, G, ℓ) where n is even and $\ell \in L^N$, $L = \{b, w\}$, such that $|\{i \in N : \ell(i) = b\}| = |\{i \in N : \ell(i) = w\}|$. Initialise the set of marked agents $M = \{\}$. Let $B(\ell) := \{i \in N : \ell(i) = b\}$ and $W(\ell) := N \setminus B(\ell)$, $\forall \ell \in L^N$.

while $|N \setminus M| > 0$ **do**

- Observe the next opinion labelling ℓ^+ .
- if** $B(\ell^+) = W(\ell)$ **then**
 - ▷ Opinion diffusion is consistent with Lemma 2. ◀
 - Select some agents $i \in W(\ell^+) \setminus M$ and $j \in B(\ell^+) \setminus M$.
 - Mark the agents (i, j) by adding them to M .
 - Change the opinions of agents i and j .
 - Update ℓ to coincide with the current network labelling.
- else**
 - Exit the algorithm and conclude that the social network is **NOT** a clique.

▷ So far, there is no evidence that contradicts G being a clique. ◀

Finish *The Windmill* and return the set of pairs of agents M .

Definition 5. *Let N be a set of agents, ℓ a labelling function and \mathcal{H} a hypothesis space. Then, $(N, \cdot, \ell)_{\mathcal{H}} := \{(N, G, \ell) : G \in \mathcal{H}\}$ is the family of labelled social networks admitted in \mathcal{H} .*

After observing a sequence of opinion diffusion steps, the campaigner refines her knowledge by discarding all inconsistent network configurations. However, once the space has reduced sufficiently, some influence relations may exist in either all or none of the feasible graphs. Take for example, a reduced hypothesis space consisting of just the two networks in Figure 1b. In it, the campaigner is certain that the agent at the top follows the agent on the left and influences the agent on the right. Still, whether the agent on the right influences the agent on the left remains uncertain. Similarly, there might be case where a set of agents collectively influences another agent. That is, either all the agents in the set are influencers, or none of them are. We call these *simultaneous influencers* and they will play a key role when establishing what the campaigner can learn from each run of *The Windmill*.

Definition 6. Let $(N, \cdot, \ell)_{\mathcal{H}}$ be a family of binary labelled social networks. Then, the agents in a subset $\Sigma \subseteq N \setminus \{k\}$ are simultaneous influencers of agent $k \in N$ if, for every $G \in \mathcal{H}$ and for all $i, j \in \Sigma$, $i \neq j$, we have that $i \in G_k^{-1}$ if and only if $j \in G_k^{-1}$.

Lemma 3. If *The Windmill* algorithm finishes without finding a contradiction to the clique belief and returns a sequence of agent pairs (i_t, j_t) , for $t = 1, \dots, n/2$, then the campaigner learns that:

- (1) agent $i_t \in G_{j_t}^{-1}$ and agent $j_t \in G_{i_t}^{-1}$, for $t = 1, \dots, n/2$,
- (2) the agents i_t and j_t , $t = 1, \dots, n/2$, are simultaneous influencers of every agent $k \in N \setminus \{i_t, j_t\}$.

PROOF. Part (1) follows from noting that the campaigner intervenes on two agents i and j with opposite opinions and that, thanks to Lemma 2, they should change opinions in the next opinion diffusion step. However, all agents who agreed with agent i (and, similarly, with agent j) before the intervention will then disagree, and vice versa. The only agent who maintains its opinion relative to agent i is agent j . This means that, aside from agent j , agent i has the same number of influencers who agree and disagree with it on the intervened and non-intervened states. Yet, agent i changes its opinion in both cases. Thus, agent j is, in fact, the tiebreaker.

For Part (2), we focus on three network labellings $\ell_0, \ell_1, \ell_{\frac{n}{2}} \in L^N$ shown in Figure 3. Note that these states correspond to the initial state and the first and last interventions of *The Windmill*. In the initial state (N, G, ℓ_0) we draw a vertical line and consider the four agents $i_1, i_2, j_1, j_2 \subseteq N$ closest to it, such that $\ell(i_1) = \ell(i_2)$ and $\ell(j_1) = \ell(j_2)$. We choose one of them, for instance $k := i_1$, to be our reference and notice that the agents in $N \setminus \{k, i_2, j_1, j_2\}$ can be grouped into two sets A_k^* and D_k^* where agents within them agree with one another in $\ell \in \{\ell_0, \ell_1, \ell_{\frac{n}{2}}\}$.




Figure 3: Three network labellings seen in *The Windmill*.

If *The Windmill* does not halt, it means that all agents change their opinions after one opinion diffusion step, including agent k . So $m_k(\ell) < 0$ for $\ell \in \{\ell_0, \ell_1, \ell_{\frac{n}{2}}\}$. This is equivalent to having more agents disagreeing with agent k than agreeing with it, which can be written in terms of A_k^* , D_k^* and whether agents j_1, j_2 and i_2 influence agent k , as follows:

$$\ell_0 : |D_k^*| + \mathbb{1}_{G_k^{-1}}(j_1) + \mathbb{1}_{G_k^{-1}}(j_2) > |A_k^*| + \mathbb{1}_{G_k^{-1}}(i_2), \quad (1)$$

$$\ell_1 : |D_k^*| + \mathbb{1}_{G_k^{-1}}(i_2) + \mathbb{1}_{G_k^{-1}}(j_1) > |A_k^*| + \mathbb{1}_{G_k^{-1}}(j_2), \quad (2)$$

$$\ell_{\frac{n}{2}} : |A_k^*| + \mathbb{1}_{G_k^{-1}}(i_2) + \mathbb{1}_{G_k^{-1}}(j_1) > |D_k^*| + \mathbb{1}_{G_k^{-1}}(j_2), \quad (3)$$

where $\mathbb{1}_{G_k^{-1}}(x)$, $x \in N$, is the indicator function that is 1 if $x \in G_k^{-1}$, and 0 otherwise. From part (1), we have that $\mathbb{1}_{G_k^{-1}}(j_1) = 1$. Moreover, $|A_k^*| = |D_k^*|$ since agent k would otherwise have maintained

its opinion after at least one of the labellings ℓ_1, ℓ_2 or $\ell_{\frac{n}{2}}$. Therefore, the system of inequalities (1), (2) and (3) simplifies into

$$1 + \mathbb{1}_{G_k^{-1}}(j_2) > \mathbb{1}_{G_k^{-1}}(i_2) \quad \text{and} \quad 1 + \mathbb{1}_{G_k^{-1}}(i_2) > \mathbb{1}_{G_k^{-1}}(j_2),$$

which holds only when $\mathbb{1}_{G_k^{-1}}(i_2) = \mathbb{1}_{G_k^{-1}}(j_2)$, showing that agents i_2 and j_2 are simultaneous influencers of agent k .

Note that these results apply for any agent $k \in N$ taken as a reference, as once *The Windmill* finishes without halting, all agents would have perceived opinion states equivalent to ℓ_0, ℓ_1 and $\ell_{\frac{n}{2}}$. Furthermore, if i_2 and j_2 are simultaneous influencers of agent k , to establish that the next pair of agents intervened on, namely (i_3, j_3) , are also simultaneous influencers, we update

$$A_k^* \leftarrow A_k^* \setminus \{i_3, j_3\} \quad \text{and} \quad D_k^* \leftarrow D_k^* \setminus \{i_3, j_3\}.$$

Replicating our analysis leads to the same system of inequalities since the indicator functions corresponding to i_2 and j_2 are always on opposite sides and, thus, do not affect the simplifications. \square

Lemma 3 describes the campaigner's knowledge after running *The Windmill* once without it taking a premature exit and returning the set $M_1 := \{(i_t, j_t) \in N \times N : i_t \neq j_t, t = 1, \dots, n/2\}$ of pairs of agents intervened on concurrently. *The Windmill* will finish after completing half a revolution of the imaginary line separating agents with opposite opinions and, at each step, the campaigner intervenes on two agents. Thus, a completed run consumes $n/2$ of *Obs* and n of *Int*. In the following Proposition 2, we will show how to leverage the information retrieved from the first run to create a second intervention set M_2 , such that the campaigner can know for sure if the underlying network is a clique or not after running *The Windmill* again but targeting the agent pairs in it.

Proposition 2. Let (N, G, ℓ) be a binary-labelled social network where n is even and the labelling $\ell \in L^N$ has balanced opinions. The campaigner can determine if G is a clique by spending n of her observation budget and $2n$ of her intervention budget.

Intuitively, to prove Proposition 2 we keep track of the agents' positions in both runs of *The Windmill*, moving the agents from the first position to the second via some permutation, and then merging the resulting simultaneous influencers. Note that there will be many permutations that allow the campaigner to infer that the network is a clique. We use the permutation $\sigma : N \rightarrow N$ that moves the agents in $D_k^* \cup \{j_1, j_2\}$ from Figure 3 one position clockwise within themselves, while leaving the others fixed. Moreover, we set the merging rule to be that: if two sets $\Sigma_1, \Sigma_2 \subseteq N$ of simultaneous agents are not disjoint, then all agents in $\Sigma := \Sigma_1 \cup \Sigma_2 \subseteq N$ are simultaneous influencers. After running *The Windmill* the second time, we obtain that M_1 and $M_2 := \sigma(M_1)$ do not share any pair of agents intervened on simultaneously and merging results in $N \setminus \{k\}$ being a set of simultaneous influencers of our reference agent k . However, because agent k at least follows its pair in M_1 (and in M_2), then it must be the case that agent k follows everyone. Consequently, G must be a clique, as agent k is arbitrary.

4 LEARNING UNDER BUDGET CONSTRAINTS

Our campaigner's approach consists of discarding all the directed graphs that do not abide by the one-step opinion diffusion observed, allowing her to retrieve any network structure given observation

and intervention budgets polynomial in the number of agents of the network. However, we are also interested in the scenario where the campaigner has a more limited budget.

We aim to explain the campaigner’s learning pipeline using a Query-Based Learning framework. To do so, Subsection 4.1 translates the observation budget directly into the number of queries that a learner (the campaigner) can ask an oracle and discusses how to reduce the current hypothesis space by establishing each agent’s feasible influencers after one opinion diffusion step. This leads to an *a posteriori* measure of the campaigner’s learning progress, which we use in Subsection 4.2 to predict the most likely answer to a campaigner’s query and, subsequently, evaluate the expected learning gain of a query compared to its intervention cost in Subsection 4.3.

4.1 Learning from a fixed observation budget

So far, we have referred to $\mathcal{H}_t \subseteq \mathcal{H}$ as the reduced hypothesis space consisting of all the networks that satisfy the observed opinion diffusion up to time $t \leq \text{Obs}$. However, we can decouple the agents in our social network and build \mathcal{H}_t as the direct product of the feasible influencers of each agent.

Definition 7. Let $(N, \cdot, \ell)_{\mathcal{H}}$ be a family of labelled social networks and suppose the campaigner observes the opinion diffusion step ℓ^+ . Then, the set of feasible influencers of agent $i \in N$ conditioned to ℓ^+ is given by

$$F(i, \ell) := \{G_i^{-1} \subset N : G \in \mathcal{H} \text{ and } \ell^+(i) = \ell_G^+(i)\}.$$

Thus, we can use Definition 7 to characterise \mathcal{H}_t for an observed sequence of one-step opinion diffusion $\ell_1^+, \dots, \ell_t^+, t \leq \text{Obs}$, as

$$\mathcal{H}_t = \{G \in \mathcal{H} : G_i^{-1} \in F(i, \ell) \text{ for all } i \in N \text{ and } \ell = \ell_1, \dots, \ell_t\}.$$

We can use this representation of \mathcal{H}_t to retrieve how many networks the campaigner can still consider to make her prediction.

Lemma 4. Let $(N, \cdot, \ell_t)_{\mathcal{H}}$ be a family of binary-labelled social networks, where \mathcal{H} is the space of all directed graphs over N . Then, the hypothesis space of networks that satisfy the opinion diffusion ℓ_t^+ , $\mathcal{H}_t \subseteq \mathcal{H}$, is such that

$$|\mathcal{H}_t| = \prod_{i \in N} \sum_{r=0}^{|A_i|} \sum_{k \in k_r} \binom{|A_i|}{r} \binom{|D_i|}{k}, \quad (4)$$

where $A_i := \{j \in N \setminus \{i\} : \ell_t(i) = \ell_t(j)\}$, $D_i := (A_i \cup \{i\})^c$ and $k_r(i) = \{0, \dots, r\}$ if $\ell_t^+(i) = \ell_t(i)$; or $k_r(i) = \{r+1, \dots, |D_i|\}$ if not.

PROOF. We build (4) by counting how many networks would admit the opinion diffusion ℓ_t^+ . Each term in the product counts the feasible influencer sets of a given agent $i \in N$. Furthermore, each term in the double sum counts the number of possible edge configurations where agent i has exactly r influencers who agree with it and k influencers who disagree. Notice that k_r is chosen to be less or equal to r if the agent does not change its opinion and strictly greater otherwise, to attain the opinion diffusion mechanism. Yet, because the campaigner does not know how many influencers agents have, she needs to consider all possibilities, that is, agent i having from 0 to $|A_i|$ influencers who agree with it and capping the possible disagreeing influencers with $|D_i|$. \square

Recall from Definition 4 that the information content gained for an additional observation is given by the ratio of the probabilities of choosing a graph in two nested hypothesis spaces. In particular, when $X \sim \text{Uniform}(\mathcal{H})$, the information gained from observing the opinion diffusion step at time $t+1$ is $\log_2(|\mathcal{H}_t|/|\mathcal{H}_{t+1}|)$.

Generalising this for any distribution \mathcal{D} over \mathcal{H} is achieved simply by taking the weighted sum of the networks encoded in (4).

4.2 Most likely query answers

A related task to estimating the underlying network is predicting the oracle’s response to an opinion diffusion query. Intuitively, the more certain the campaigner is about the underlying network, the more accurate her prediction for the subsequent labelling will be. Moreover, in the case of synchronous majority, opinion transitions rely solely on an agent’s influencers. Therefore, we will calculate the probability of the next network labelling based on feasible influencer sets as discussed in Subsection 4.1.

Definition 8. Let ℓ_1, \dots, ℓ_t , be a sequence of labellings over a social network (N, G) for some $G \in \mathcal{H}$. Then, the probability that agent i does not change its opinion, after observing the labelling $\ell \in L^N$, conditioned on the observations of $\ell_1^+, \dots, \ell_t^+$, is given by

$$\mathbb{P}(\ell^+(i) = \ell(i) \mid \ell_1, \dots, \ell_t) = \frac{|\{G_i^{-1} \in \bigcap_{k=1}^t F(i, \ell_k) : \langle \mathbf{a}_i, \mathbf{g}_i \rangle \geq 0\}|}{|\bigcap_{k=1}^t F(i, \ell_k)|}, \quad (5)$$

where \mathbf{g}_i and \mathbf{a}_i are, respectively, the vector representations of the influencer set G_i^{-1} and the agents that agree with agent i in ℓ , i.e.

$$\mathbf{g}_{ij} = \begin{cases} 1 & \text{if } j \in G_i^{-1}, \\ 0 & \text{if } j \notin G_i^{-1}, \end{cases} \text{ and } \mathbf{a}_{ij} = \begin{cases} 1 & \text{if } \ell(j) = \ell(i), i \neq j, \\ 0 & \text{if } j = i, \\ -1 & \text{if } \ell(j) \neq \ell(i), i \neq j. \end{cases} \quad (6)$$

Naturally, we would also like to know how the hypothesis space shrinks based on how similar are the labels of the opinion diffusion steps observed. This means, how can we manipulate the limits of the double sums in (4) to find which terms are present in more than one network labelling. In the following Lemma 5, we show how to do this for two adjacent network labellings.

Lemma 5. Let $(N, \cdot, \ell)_{\mathcal{H}}$, $\ell \in \{\ell_1, \ell_2\}$, be two families of binary-labelled social networks where ℓ_1 and ℓ_2 differ only on agent $i \in N$. Then, the number of feasible influencer sets of an agent $j \in N \setminus \{i\}$, such that $\ell_1(i) = \ell_1(j)$, $\ell_1^+(j) = \ell_1(j)$ and $\ell_2^+(j) = \ell_2(j)^c$, is

$$|F(j, \ell_1) \cap F(j, \ell_2)| = \sum_{r=0}^{|A_j^*|} \binom{|A_j^*|}{r} \binom{|D_j^*|+1}{r+1}, \quad (7)$$

where $A_j^* := \{k \in N \setminus \{i, j\} : \ell(k) = \ell(j), \ell \in \{\ell_1, \ell_2\}\}$ and $D_j^* := (A_j^* \cup \{i, j\})^c$. Moreover, there is no graph in $G \in \mathcal{H}$ such that $\ell_1(i) = \ell_1(j)$, $\ell_1^+(j) = \ell_1(j)^c$ and $\ell_2^+(j) = \ell_1(j)$.

To obtain (7) we isolate agent i from the set of agents who agree and disagree with agent j in (4). Note that $A_j^* = A_j(\ell) \setminus \{i\}$ and $D_j^* = D_j(\ell) \setminus \{i\}$, for $\ell \in \{\ell_1, \ell_2\}$. In this way, we can break the sums distinguishing the terms where $i \in G_j^{-1}$ and where $i \notin G_j^{-1}$. Then, to establish the elements of $F(j, \ell_1) \cap F(j, \ell_2)$ becomes a matter of

replacing the limits in k_r according to the opinion diffusion ℓ_1^+ and ℓ_2^+ , and finding where they match. For the first scenario, we obtain

$$|F(j, \ell_1) \cap F(j, \ell_2)| = \sum_{r=0}^{|A_j^*|} \underbrace{\binom{|A_j^*|}{r} \binom{|D_j^*|}{r}}_{|G_j^{-1}| \text{ is even}} + \underbrace{\binom{|A_j^*|}{r} \binom{|D_j^*|}{r+1}}_{|G_j^{-1}| \text{ is odd}},$$

which is equivalent to (7) by applying Pascal’s Identity. Following the same logic for the second scenario leads to $|F(j, \ell_1) \cap F(j, \ell_2)| = 0$, which means that there are no sets of influencers that could simultaneously change agent j ’s opinion after ℓ_1 but not after ℓ_2 .

Remark 4. *Lemma 5 provides an alternative stronger and combinatorial proof for Lemma 1. It shows that not all opinion diffusion steps are permissible under the synchronous majority update rule. In particular, we have that if agent $j \in N$ changes its opinion after $\ell_1 \in L^N$, then it will also change its opinion in any adjacent network that differs on an agent that agreed with it in ℓ_1 .*

4.3 Information Gain versus Intervention Cost

In our setup, the cost incurred by the campaigner for placing a specific query is the number of agents intervened on. But how much information can we gain with some extra cost? To this end, we formulate the notion of expected information gain.

Theorem 2. *Let ℓ_1, \dots, ℓ_t , be a sequence of labellings over a social network (N, G) for some $G \in \mathcal{H}$. If every graph in \mathcal{H} is equally likely, then the expected information gained from observing the opinion diffusion of the labelling $\ell \in L^N$ is*

$$\sum_{s \in L^N} \log_2 \left(\frac{|\mathcal{H}_t|}{|\mathcal{H}_t \cap \mathcal{H}_s^+|} \right) \prod_{i \in N} \mathbb{P}(\ell^+(i) = s(i) \mid \ell_1, \dots, \ell_t), \quad (8)$$

where $\mathcal{H}_t := \mathcal{H}_1 \cap \dots \cap \mathcal{H}_t$ is the hypothesis space consistent with the observation up to time t , \mathcal{H}_s^+ is the set of networks consistent with the opinion diffusion step $\ell^+(i) = s(i), \forall i \in N$, and we calculate $\mathbb{P}(\ell^+(i) = s(i) \mid \ell_1, \dots, \ell_t)$ using (5), for every $i \in N$.²

PROOF. In (8), $\log_2 (|\mathcal{H}_t|/|\mathcal{H}_t \cap \mathcal{H}_s^+|)$ is the information gained, as per Definition 4, when the opinion diffusion step ℓ^+ results in the network labelling s , given that all networks in \mathcal{H} are uniformly distributed. Moreover, we take the product over the conditional probabilities $\mathbb{P}(\ell^+(i) = s(i) \mid \ell_1, \dots, \ell_t)$ to obtain the joint probability of observing the transition from ℓ to s , since the agents in N change (or do not change) their opinion independently from each other. Finally, summing over all the possible subsequent network labellings $s \in L^N$ yields the desired conditional expectation. \square

Theorem 2 allows the campaigner to quantify the expected information gained from a specific query. Thus, a direct greedy approach to identify the optimal next query would be to calculate the expected information gained from a pool of possible next queries, while considering the intervention cost incurred, given by the Hamming distance between the current opinion state and the desired query. However, how to efficiently explore the exponential number of potential next queries remains as an open question.

² We choose between $\mathbb{P}(\ell^+(i) = \ell(i) \mid \ell_1, \dots, \ell_t)$ or its complement depending on whether $\ell(i) = s(i)$ or $\ell(i) = s(i)^c$.

5 DISCUSSION

We study the problem of a campaigner with budget constraints aiming to learn a social network structure with underlying synchronous majority dynamics. We provide upper bounds on the campaigner’s observation and intervention budgets for her to identify the exact social network *without any information on the network class*. We also improve these bounds for the case of clique classes, arguably the most difficult structures to alter by intervention. The study of other classes of networks may reveal further subtleties for the design of the campaigner’s strategy. While it is not hard to see that classes such as trees and simple cycles are no more difficult to learn than even cliques, what happens with other structural information about the network is far from clear.

When proving our results we observed that fast convergence does not aid the campaigner’s discovery. In the case of odd cliques, for example, one-step convergence meant that campaigner had to use $n/2 - 1$ intervention budget to “reset” the network state to maximise the information gain. Furthermore, when addressing the campaigner’s learning task with a fixed budget, we introduced a combinatorial representation of potential networks through sets of feasible influencers, simplifying the exploration of the hypothesis space from super-exponential to exponential. So, although the results are generally very positive from the campaigner’s point of view, this relies on her having sufficient budget for exact learning. Not having this severely restricts the learning task and improves the network resilience. What is more, in our analysis, we took a worst-case approach and allowed the campaigner to intervene on any node and observe the entire network. These assumptions are rarely met in practice and real-world networks are, arguably, much less easy to access, manipulate and observe.

Another important avenue for future work is going beyond majority dynamics. Each threshold-based update will display different behaviours - think of unanimous updates for instance - and it is important to understand whether fixed given thresholds can be dealt with by directly adapting our current results.

When rules are not deterministic, such as on an independent cascade model [16], our constructions for exact learning no longer work. In these contexts reasoning about information gain is still possible, although frameworks such as PAC learning [15] could provide a more suitable toolbox. We could investigate the approximate estimation of a network to relax the exact learning condition. For example, we could say that G is ε -learnable for $\varepsilon > 0$, if there exists some t_ε such that $\mathbb{P}(X = G \mid \mathcal{H}_t) > 1 - \varepsilon$ for all $t \geq t_\varepsilon$, to include PAC learning requirements in a generalised version of our framework.

ACKNOWLEDGMENTS

DC is supported in part by the Engineering and Physical Sciences Research Council [EP/X03027X/1]. LE acknowledges the support of the Centre for Discrete Mathematics and its Applications (DIMAP) at the University of Warwick and of the Engineering and Physical Sciences Research Council through the Mathematics of Systems II Centre for Doctoral Training at the University of Warwick [EP/S022244/1]. PT acknowledges the support of the Leverhulme Trust for the Research Grant RPG-2023-050 and the TAILOR Connectivity Fund (Agreement 29).

REFERENCES

- [1] Vincenzo Auletta, Ioannis Caragiannis, Diodato Ferraioli, Clemente Galdi, and Giuseppe Persiano. 2015. Minority Becomes Majority in Social Networks. In *Web and Internet Economics*, Evangelos Markakis and Guido Schäfer (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 74–88.
- [2] Vincenzo Auletta, Diodato Ferraioli, and Gianluigi Greco. 2019. On the Complexity of Opinion Consensus under Majority Dynamics. In *Proceedings of the 20th Italian Conference on Theoretical Computer Science, ICTCS September 9–11, 2019 (CEUR Workshop Proceedings, Vol. 2504)*, Alessandra Cherubini, Nicoletta Sabadini, and Simone Tini (Eds.). CEUR-WS.org, Como, Italy, 104–109.
- [3] Vincenzo Auletta, Diodato Ferraioli, and Gianluigi Greco. 2020. On the complexity of reasoning about opinion diffusion under majority dynamics. *Artificial Intelligence* 284 (2020), 103288. <https://doi.org/10.1016/j.artint.2020.103288>
- [4] Jacques Bara, Omer Lev, and Paolo Turrini. 2022. Predicting voting outcomes in the presence of communities, echo chambers and multiple parties. *Artif. Intell.* 312 (2022), 103773. <https://doi.org/10.1016/j.artint.2022.103773>
- [5] Robert Brederick and Edith Elkind. 2017. Manipulating Opinion Diffusion in Social Networks. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*. AAAI Press, Melbourne, Australia, 894–900. <https://doi.org/10.24963/ijcai.2017/124>
- [6] Markus Brill, Edith Elkind, Ulle Endriss, and Umberto Grandi. 2016. Pairwise Diffusion of Preference Rankings in Social Networks. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 9–15 July 2016*, Subbarao Kambhampati (Ed.). IJCAI/AAAI Press, New York, NY, USA, 130–136. <http://www.ijcai.org/Abstract/16/026>
- [7] Matteo Castiglioni, Diodato Ferraioli, Giulia Landriani, and Nicola Gatti. 2019. Election Manipulation on Social Networks with Messages on Multiple Candidates. *CoRR* abs/1902.03779 (2019), 10. [arXiv:1902.03779](http://arxiv.org/abs/1902.03779) <http://arxiv.org/abs/1902.03779>
- [8] Wei Chen, Yajun Wang, and Siyu Yang. 2009. Efficient Influence Maximization in Social Networks. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (Paris, France) (KDD '09)*. Association for Computing Machinery, New York, NY, USA, 199–208. <https://doi.org/10.1145/1557019.1557047>
- [9] Dmitry Chistikov, Grzegorz Lisowski, Mike Paterson, and Paolo Turrini. 2020. Convergence of Opinion Diffusion is PSPACE-Complete. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 05 (Apr. 2020), 7103–7110. <https://doi.org/10.1609/aaai.v34i05.6197>
- [10] Argyrios Deligkas, Eduard Eiben, Tiger-Lily Goldsmith, and George Skretas. 2023. Being an Influencer is Hard: The Complexity of Influence Maximization in Temporal Graphs with a Fixed Source. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems (London, United Kingdom) (AAMAS '23)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2222–2230.
- [11] David Easley and Jon Kleinberg. 2010. *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*. Cambridge University Press, Cambridge, United Kingdom. <https://doi.org/10.1017/CBO9780511761942>
- [12] Manuel Gomez-Rodriguez, David Balduzzi, and Bernhard Scholkopf. 2011. Uncovering the Temporal Dynamics of Diffusion Networks. In *Proceedings of the 28th International Conference on International Conference on Machine Learning (Bellevue, Washington, USA) (ICML '11, Vol. 28)*. Omnipress, Madison, WI, USA, 561–568.
- [13] Manuel Gomez-Rodriguez, Jure Leskovec, and Andreas Krause. 2010. Inferring networks of diffusion and influence. *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining 5* (2010), 1–37. Issue 4. <https://api.semanticscholar.org/CorpusID:2327990>
- [14] Michel Grabish and Agnieszka Rusinowska. 2010. A model of influence in a social network. *Theory and Decisions* 69 (2010), 69–96.
- [15] Michael J. Kearns and Umesh V. Vazirani. 1994. *An Introduction to Computational Learning Theory*. MIT Press, Cambridge, MA, USA.
- [16] 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 (Washington, D.C.) (KDD '03)*. Association for Computing Machinery, New York, NY, USA, 137–146. <https://doi.org/10.1145/956750.956769>
- [17] David Kempe, Jon M. Kleinberg, and Éva Tardos. 2005. Influential Nodes in a Diffusion Model for Social Networks. In *Automata, Languages and Programming, 32nd International Colloquium, ICALP July 11–15, 2005, Proceedings (Lecture Notes in Computer Science, Vol. 3580)*, Luis Caires, Giuseppe F. Italiano, Luis Monteiro, Catuscia Palamidessi, and Moti Yung (Eds.). Springer, Lisbon, Portugal, 1127–1138. https://doi.org/10.1007/11523468_91
- [18] Fang Kong, Jize Xie, Baoxiang Wang, Tao Yao, and Shuai Li. 2023. Online Influence Maximization under Decreasing Cascade Model. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems (London, United Kingdom) (AAMAS '23)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2197–2204.
- [19] Siyu Lei, Silviu Maniu, Luyi Mo, Reynold Cheng, and Pierre Senellart. 2015. Online Influence Maximization. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (Sydney, NSW, Australia) (KDD '15)*. Association for Computing Machinery, New York, NY, USA, 645–654. <https://doi.org/10.1145/2783258.2783271>
- [20] Jure Leskovec, Andreas Krause, Carlos Guestrin, Christos Faloutsos, Jeanne VanBriesen, and Natalie Glance. 2007. Cost-Effective Outbreak Detection in Networks. In *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (San Jose, California, USA) (KDD '07)*. Association for Computing Machinery, New York, NY, USA, 420–429. <https://doi.org/10.1145/1281192.1281239>
- [21] David J. C. MacKay. 2003. *Information Theory, Inference, and Learning Algorithms*. Copyright Cambridge University Press, Cambridge, United Kingdom.
- [22] Seth A. Myers and Jure Leskovec. 2010. On the Convexity of Latent Social Network Inference. In *NIPS. Curran Associates Inc.* 57 Morehouse Lane Red Hook NY United States, Vancouver, British Columbia, Canada, 1741–1749. <https://api.semanticscholar.org/CorpusID:6812407>
- [23] Harikrishna Narasimhan, David C Parkes, and Yaron Singer. 2015. Learnability of Influence in Networks. In *Advances in Neural Information Processing Systems*, C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett (Eds.), Vol. 28. Curran Associates, Inc., Montreal, Canada. https://proceedings.neurips.cc/paper_files/paper/2015/file/4a2ddf148c5a9c4215a529e8cbdcc06-Paper.pdf
- [24] Praneeth Netrapalli and Sujay Sanghavi. 2012. Learning the Graph of Epidemic Cascades. In *Proceedings of the 12th ACM SIGMETRICS/PERFORMANCE Joint International Conference on Measurement and Modeling of Computer Systems (London, England, UK) (SIGMETRICS '12)*. Association for Computing Machinery, New York, NY, USA, 211–222. <https://doi.org/10.1145/2254756.2254783>
- [25] Claude Sammut and Geoffrey I. Webb (Eds.). 2017. *Encyclopedia of Machine Learning and Data Mining* (2 ed.). Springer, New York. <https://doi.org/10.1007/978-1-4899-7687-1>
- [26] Youze Tang, Yanchen Shi, and Xiaokui Xiao. 2015. Influence Maximization in Near-Linear Time: A Martingale Approach. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data (Melbourne, Victoria, Australia) (SIGMOD '15)*. Association for Computing Machinery, New York, NY, USA, 1539–1554. <https://doi.org/10.1145/2723372.2723734>
- [27] Yiyi Zhao, Gang Kou, Yi Peng, and Yang Chen. 2018. Understanding influence power of opinion leaders in e-commerce networks: An opinion dynamics theory perspective. *Information Sciences* 426 (2018), 131–147. <https://doi.org/10.1016/j.ins.2017.10.031>