Maximizing the Spread of an Opinion when Tertium Datur Est

Vincenzo Auletta
Università degli Studi di Salerno, Fisciano (SA), Italy
auletta@unisa.it

Diodato Ferraioli
Università degli Studi di Salerno, Fisciano (SA), Italy
dferraioli@unisa.it

Valeria Fionda
Università della Calabria, Rende (CS), Italy
valeria.fionda@unical.it

Gianluigi Greco
Università della Calabria, Rende (CS), Italy
gianluigi.greco@unical.it

ABSTRACT

Opinion diffusion has been largely studied in the literature on settings where the opinion whose spread has to be maximized, say *white*, competes against one opinion only, say *black*. For instance, for diffusion mechanisms modeled in terms of best response dynamics over *majority* agents (who change their opinion as to conform it to the majority of their neighbors), it is known that the spread can be maximized via certain greedy dynamics that can be computed in polynomial time. This setting is precisely the one considered in the paper. However, differently from earlier literature, it is assumed that one further opinion, say *gray*, is available to the agents. Moving from the observation that, with the third alternative to hand, greedy dynamics can dramatically fail to maximize the spread of opinion *white*, the paper then embarks in thorough computational, algorithmic and experimental studies.

The picture that emerges is totally different from what is known for the case when two opinions are available only.

Indeed, it is shown that greedy dynamics can dramatically fail in maximizing the spread. In particular, deciding whether there exists a dynamics that can spread the opinion *white* to at least k agents or can reach a consensus is shown to be intractable, formally NP-hard. On the other hand, islands of tractability based on certain structural properties of the interaction graph are identified. Finally, experimental results are discussed, which shed lights on opinion diffusion in real social networks.

The final version of this paper is available at https://dl.acm.org/doi/pdf/10.5555/3306127.3331823.

1 INTRODUCTION

Reasoning about social networks is a topic that has been attracting attention in the research literature since the late seventies, when a number of models have been proposed to understand how opinions can form and diffuse because of the existence of social ties among the individuals [19, 27, 33]. By abstracting from their peculiarities, these models as well as most of their subsequent refinements [20, 30] propose high-level formalizations of the social environments, where the mental state of each individual is encoded as a variable taking either binary values, say *black* and *white*, or real values in the interval $[0, 1]$ providing a scale (e.g., probabilistic) for defining intermediate mental states between the two extremes.

In fact, enhancing the expressiveness of such models for opinion diffusion is an ambitious goal of the artificial intelligence community, where logic-based frameworks are increasingly adopted to reason about the diffusion of complex opinions, such as multiple binary issues, linear orders, or belief bases. Logic formalizations currently include models based on (variants of) *propositional logic* [1, 4, 13, 16, 21, 26], *logic programming* [40, 41], *epistemic logic* [8, 37, 39], *fuzzy logic* [29, 42], and *dynamic modal logic* [18, 35], just to name a few.

In this paper, we continue along the line of research of enhancing the expressiveness of some basic diffusion models, while taking however a minimalistic perspective: Rather than using complex theories to model the available opinions and the mental states of the individuals, we just assume that one further opinion, say *gray*, is available to them in addition to *white* and *black*.[1] On the one hand, the environments we deal with are meaningful in their own; for instance, with our modeling approach, one can reason about opinion diffusion in political elections, by focusing on the (social-media) impact of some emerging party in addition to the more traditional Left and Right ones. On the other hand, the negative results (e.g., computational intractability) we shall derive in our setting immediately apply to more involved models and shed lights on their core properties.

So, in a nutshell, the paper addresses the following question: *What happens to the dynamics of social environments when three opinions are available to the individuals?*

We will provide from analytic, computational, and experimental viewpoints an answer to this question for the a specific "game-theoretic" setting [3, 6, 7, 12, 15, 17, 25], where individuals are rational agents that *asynchronously* change their current opinion—modeled as a discrete variable with three possible values—as to conform it to the *majority* of their neighbors in the network.

The picture that emerges is totally different from what is known for the case when two opinions are available only:

- First, we show that dynamics where we greedily prefer to change to *white* the opinion of the agents do not necessarily maximize the spread of that opinion. Indeed, differently from the two-opinions scenario [12], optimal propagations can require interleaving the diffusion of *white* with the diffusion of one of the two other opinions. Moreover, optimal dynamics over a set N of agents can require that some agent changes

[1] Note that in this work assume *gray* to be a completely different opinion with respect to *white* and *black*, and it is not in the "middle" of *white* and *black* (e.g., it does not coincide with $1/2$ in the setting where opinions range in the interval $[0, 1]$). The choice of using *gray* (in place of a different color) is motivated by coloring of our figures.

her opinion $O(|N|)$ times, which is rather uprising as in the two-opinions setting one can focus, w.l.o.g., on dynamics where each agent changes her opinion at most twice.
- Second, we show that — as a consequence of the "non-monotone" behavior of optimal dynamics — the problem of maximizing the spread of *white* is intractable, formally NP-hard, even if we just ask for deciding whether consensus on *white* can be reached among all the agents.
- Motivated by the above bad news, we then look for islands of tractability based on some structural properties of the networks as well as on fixing a constant bound on the maximum number of times each agent can change her opinion.[2] Notably, tractability is shown to hold not only over majority agents, but also on more general *threshold-majority* ones, i.e., very roughly, on agents that are influenced by their neighbors only if the number of them that already agreed on some opinion is above some given threshold.
- Even if an analytical description of the optimal dynamics is impossible due to our complexity results, it would still be interesting to gain insight on the best way to maximize the spread of *white* on a network. For this reason, we analyze dynamics with three opinions from the experimental/practical viewpoint. In particular, our focus is to understand how the initial distribution of remaining opinions influences the number of agents with opinion *white* at the stable state. We indeed show that the number of agents to which opinion *white* can spread significantly grows when the initial fractions of agents that hold opinion *black* and *gray* are almost equal. This kind of *entropic gain* emerges on a number of real-world networks, for instance on FaceBook.

We stress that most of our results can be extended to more than three opinions. Clearly, this is the case for negative results. As for tractable cases, it turns out that our proof can be generalized to any *fixed* number of opinions with only minor modifications.

2 FORMAL FRAMEWORK

In this section, we introduce the basic notions and notations we shall use throughout the rest of the paper.

Social Environments. Let $\mathcal{N}$ be a given set of agents. A *social environment* S is defined as a tuple $\langle O, N, E, \{t_x\}_{x\in N}\rangle$ where: $O \subseteq \{white, black, gray\}$ is a set of opinions; (N, E) is an undirected graph, shortly denoted by $G(S)$ and encoding the interactions of a set $N \subseteq \mathcal{N}$ of agents; and, for each agent $x \in N$, t_x is a *threshold function* mapping every opinion $op \in O$ to the number $t_x(op) \in \{0, \ldots, |\delta(x)|\}$, where $\delta(x)$ is the set $\{y \mid \{x, y\} \in E\}$ of the *neighbors* of x in S.

An *opinion profile* o for S is a function $o : N \to O$ mapping agents to opinions. Agents express cardinal preferences over opinion profiles as follows. For any set $C \subseteq N$ of agents and every $op \in O$, let $C_{op/o}$ (shortly C_{op}, if the profile o is clearly understood) be the set of all agents of C with opinion op in the profile o. Then, the *utility function* u_x of an agent $x \in N$ is such that, for each profile o_{-x} of agents different from x, and for every $op \in O$,

$$u_x(op, o_{-x}) = \begin{cases} 0, & \text{if } |\delta(x)_{op}| \leq t_x(op), \\ |\delta(x)_{op}| & \text{otherwise,} \end{cases}$$

where $\delta(x)_{op}$ denotes the number of neighbors of x holding opinion op, i.e., $\delta(x)_{op} = \delta(x) \cap C_{op/o}$.

Agents in S are called *threshold-majority.* Roughly speaking, a threshold majority agent keeps a threshold for every opinion, and always prefers to adopt the opinion that is supported by the majority of neighbors among the ones that pass the corresponding thresholds. If no opinion passes the threshold, we assume that the agent prefers to keep her current opinion. In general, we assume that ties that involve the current opinion are always broken in favor of this opinion. No tie-breaking rule is defined for the case that ties do not involve the current opinion.

Clearly, threshold-majority agents extend standard *majority* ones, that simply take the opinion supported by most of the neighbors: indeed, a majority agent x is such that $t_x(op) = 0$ for each $op \in O$. We stress that all our negative results will hold not only for threshold-majoirty agents, but also for the simplex setting with majority agents.

Moreover, it is immediate to see that our definition also extends a natural generalization of the classic threshold model [27, 31, 38] to more than two spreading opinions, named *separated threshold model* [11], the main difference being that in the latter model agents are assumed to change their mind only once, whereas in our model the number of times agents are allowed to update their opinion is in principle unlimited.

Finally, our definition also encompasses and generalizes the *discrete preference games* model of discrete opinion dynamics, as defined in [15, 22]: here, agents are provided with private beliefs and opinions are chosen in order to minimize both the distance with neighbors' opinions and the distance with their own beliefs. This behavior can be immediately in our framework, by choosing thresholds opportunely based on the belief.

Dynamics and Convergence. An agent $x \in N$ is *stable* in o if $u_x(o') \leq u_x(o)$ holds for each profile o′ such that $o'(y) = o(y)$, for each $y \in N \setminus \{x\}$. The profile o is stable if all agents are stable in it. We consider an asynchronous model where, at each time instant, precisely one agent that is not stable changes her opinion. In particular, a *dynamics* for S is a sequence of profiles $o_0, \ldots, o_k$ such that o_{h+1}, for each $h \in \{0, \ldots, k-1\}$, is obtained from o_h by setting the opinion of an agent x that is not stable in o_h to an opinion that makes x stable in o_{h+1} — that is, we consider a *best response* dynamics, where at each time step the selected agent adopts the opinion that miximizes her utility. The dynamics is said to be *converging* if o_k is stable. Moreover, it will be shortly denoted as $o_0 \rightsquigarrow o_k$, if we are interested in the initial and final profiles only. Note that converging dynamics are guaranteed to exist for social environments with majority-threshold agents.

THEOREM 2.1. *Let $S = \langle O, N, E, \{t_x\}_{x\in N}\rangle$ be a social environment. Then, for each profile o for S, there is a converging dynamics starting at o and whose length is at most $3|E|$.*

PROOF. Let $S = \langle O, N, E, \{t_x\}_{x\in N}\rangle$ be a social environment and, for each profile o, define $\Phi(o) = |\{(x, y) \in E \mid o(x) = o(y)\}| -$

[2] The latter requirement naturally emerges in real-world scenarios, where individuals tend to be reluctant to change their opinion, after they have already publicly expressed one.

$\sum_{x \in N} t_x(\mathrm{o}(x))$. Observe that $-2|E| = -\sum_{x \in N} |\delta(x)| \leq \Phi(\mathrm{o}) \leq |E|$. Consider any dynamics having the form $\mathrm{o}_0, ..., \mathrm{o}_k$ and note that, for each $h \in \{0, ..., k-1\}$, $\Phi(\mathrm{o}_{h+1}) - \Phi(\mathrm{o}_h) \geq 1$ holds. Indeed, if $x \in N$ is the agent that changes her opinion in o_h, then we know that

$$\Phi(\mathrm{o}_{h+1}) - \Phi(\mathrm{o}_h) = \Big(|\delta(x)_{\mathrm{o}_{h+1}(x)/\mathrm{o}_{h+1}}| - t_x(\mathrm{o}_{h+1}(x))\Big) - \Big(|\delta(x)_{\mathrm{o}_h(x)/\mathrm{o}_h}| - t_x(\mathrm{o}_h(x))\Big) \geq 1,$$

where the last inequality follows by our definitions of stable and threshold-majority agents. Hence, Φ is a potential function that monotonically decreases at each step. The maximum number of steps is bounded by the difference between the maximum and the minimum value that Φ can take. □

Since we are interested in maximizing the spread of opinion *white*, we denote by $\max_\varrho(\mathrm{o}, S)$ the maximum number of agents that hold opinion *white* over all profiles o^* such that $\mathrm{o} \rightsquigarrow \mathrm{o}^*$ is a dynamics where each agent changes her opinion at most ϱ times. If we omit the subscript and write $\max(\mathrm{o}, S)$, then we just mean that no bound is considered on the number of times each agent can change her opinion.

3 ANALYSIS OF DYNAMICS OVER THREE OPINIONS

In this section we analyze the behavior of the dynamics maximizing the spread of opinion *white* in social environments with three opinions and we contrast such behavior with that characterizing environments with two opinions only.

3.1 Greedy Dynamics

A dynamics $\mathrm{o}_0, ..., \mathrm{o}_k$ is *greedy* if for each $h \in \{0, ..., k-1\}$, whenever there is an agent $x \in N$ that is not stable in o_h and that can become stable by taking opinion *white*, then $|N_{white/\mathrm{o}_{h+1}}| > |N_{white/\mathrm{o}_h}|$ holds. For an environment S, let us then define the *Price of Greediness* $\mathrm{PoG}(S)$ as the ratio between $\max_{\mathrm{o}}(\mathrm{o}, S)$ and the maximum number of agents that hold opinion *white* over all stable profiles o^* that are reachable from o via some greedy dynamics.

For environments S defined over majority agents (but the same is easily seen to hold on threshold-majority ones) and with two alternative opinions, it is known that $\mathrm{PoG}(S) = 1$ always holds [12]. Instead, for environments S where one further opinion is allowed, we next show that $\mathrm{PoG}(S)$ can grow arbitrarily. The result is somehow surprising and holds even on majority agents.

Theorem 3.1. *There is a class of environments $\{S_n\}_{n>0}$ over majority agents such that* $\lim_{n\to\infty} \mathrm{PoG}(S_n) = \infty$.

Proof. Let $S_n = \langle O, N, E, \{t_x\}_{x \in N} \rangle$ to be a social environment where $O = \{white, black, gray\}$ and $\mathrm{G}(S_n)$ is the graph depicted in Figure 1.(a) with $N = \{a, b, c, d, e, f\} \cup \{h_1, ..., h_n\}$. Let o be the initial configuration where each agent holds the opinion corresponding to the coloring (either *white*, or *black*, or *gray*) of the nodes in the figure. Assume that all agents are majority ones.

Consider the following dynamics $\mathrm{o} \rightsquigarrow \mathrm{o}^*$, illustrated in Figure 1.(b): First, c adopts opinion *gray* (hence the dynamics is not greedy); then, all agents adopt opinion *white* in the following order: $b, c, d, h_1, ..., h_{n-1}, h_n$. Note that this dynamics is feasible for majority agents. In fact, every time an agent changes her opinion, she leaves an opinion that is not supported by the majority of her neighbors in favor of an opinion supported by the majority of her neighbors (with ties being resolved precisely in favor of *white*). Eventually, an opinion profile is reached where all agents hold opinion *white*. Hence, we have $\max(\mathrm{o}, S_n) = n + 6$.

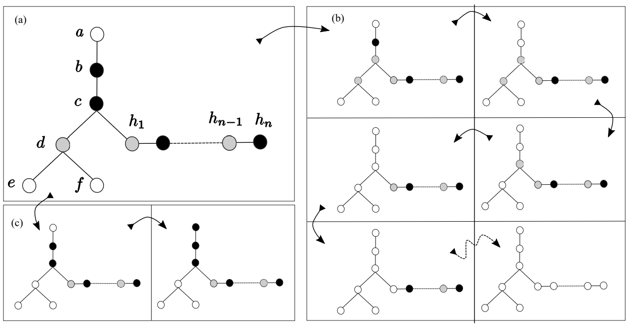

Figure 1: Environment S_n and dynamics in Theorem 3.1.

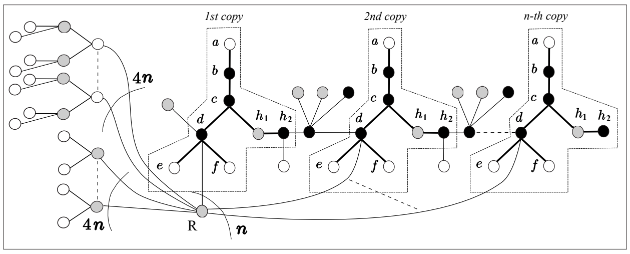

Figure 2: Environment $S_n^\star$ and profile $\mathrm{o}_n^\star$ in Theorem 3.2.

On the other hand, observe that any greedy dynamics $\mathrm{o} \rightsquigarrow \tilde{\mathrm{o}}$ must start with agent d changing her opinion to *white*, since this is the only agent interested in adopting opinion *white*. Let o' be the configuration after this change, which is graphically illustrated in Figure 1.(c). From this configuration, no other agent is interested in changing her opinion to *white*, whatever greedy dynamics we consider. Rather, agent a will eventually adopt opinion *black* (possibly after that some of the agents in $\{h_1, ..., h_n\}$ changed her opinion to *black* or *gray*. Hence, we have $\widetilde{\max}(\mathrm{o}, S_n) = 3$ and the result clearly follows. □

3.2 Dynamics with Bounds on the Number of Changes

Another appealing property of environments with two opinions is that, to maximize the spread of *white*, one can focus, w.l.o.g., on (greedy) dynamics where each agent changes her opinion at most twice [12]; that is, $\max(\mathrm{o}, S) = \max_2(\mathrm{o}, S)$ holds, for each initial profile o and environment S over two opinions. So, one might be tempted to guess that, when moving to environments with three opinions, changing the opinion of each agent at most three times might suffice. We next show that this is not the case.

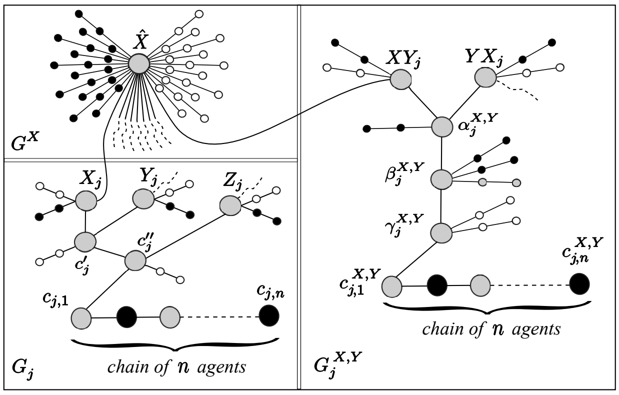

Figure 3: Gadgets exploited in the proof of Theorem 4.1.

Theorem 3.2. *There is a class of pairs of profiles and environments* $\{o_n^\star, S_n^\star\}_{n>0}$ *over majority agents where, denoting by* ϱ_n *the minimum value for which* $\max_{\varrho_n}(o_n^\star, S_n^\star) = \max(o_n^\star, S_n^\star)$, *it holds that* $\lim_{n\to\infty} \varrho_n = \infty$.

Proof Sketch. Consider the environment S_2 in the proof of Theorem 3.1, built on agents a, b, c, d, e, f, h_1, and h_2.

We build $S_n^\star$ by arranging n distinct copies of S_2 as illustrated in Figure 2, and we define $o_n^\star$ correspondingly to the coloring of the nodes—note that all copies of agent d are *black*, rather than *gray* (as in Figure 1). Moreover, all copies of agent d are connected to the novel node R, which has a total of $9n$ connections. In this construction, the crucial property is that, to maximize the spread of opinion *white* in the i-th copy of S_2, for each $i \in \{1, .., n\}$, agent d must first change her opinion to *gray*—for propagating this opinion to c—and then change it again to *white*. But, the first flip requires that R holds opinion *gray* and the whole (i-1)-th copy of S_2 has already converged to *white*, so that *gray* can be propagated via the gadget connecting the (i-1)-th copy of h_2 with the i-th copy of d. The subsequent flip (to *white*) requires instead that R holds opinion *white*.

Hence, to orchestrate the required changes on the various copies of d, agent R has to change her opinion $2n - 1$ times, which is possible due to her further $4n + 4n$ connections. The result follows since dynamics where R changes her opinion less than $2n - 1$ times do not lead to consensus on *white*. □

4 COMPLEXITY RESULTS: HARD AND EASY NETWORKS

In this section, we consider the problem of maximizing the spread of opinion *white* from the computational perspective. In the analysis we assume a standard encoding for the input which takes polynomial space in $|N|$, and we assume that threshold functions can be computed in polynomial time.

4.1 Hardness Results

Let Max-Spread[ϱ] be the problem that, given a social environment S, an initial configuration o, and a natural number τ, asks whether $\max_\varrho(o, S) \geq \tau$ holds. For environments over two opinions, Max-Spread[ϱ] is clearly tractable, since we have just to implement a greedy dynamics. We next show that a complexity blow-up occurs if one further opinion is allowed, even by focusing on majority agents.

Theorem 4.1. *For each natural number* $\varrho \geq 1$, Max-Spread[ϱ] *is* NP*-complete. Hardness holds even on environments defined over majority agents.*

Proof Sketch. Membership in NP follows easily after Theorem 2.1. Concerning the hardness, let $\Phi = c_1 \wedge \cdots \wedge c_m$ be a Boolean formula in conjunctive normal form such that, for each $j \in \{1, ..., m\}$, c_j is a clause consisting of the disjunction of precisely three variables (all occurring positively) taken from a set $\mathcal{V}$. The problem one-in-three positive 3-SAT of deciding whether there is a truth assignment σ such that, for each clause c_j, precisely one variable in c_j evaluates to true is known to be NP-hard [23]. W.l.o.g., we assume that each variable occurs precisely in three clauses.

Given the formula Φ, consider the environment S_Φ where all agents are majority ones and the graph $G(S_\Phi) = (N, E)$ is built as follows. For each variable $X \in \mathcal{V}$, N contains a distinguished agent $\hat{X}$ and the subgraph G^X depicted in Figure 3—showing also the initial profile o_Φ. Note that $\hat{X}$ is connected to the agents X_j and XY_j, for each clause c_j where X occurs together with Y. In their turn, these agents belong to the subgraphs G_j and $G_j^{X,Y}$. So, in addition to 10 "anonymous" *black* neighbors and 10 "anonymous" *white* neighbors, $\hat{X}$ has 9 *gray* neighbors in subgraphs having the form G_j and $G_j^{X,Y}$, and corresponding to the three clauses c_j where X occurs and to the other two variables in c_j.

Now, each subgraph involves a chain, for a total of $4m$ chains each one over n nodes. In fact, it is easy to define n to be large enough ($n = 21|\mathcal{V}| + 96m + 1$ suffices), so that $\max(o_\Phi, S_\Phi) \geq 4mn$ holds if, and only if, all the nodes in these chains will adopt opinion *white* in a converging dynamics. So, to conclude the proof, let us analyze how such converging dynamics look like. First, note that $\hat{X}$ is the only unstable agent and every dynamics must originate from $\hat{X}$ changing her opinion, with *white* (resp., *black*) meant to encode that X evaluates to true (resp., false). The adopted opinion is consistently propagated to X_j and XY_j.

Then, the subgraph G_j acts as an evaluator for the corresponding clause c_j, so that $c_{j,1}$ can eventually get opinion *white* if, and only if, one agent corresponding to a variable in c_j had previously adopted opinion *white*. Moreover, the subgraph $G_j^{X,Y}$ is such that $c_{j,1}^{X,Y}$ can eventually get opinion *white* if, and only if, at least one agent between $\hat{X}$ and $\hat{Y}$ had previously adopted opinion *black*. Eventually, when the first node in a chain adopts opinion *white*, the whole chain can adopt the same opinion, too. Therefore, $\max(o_\Phi, S_\Phi) \geq 4mn$ holds if, and only if, one-in-three positive 3-SAT has a positive answer over Φ. To conclude, note that in the dynamics every agent changes her opinion at most once. Hence, the reduction holds for each $\varrho \geq 1$. □

In the above proof, the careful reader might have noticed that we used a number of gadgets defined over ("anonymous") agents that are already stable with their initial opinion *black* or *gray* and never changes their mind. With some additional efforts, we can

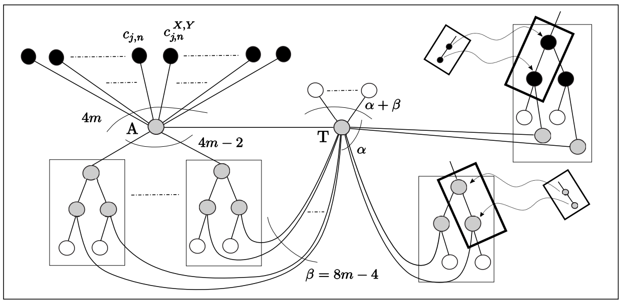

Figure 4: Construction in the proof of Theorem 4.2.

modify these gadgets and show the NP-hardness of the Complete-Spread[ϱ] problem defined as the restriction of Max-Spread[ϱ] to the instances where $\tau = |N|$ (i.e., where all agents have to adopt *white*).

Theorem 4.2. *For each $\varrho \geq 2$, Complete-Spread[ϱ] is NP-complete. Hardness holds even on environments defined over majority agents.*

Proof Sketch. We extend the reduction in the proof of Theorem 4.1 as follows. Given the Boolean formula Φ over the variables in $\mathcal{V}$, we build an environment $\hat{S}_\Phi$ by adding into S_Φ the agents reported in the subgraph illustrated in Figure 4 (as usual, the coloring in the figure reflects the opinion in an initial profile, say $\hat{o}_\Phi$). In $\hat{S}_\Phi$, all agents having the form $c_{j,k}$ or $c_{j,k}^{X,Y}$ that initially hold opinion *black* (for a total of $4m$ agents) are connected to the novel agent A, which is connected to the novel agent T and to $4m-2$ novel gadgets (from which $\beta = 8m-2$ agents with opinion *gray* are in their turn connected to T). Moreover, further connections to some novel agents are added to those agents that have been depicted in Figure 3 as small circles and are initially colored *black* or *gray* (see again Figure 4). In particular, the pairs of *black* (resp., *gray*) small circles in Figure 3 are $10|\mathcal{V}| + 18m$ (resp., $3m$), so that in the figure we have $\alpha = 2 \times (10|\mathcal{V}| + 18m + 3m)$. Eventually, note that agent T has $\alpha + \beta$ novel neighbors whose initial opinion is *white*.

Our construction does not alter the properties we have discussed in the proof of Theorem 4.1. So, it can be checked that consensus on *white* can be achieved if, and only if, agent T can change her opinion to *white* in some dynamics. But, this can happen if, and only if, agent A can previously adopt opinion *white*; hence if, and only if, all the *black* neighbors of A can previously adopt opinion *white*, too—which we know from the proof of Theorem 4.1 can happen if, and only if, one-in-three positive 3-SAT has a positive answer over Φ. Note, in particular, that such converging dynamics requires that some agent changes her opinion twice. □

Note that $\varrho \geq 2$ is crucial above, because Complete-Spread[1] can be solved via a simple polynomial-time algorithm, based on greedily spreading *white* as long as this is possible (and then checking if consensus is achieved).

4.2 Islands of Tractability

We now look for environments on which Max-Spread[ϱ] (hence, Complete-Spread[ϱ]) is tractable, by focusing on (nearly-)*acyclic* interaction graphs and imposing constant bounds on the maximum number ϱ of times each agent can change her opinion. Note that structural tractability results for diffusion-related problems have been already established in the literature, via ad-hoc algorithmic approaches [9, 12]. Our proof takes a different perspective and shows that expressibility in *Monadic Second Order (MSO) logic* is the most intimate reason for the tractability. We envisage that our method, with minor modifications, can be used to establish similar results over other problems and different models of diffusion.

Preliminaries on treewidth and MSO. A *tree decomposition* of a graph $G = (V, E)$ is a pair $\langle T, \chi \rangle$, where T is a tree and χ is a labeling function assigning to each vertex p in T a set of nodes $\chi(p) \subseteq V$, such that the following conditions are satisfied: (1) for each node $x \in V$, there exists p in T such that $x \in \chi(p)$; (2) for each edge $\{x, y\} \in E$, there exists p in T such that $\{x, y\} \subseteq \chi(p)$; and, (3) for each node $x \in V$, the subgraph of T induced by all nodes p such that $x \in \chi(p)$ is connected. The *width* of $\langle T, \chi \rangle$ is the number $\max_{p \in T}(|\chi(p)| - 1)$. The *treewidth* of G, denoted by $tw(G)$, is the minimum width over all its tree decompositions. Treewidth is known to be a true generalization of acyclicity: G is acyclic if, and only if, $tw(G) = 1$.

A *finite structure* $\mathcal{D}$ consists of a finite domain D and relations $R_1, \ldots, R_k$. Each relation R_i consists of a set of tuples $(c_1, \ldots, c_{a_i})$, where $c_j \in D$, for each $1 \leq j \leq a_i$. Monadic Second Order (MSO) enhances the expressiveness of First Order logic by allowing the use of set variables, of the membership relation $\in$, and of the quantifiers $\exists$ and $\forall$ over set variables. In addition, it is often convenient to use symbols like $\subseteq, \subset, \cap, \cup, \rightarrow$, and $\leftrightarrow$ with their usual meaning, as abbreviations. When an MSO formula ϕ is evaluated over a finite structure $\mathcal{D}$, the relation symbols of ϕ are interpreted as the corresponding relations of $\mathcal{D}$ and the variables of ϕ range over the domain D. The fact that an MSO formula ϕ holds over $\mathcal{D}$ is denoted by $\mathcal{D} \models \phi$.

In the following, if $\Phi(\mathtt{W})$ is an MSO formula where $\mathtt{W}$ is the only *free*, i.e., not quantified, (set) variable occurring in it, then for each $W \subseteq D$, we denote by $\Phi(W)$ the MSO formula, without free variables, where $\mathtt{W}$ is replaced by W.

For instance, the property that a set $\mathtt{W}$ of nodes in a graph $G = (N, E)$ — viewed as a structure over domain N and where E is encoded as the relation edge — is a *vertex cover* can be expressed via the MSO formula:

$$vc(\mathtt{W}) \equiv \forall y, (y \in \mathtt{W}) \vee (\exists x \in \mathtt{W} \wedge \mathsf{edge}(x, y)).$$

Then, $G \models vc(W)$ if, and only if, $W \subseteq N$ is a vertex cover.

MSO Encoding. We are now ready to claim that we can encode Max-Spread[ϱ] in MSO.

Theorem 4.3. *Let $\varrho, \delta > 0$ be two natural numbers. Then, there exists an MSO formula $\Phi_{\varrho,\delta}(\mathtt{W})$ and a polynomial-time computable function f that, given any environment $S = \langle O, N, E, \{t_x\}_{x \in N} \rangle$ with $\max_{x \in N} |\delta(x)| \leq \delta$ and a profile o for it, returns a finite structure $f(o, S)$ such that, for each $W \subseteq N$, the following are equivalent:*

1. *$f(o, S) \models \Phi_{\varrho,\delta}(W)$;*
2. *there is a dynamics $o \rightsquigarrow o^*$ for S where each agent changes her opinion at most ϱ times and $N_{white/o^*} = W$.*

$$\Phi_{\varrho,\delta}(\mathtt{W}): \quad \exists_{i,j}\mathtt{order}_{i,j}\ \exists_{op,i}\mathtt{holds}_{op,i}\ \ \exists_{op}\mathtt{final}_{op}\ \bigwedge_{i,j,h}\psi^h_{i,j} \wedge \bigwedge_{op,i,h}\phi^h_{op,i} \wedge \bigwedge_{op,h}\tau^h_{op} \wedge \xi(\mathtt{W})$$

$$\begin{array}{rl}
\psi^1_{i>0,j\geq 0}: & \forall w\ \mathtt{order}_{i,j}(w) \rightarrow \exists x,y\ \mathtt{from}(x,w) \wedge \mathtt{to}(y,w) \wedge \neg\bigvee_{j'\neq j}\mathtt{order}_{i,j'}(w) \wedge \bigvee_{op}\mathtt{holds}_{op,i}(x) \\
\psi^2_{1,j>0}: & \forall w,\bar{w}\ \mathtt{order}_{1,j}(w) \wedge \mathtt{inverse}(w,\bar{w}) \rightarrow \mathtt{order}_{j,0}(\bar{w}) \\
\psi^2_{i>1,0}: & \forall w,\bar{w}\ \mathtt{order}_{i,0}(w) \wedge \mathtt{inverse}(w,\bar{w}) \rightarrow \mathtt{order}_{i-1,0}(w) \\
\psi^2_{i>1,j>0}: & \forall w,\bar{w}\ \mathtt{order}_{i,j}(w) \wedge \mathtt{inverse}(w,\bar{w}) \rightarrow \neg(\mathtt{order}_{j,i-1}(\bar{w}) \leftrightarrow \mathtt{order}_{i-1,j}(w))
\end{array}$$

$$\begin{array}{rl}
\phi^1_{op,i>0}: & \forall x\ \mathtt{holds}_{op,i}(x) \rightarrow \neg\bigvee_{op'\neq op}\mathtt{holds}_{op',i}(x) \wedge \bigvee_{\hat{op}\neq op}(\mathtt{holds}_{\hat{op},i-1}(x) \wedge \mathtt{better}_{op,\hat{op},i}(x)) \\
\phi^2_{op,i>0}: & \forall x\ \mathtt{holds}_{op,i}(x) \rightarrow \neg\bigvee_{op'\neq op}\mathtt{better}_{op',op,i}(x) \\
\phi^3_{op,i>0}: & \forall w,x\ \mathtt{holds}_{op,i}(x) \wedge \mathtt{from}(x,w) \rightarrow \bigvee_j \mathtt{order}_{i,j}(w)
\end{array}$$

$$\begin{array}{rl}
\tau^1_{op}: & \forall x\ \mathtt{final}_{op}(x) \rightarrow \mathtt{holds}_{op,\varrho}(x) \vee \bigvee_{i\in\{0,\dots,\varrho-1\}}(\mathtt{holds}_{op,i}(x) \wedge \neg\bigvee_{op'}\mathtt{holds}_{op',i+1}(x)) \\
\tau^2_{op}: & \forall x\ \mathtt{final}_{op}(x) \rightarrow \neg\bigvee_{\hat{op}\neq op}(\mathtt{better}_{\hat{op},op,\infty}(x))
\end{array}$$

$$\xi(\mathtt{W}): \ \forall x\ \bigvee_{op}\mathtt{final}_{op}(x) \wedge (x \in \mathtt{W} \leftrightarrow \mathtt{final}_{white}(x))$$

$$\begin{array}{lcl}
\mathtt{inverse}(w,\bar{w}) & \equiv & \exists x,y\ \mathtt{from}(x,w) \wedge \mathtt{to}(y,w) \wedge \mathtt{to}(y,\bar{w}) \wedge \mathtt{from}(x,\bar{w}) \\
\mathtt{better}_{op',op,i}(x) & \equiv & \bigvee_{1\leq d\leq\delta,\ 1\leq d'\leq\delta}(\mathtt{threshold}_{op',d'}(x) \wedge \mathtt{threshold}_{op,d}(x) \rightarrow \\
 & & \qquad \bigvee_{d<\beta<\beta',\ d'<\beta'\leq\delta}(\mathtt{count}_{\beta,op,i}(x) \wedge \mathtt{count}_{\beta',op',i}(x)) \vee \\
 & & \qquad \bigvee_{\beta\leq d,\ d'<\beta'\leq\delta}(\mathtt{count}_{\beta,op,i}(x) \wedge \mathtt{count}_{\beta',op',i}(x))) \\
\mathtt{count}_{\gamma,op,i}(x) & \equiv & \mathtt{atleast}_{\gamma,op,i}(x) \wedge \neg\mathtt{atleast}_{\gamma+1,op,i}(x) \\
\mathtt{atleast}_{\gamma,op,i}(x) & \equiv & \exists w_1,\dots,w_\gamma,x,y_1,\dots,y_\gamma\ \bigwedge_{1\leq h\leq\gamma}(\mathtt{from}(x,w_h) \wedge \mathtt{to}(y_h,w_h)) \wedge \\
 & & \qquad \bigwedge_{1\leq h\leq\gamma}\bigvee_{j_h}(\mathtt{holds}_{op,j_h}(y_h) \wedge \mathtt{order}_{i,j_h}(w_h)) \\
\mathtt{atleast}_{\gamma,op,\infty}(x) & \equiv & \exists w_1,\dots,w_\gamma,x,y_1,\dots,y_\gamma\ \bigwedge_{1\leq h\leq\gamma}(\mathtt{from}(x,w_h) \wedge \mathtt{to}(y_h,w_h)) \wedge \bigwedge_{1\leq h\leq\alpha}\mathtt{final}_{op}(y_h)
\end{array}$$

Figure 5: MSO encoding for maximizing the spread of opinion *white*.

The first ingredient of the encoding is the definition of the finite structure $f(\mathrm{o},S) = \mathcal{D}_1 \cup \mathcal{D}_2 \cup \mathcal{D}_3$ over the domain $N \cup \{e_{x,y}, e_{y,x} \mid \{x,y\} \in E\}$ and such that:

- $\mathcal{D}_1 = \{\mathtt{from}(x,e_{x,y}), \mathtt{to}(y,e_{x,y}), \mathtt{from}(y,e_{y,x}), \mathtt{to}(x,e_{y,x}) \mid \{x,y\} \in E\}$; that is, we consider a directed version of E and $\mathcal{D}_1$ encodes such directed edges by using two incidence relations over its endpoints;
- $\mathcal{D}_2 = \{\mathtt{threshold}_{op,d}(x) \mid x \in N \wedge op \in O \wedge t_x(op) = d\}$; that is, $\mathcal{D}_2$ encodes the thresholds of the agents over the possible opinions;
- $\mathcal{D}_3 = \{\mathtt{holds}_{\mathrm{o}(x),0}(x) \mid x \in N\}$; that is, $\mathcal{D}_3$ encodes the opinion that each agent initially (at step "0") holds in o.

On top of $f(\mathrm{o},S)$, we now define a MSO formula that mimics the dynamics for S starting at o. In fact, we would like to define a *fixed* formula, which must be independent on S and o. This makes the encoding rather challenging: In particular, the formula cannot explicitly encode the sequence of changes occurring in the dynamics, because the length of this sequence is in general not constant—indeed, it is linear in $|N|$ (cf. Lemma 2.1). To address this issue, our solution approach is to encode the dynamics implicitly, by relating the changes of each agent with the changes of her neighbors only (i.e., we encode different local views—one per agent—of the sequence, from which global information can be reconstructed) and by defining an encoding $\Phi_{\varrho,\delta}(\mathtt{W})$ that is parametric w.r.t. the maximum number ϱ of changes that every agent can perform and w.r.t. the maximum number δ of neighbors over each agent.[3] In more details, we consider the following MSO formula by referring to Figure 5 for details on the subformulas occurring in it:

$$\Phi_{\varrho,\delta}(\mathtt{W}) \equiv \exists_{i,j}\mathtt{order}_{i,j}\ \exists_{op,i}\mathtt{holds}_{op,i}\ \ \exists_{op}\mathtt{final}_{op}$$
$$\bigwedge_{i,j,h}\psi^h_{i,j} \wedge \bigwedge_{op,i,h}\phi^h_{op,i} \wedge \bigwedge_{op,h}\tau^h_{op} \wedge \xi(\mathtt{W})$$

Let us discuss the structure of $\Phi_{m,\delta}(\mathtt{W})$. Concerning the variables, note that in addition to the free set variable W (whose role will be explained later when discussing the subformula $\xi(\mathtt{W})$), $\Phi_{m,\delta}(\mathtt{W})$ uses the following sets variables:

- $\mathtt{order}_{i,j}$: For every $i \in \{1,\dots,\varrho\}$ and $j \in \{0,\dots,\varrho\}$, this is a set variable that is existentially quantified in $\Phi_{m,\delta}(\mathtt{W})$—so we have a total of $\varrho \times (\varrho+1)$ set variables of this kind. Its intended meaning is that w belongs to $\mathtt{order}_{i,j}$ only if w corresponds to an edge from $x \in N$ to $y \in N$, and x makes her i-th change in the dynamics when her adjacent node $y \in N$ has already made her j-th change and before her $(j+1)$-th change. In particular, variables $\mathtt{order}_{i,0}$ encode the fact that x makes her i-th change when y has not yet changed the opinion she initially holds.
- $\mathtt{holds}_{op,i}$: For every $op \in O$ and $i \in \{1,\dots,\varrho\}$, this is again an existentially quantified set variable—so we have a total of $|O|\times\varrho$ set variables of this kind. Its intended meaning is that $x \in N$ belongs to $\mathtt{holds}_{op,i}$ only if she gets opinion $op \in O$ at her i-th change. Note that the initial opinion of the agents are encoded in $\mathcal{D}_{S,\mathrm{o}}$ via the relation symbol $\mathtt{holds}_{op,0}$.
- $\mathtt{final}_{op}$: For eery $op \in O$, this existentially quantified set variable is meant to define to opinion that is eventually hold by each agent at the end of the dynamics. Note that each agent must be stable in the resulting profile.

[3] So, in order for $\Phi_{\varrho,\delta}(\mathtt{W})$ being independent on S and o, ϱ and δ must actually be constants.

Consider now the first group of subformulas shown in Figure 5. Rules $\psi^1_{i,j}$ (for $i > 0$ and $j \geq 0$) define the semantics of an element w belonging to $\mathtt{order}_{i,j}$: w must actually encode an edge from some node $x \in N$ to some node $y \in N$, it does not simultaneously belong to $\mathtt{order}_{i,j'}$ for some $j' \neq j$, and there must be an opinion op that x holds after her i-th change. The remaining rules define the constrains relating the order of activation between nodes x and y. For instance, $\mathtt{order}_{1,j}$ encodes that the first change of x occurred between the j-th and $(j+1)$-th change of y; hence, rule $\psi^2_{1,j}$ (with $j > 0$) specifies that the j-th change of y occurred after the first change of x. Note that these rules require to reason about the reverse edge $\bar{w}$, which is characterized via the macro, i.e., a formula meant to be substituted inline, `inverse`.

The second group of rules defines, instead, the intended meaning of the elements x belonging to $\mathtt{holds}_{op,i}$: At her i-th change, x can get precisely one opinion op, which moreover she prefers to the opinion $\hat{op} \neq op$ she holds before ($\phi^1_{op,i}$) and to any other possible opinion $op' \neq op$ she might get ($\phi^2_{op,i}$)—note that relationships of preferences are encoded via the macro `better`, which in turn uses `count` and `atleast`, all of them depending on the maximum degree δ. Moreover, rules $\phi^3_{op,i}$ prescribe that the i-th change of x is properly related to some j-th change of each of her neighbors.

The third group of rules define the semantics of $\mathtt{final}_{op}$ based on that of $\mathtt{holds}_{op,i}$, prescribing that agents will not subsequently change (there note that saturation after ϱ changes) and that they are stable in the resulting final profile. In particular, the stabling condition (τ^2_{op}) again uses the macro `better` this time with the symbolic value $i = \infty$; the macro is identical to the case where i is a natural number, except the definition of `atleast` which explicitly refer to the final profile.

Eventually, the description of the formula is completed by noticing that $\xi(\mathtt{W})$ forces the final configuration to be defined for each agent, and prescribes that `W` is the set of agents that will eventually hold opinion *white*.

Putting it All Together. For any pair of natural numbers k and δ, let $\mathcal{G}_{k,\delta}$ be the set of all graphs whose *treewidth* is bounded by k and whose degree is bounded by δ. Observe that since treewidth generalizes acyclicity, then $\mathcal{G}_{1,\delta}$ is the set of all acyclic graphs with degree bounded by δ.

Theorem 4.4. *Let $\varrho, k, \delta > 0$ be fixed natural numbers. Computing a dynamics where each agent changes her opinion at most ϱ times and maximizing the spread of white is feasible in polynomial time on all environments S with $\mathrm{G}(S) \in \mathcal{G}_{k,\delta}$.*

Proof Sketch. Let S be an environment with $\mathrm{G}(S) \in \mathcal{G}_{k,\delta}$ and let o be a profile for it. In polynomial time, we can build the finite structure $f(\mathrm{o}, S)$ in Theorem 4.3. Now, observe that the so-called *Gaifman graph* of $f(\mathrm{o}, S)$ —which is the undirected graph whose vertices are the elements of the domain and there is an edge between any pair of elements occurring in the same tuple—precisely coincides with the graph $\mathrm{G}(S)$. Hence, its treewidth is bounded by k and we can apply a master theorem for MSO formulas [2], guaranteeing that we can compute in polynomial time the set $W \subseteq N$ of nodes with maximum cardinality such that $f(\mathrm{o}, S, \mathrm{o}) \models \Phi_{\varrho,\delta}(W)$ (and also reconstruct the optimal dynamics by inspecting the values that the existentially quantified set variables assumed on the maximum). The result then follows by Theorem 4.3. ◻

5 SPREAD MAXIMIZATION IN PRACTICE

In this section, we reconsider the problem of maximizing the spread of opinion *white* from an experimental/practical viewpoint by analyzing some real-world social networks. Indeed, the hardness results described above defeat any hope to solve this problem optimally. However, we will see that there are initial distributions of the opinions that enable a larger diffusion of the *white* opinion on the networks. Interestingly, these distributions are characterized by a larger entropy among the remaining opinions. This *entropic gain* then suggests that effective spread maximization procedures can be run whenever the network exhibit such a property.

Moreover, even if we showed that in some settings the performances of greedy dynamics can be dramatically reduced by the availability of a third opinion, this is not usually not the case in the real world networks: indeed, we show that the number of agents with *white* opinion degrades as soon as the considered dynamics differs from a greedy dynamics.

Experimental Setting. We considered a benchmark consisting of 15 graph datasets, whose main features are summarized in Figure 6, which report: the name of the dataset, the number of nodes, the number of edges, the assortativity (r_{kk}), the average degree ($\langle k \rangle$), the maximum degree (k^*) and the coefficient (α) of the power law distribution that better approximates the degree distribution of the network and computed according to a well-known distance measure [10] — on Facebook (fb) and Deezer (dz), some relevant subgraphs are also considered.

The datasets fb-Art, fb-Ath, fb-Com, fb-Gov, fb-NS, fb-Pol, fb-PF, fb-TvS have been extracted from the Facebook (fb) dataset by considering respectively artists', athletes', companies', government's, new sites', politicians', public figures' and TV shows' pages only. The datasets dz-HR, dz-HU and dz-RO have been extracted from the Deezer dataset by considering the friendships networks of users in Croatia, Hungary, and Romania, respectively.

Each graph dataset is viewed as the interaction graph underlying a social environment with three opinions and populated by majority agents. Unless stated otherwise, opinion *white* has been initially assigned to the 5% of agents/nodes having the maximum degree, and greedy strategies to propagate opinion *white* are considered such that, whenever no agent can change her opinion to *white*, a random agent that can change her opinion to *black* or *gray* is selected. Results are reported as the average on three runs.

Results. In the first experiment, (in addition to the above initialization for *white*), *black* has been assigned to a number of agents, selected at random, in a way that their percentage varies from 2.5% to 47.5%. The remaining agents initially hold opinion *gray*. Dynamics are then simulated and results (percentage of agents that will eventually adopt opinion *white*) are reported in the plot on the left of Figure 6. Note that the more the percentage of nodes *gray* and *black* is initially balanced, the more *white* spreads in the network.

To stress this phenomenon, in a second experiment, we considered the most balanced setting and we measured a parameter

Network	#nodes	#edges	r_{kk}	$\langle k \rangle$	k^*	α	eg	$\langle W, k\rangle^*$	$0.75 \cdot \langle W, k\rangle^*$	$0.5 \cdot \langle W, k\rangle^*$
youtube [43]	1134890	2987624	-0.0369	5.265	28754	1.9	1.16	73.52%	72.67%	71.91%
fb [36]	134873	1380293	0.0740	20.462	1469	1.3	1.89	74.07%	65.26%	45.28%
deezer [36]	143884	846915	0.3320	11.772	420	1.3	1.02	38.21%	38.94%	39.72%
dblp [43]	317080	1049866	0.2665	6.622	343	1.5	2.81	48.03%	45.88%	42.36%
fb-Art	50521	819306	-0.019	32.43	1469	1.2	1.10	98.70%	98.75%	98.83%
fb-Ath	13868	86858	-0.027	12.52	468	1.3	1.31	92.44%	93.04%	93.57%
fb-Com	14120	52310	0.014	7.40	215	1.4	1.83	58.79%	59.69%	60.23%
fb-Gov	7058	89455	0.029	25.34	697	1.2	1.71	68.56%	65.19%	60.42%
fb-NS	27930	206259	0.022	14.76	678	1.3	1.54	84.50%	84.41%	86.24%
fb-Pol	5908	41729	0.018	14.12	323	1.3	1.84	48.68%	47.89%	23.35%
fb-PF	11573	67114	0.202	11.59	326	1.4	4.92	71.35%	70.11%	71.39%
fb-TvS	3895	17262	0.561	8.86	126	1.3	1.77	18.13%	18.89%	15.99%
dz-HR	54573	498202	0.197	18.26	420	1.2	9.25	99.20%	99.22%	99.23%
dz-HU	47538	222887	0.207	9.38	112	1.2	9813	98.13%	98.18%	7.95%
dz-RO	41773	125826	0.114	6.02	112	1.4	162.76	83.01%	83.92%	17.36%

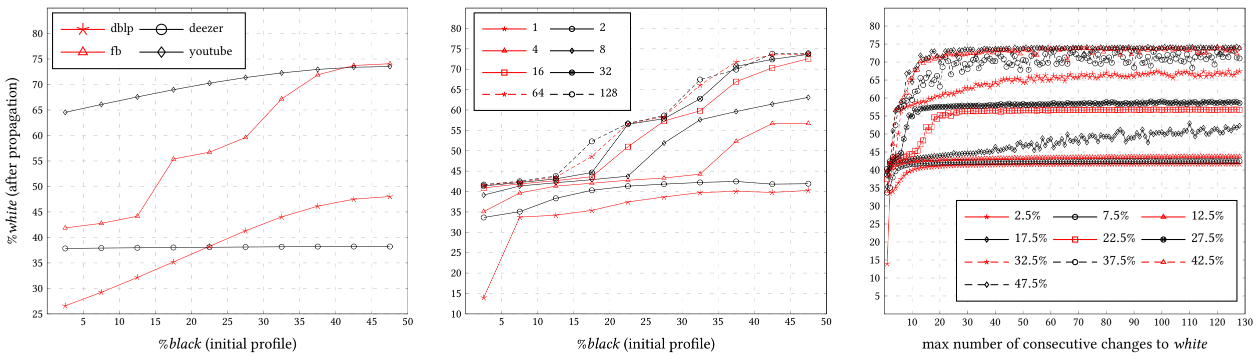

Figure 6: Dataset characteristics and summary of results.

eg, which we call *entropic gain* and which is defined as the ratio between the number of nodes that will eventually hold opinion *white* after a converging dynamics involving three opinions, and the number of nodes that will eventually hold opinion *white* after a converging dynamics starting from the same initial profile but where all opinions *gray* are changed to *black* (so that propagation involves two opinions only). For the datasets we have considered, the entropic gain is reported in the eighth column of the table, whereas in the column $\langle W, k\rangle^*$ we report the percentage of nodes that hold opinion *white* after the propagation.

In a third experiment, we contrasted the percentage reported in the column $\langle W, k\rangle^*$ with that obtained by lowering the correlation between attribute degree with respect to opinion *white* to the 75% and 50% of the maximum value (see the last two columns of the table). In particular, to decrease the correlation, we repeatedly chose at random pair of nodes x and y, with x having attribute *white*, and swap their attributes if the degree of y is greater than the degree of x. Note that the results, in most of the networks, do not vary significantly w.r.t. the different values of the correlation, which means that the initial seed of 5% of *white* nodes suffices for leading to the given values of propagations (rather) independently on how such nodes are selected.

In the last experiment, we considered different propagation strategies obtained by constraining between 1 and 128 the number of consecutive changes to *white*—in a greedy strategy there is no bound and we change agents' opinions to *white* as long as this is possible. Results are reported in the remaining two plots of Figure 6, where in particular, in the plot on the right, the series are the initial percentages of agents having opinion *black*. The results evidence that the performances of the propagation saturate even with rather low values for the above constraint.

6 CONCLUSION

We have analyzed opinion diffusion according a specific game-theoretic dynamics with three alternative opinions. Our results are foundational and define the landscape from which other relevant problems in social network analysis can be studied. For instance, an interesting avenue of further research is to study, within our setting, the Target Set Selection problem [14, 32], manipulation problems [3, 5–7, 12], as well as problems of diffusions where activated agents can express negative feedbacks [24, 34].

Moreover, from the technical viewpoint, it would be interesting to study whether opinion diffusion is tractable on larger classes than $\mathcal{G}_{k,\delta}$ and without any bound on ϱ. In fact, the bound on the degree might be avoided by mapping the MSO encoding to a CSP instance (indeed, note that there are no universally quantified variables) and then using the techniques in [28]. However, it is challenging to identify non-trivial classes of environments where, w.l.o.g., ϱ can be fixed within some given bound.

Finally, it would be of extreme interest to extend the analysis of this work to different models of opinion diffusion.

REFERENCES

[1] Erman Acar, Gianluigi Greco, and Marco Manna. 2017. Group Reasoning in Social Environments. In *Proc. of AAMAS'17*. 1296–1304.

[2] Stefan Arnborg, Jens Lagergren, and Detlef Seese. 1991. Easy Problems for Tree-decomposable Graphs. *J. Algorithms* 12, 2 (1991), 308–340.

[3] Vincenzo Auletta, Ioannis Caragiannis, Diodato Ferraioli, Clemente Galdi, and Giuseppe Persiano. 2015. Minority Becomes Majority in Social Networks. In *Proc. of WINE-15*. 74–88.

[4] Vincenzo Auletta, Ioannis Caragiannis, Diodato Ferraioli, Clemente Galdi, and Giuseppe Persiano. 2016. Generalized Discrete Preference Games. In *Proc. of IJCAI-16*, Vol. 16. 53–59.

[5] Vincenzo Auletta, Ioannis Caragiannis, Diodato Ferraioli, Clemente Galdi, and Giuseppe Persiano. 2017. Information Retention in Heterogeneous Majority Dynamics. In *Proc. of WINE'17*. 30–43.

[6] Vincenzo Auletta, Ioannis Caragiannis, Diodato Ferraioli, Clemente Galdi, and Giuseppe Persiano. 2017. Robustness in Discrete Preference Games. In *Proc. of AAMAS '17*. 1314–1322.

[7] Vincenzo Auletta, Diodato Ferraioli, and Gianluigi Greco. 2018. Reasoning about Consensus when Opinions Diffuse through Majority Dynamics. In *Proc. of IJCAI-18*. 49–55.

[8] Alexandru Baltag, Zoé Christoff, Rasmus K. Rendsvig, and Sonja Smets. 2018. Dynamic Epistemic Logics of Diffusion and Prediction in Social Networks. *Studia Logica* (2018).

[9] Oren Ben-Zwi, Danny Hermelin, Daniel Lokshtanov, and Ilan Newman. 2011. Treewidth governs the complexity of target set selection. *Discrete Optimization* 8, 1 (2011), 87 – 96.

[10] A. Bhattacharyya. 1943. On a measure of divergence between two statistical populations defined by their probability distributions. *Bulletin of the Calcutta Mathematical Society* 35 (1943), 99–109.

[11] Allan Borodin, Yuval Filmus, and Joel Oren. 2010. Threshold models for competitive influence in social networks. In *Proc. of WINE-10*. 539–550.

[12] Robert Bredereck and Edith Elkind. 2017. Manipulating Opinion Diffusion in Social Networks. In *Proc. of IJCAI'17*. 894–900.

[13] Markus Brill, Edith Elkind, Ulle Endriss, and Umberto Grandi. 2016. Pairwise Diffusion of Preference Rankings in Social Networks. In *Proc. of IJCAI-16*. 130–136.

[14] Ning Chen. 2009. On the Approximability of Influence in Social Networks. *SIAM J. Discrete Math.* 23, 3 (2009), 1400–1415.

[15] Flavio Chierichetti, Jon Kleinberg, and Sigal Oren. 2018. On discrete preferences and coordination. *J. of Computer and System Sciences* 93 (2018), 11 – 29.

[16] Laurence Cholvy. 2016. Influence-Based Opinion Diffusion. In *Proc. of AAMAS-16*. 1355–1356.

[17] Zoé Christoff and Davide Grossi. 2017. Stability in Binary Opinion Diffusion. In *Proc. of LORI-17*. 166–180.

[18] Zoé Christoff and Jens Ulrik Hansen. 2015. A Logic for Diffusion in Social Networks. *J. of Applied Logic* 13, 1 (2015), 48–77.

[19] Morris H. DeGroot. 1974. Reaching a Consensus. *J. Amer. Statist. Assoc.* 69, 345 (1974), 118–121.

[20] David Easley and Jon Kleinberg. 2010. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press.

[21] Ulle Endriss (Ed.). 2017. *Trends in Computational Social Choice*. AI Access.

[22] Diodato Ferraioli, Paul W. Goldberg, and Carmine Ventre. 2016. Decentralized dynamics for finite opinion games. *Theor. Comput. Sci.* 648 (2016), 96–115.

[23] Michael R. Garey and David S. Johnson. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co.

[24] Aristides Gionis, Evimaria Terzi, and Panayiotis Tsaparas. 2013. Opinion maximization in social networks. In *Proc. of SIAM ICDM-13*. 387–395.

[25] Umberto Grandi, Emiliano Lorini, Arianna Novaro, and Laurent Perrussel. 2017. Strategic Disclosure of Opinions on a Social Network. In *Proc. of AAMAS'17*. 1196–1204.

[26] Umberto Grandi, Emiliano Lorini, and Laurent Perrussel. 2015. Propositional Opinion Diffusion. In *Proc. of AAMAS-15*. 989–997.

[27] Mark Granovetter. 1978. Threshold models of collective behavior. *American journal of sociology* 83, 6 (1978), 1420–1443.

[28] Gianluigi Greco, Francesco Lupia, and Francesco Scarcello. 2015. Structural Tractability of Shapley and Banzhaf Values in Allocation Games. In *Proc. of IJCAI-15*. 547–553.

[29] Jens Ulrik Hansen. 2015. Reasoning about opinion dynamics in social networks. *Journal of Logic and Computation* (2015).

[30] Matthew O. Jackson. 2008. *Social and Economic Networks*. Princeton University Press, Princeton, NJ, USA.

[31] David Kempe, Jon Kleinberg, and Éva Tardos. 2003. Maximizing the spread of influence through a social network. In *Proc. of KDD-03*. ACM, 137–146.

[32] David Kempe, Jon Kleinberg, and Éva Tardos. 2005. Influential Nodes in a Diffusion Model for Social Networks. In *Proc. of ICALP'05*. 1127–1138.

[33] Keith Lehrer and Carl Wagner. 1981. Rational Consensus in Science and Society: A Philosophical and Mathematical Study. 24 (1981).

[34] Xinyue Liu, Xiangnan Kong, and Philip S Yu. 2018. Active Opinion Maximization in Social Networks. In *Proc. of ACM KDD-18*. 1840–1849.

[35] Eric Pacuit and Rohit Parikh. 2005. The Logic of Communication Graphs. In *Declarative Agent Languages and Technologies II*. Vol. 3476. Springer, 256–269.

[36] Benedek Rozemberczki, Ryan Davies, Ryan Sarkar, and Charles A. Sutton. 2018. GEMSEC: Graph Embedding with Self Clustering. *CoRR* (2018).

[37] Katsuhiko Sano and Satoshi Tojo. 2013. Dynamic Epistemic Logic for Channel-Based Agent Communication. In *Logic and its Applications*. Springer, 109–120.

[38] Thomas C Schelling. 2006. *Micromotives and macrobehavior*. WW Norton & Company.

[39] Jeremy Seligman, Fenrong Liu, and Patrick Girard. 2011. Logic in the Community. In *Logic and Its Applications*. Springer, 178–188.

[40] Paulo Shakarian, Abhinav Bhatnagar, Ashkan Aleali, Elham Shaabani, and Ruocheng Guo. 2015. *Logic Programming Based Diffusion Models*. Springer, 49–73.

[41] Paulo Shakarian, Matthias Broecheler, V. S. Subrahmanian, and Cristian Molinaro. 2013. Using Generalized Annotated Programs to Solve Social Network Diffusion Optimization Problems. *ACM Trans. Comput. Logic* 14, 2, Article 10 (2013), 10:1–10:40 pages.

[42] Li-Xin Wang and Jerry M Mendel. 2016. Fuzzy Opinion Networks: A Mathematical Framework for the Evolution of Opinions and Their Uncertainties Across Social Networks. *IEEE Trans. Fuzzy Systems* 24, 4 (2016), 880–905.

[43] Jaewon Yang and Jure Leskovec. 2012. Defining and Evaluating Network Communities Based on Ground-Truth.. In *Proc. of ICDM-12*. 745–754.