How to Mitigate Disagreement and Polarization in Opinion Formation Processes on Social Networks

Vincenzo Auletta[1], Diodato Ferraioli[1] and Grazia Ferrara[1,*]

[1]Università degli Studi di Salerno, Italy

Abstract
In recent years, we have been observing an important increase of the phenomena of disagreement and polarization in our society. Among the factors that have been recognized as causes of these social processes, there are the algorithms adopted by the social media platforms to select news and messages to present to their users that are designed to increase their social engagement. Since these phenomena may have very dangerous and destructive effects in terms of social cohesion, it is of great interest to design methods that can be adopted by social media platforms in order to mitigate disagreement and polarization in the process of opinion formation. In this work, we propose mitigation methods based on seeding. Seeding is largely used in viral marketing and opinion diffusion campaigns, and it consists of injecting information into the network by some influential nodes called seeds. We propose using information campaigns starting from seeds that were opportunely chosen to mitigate disagreement and polarization in the network. We consider two different scenarios: in the first one we assume that the whole graph of the social network is known and we present an efficient greedy-based heuristics to select a given number of seeds in order to minimize disagreement and polarization; in the second case, we assume that the social graph is unknown and we present an online learning algorithm that can be used to learn the graph while the opinion diffusion dynamics is run. Finally, in order to evaluate and demonstrate the functionality of our framework, we present some experimental results on the performance of our algorithms on a comprehensive collection of synthetic and real-world networks.

Keywords
Opinion Formation, Online Learning, Social Networks, Seeding

1. Introduction

The formation of opinions represents a social process through which individuals develop beliefs, attitudes and judgments about a specific topic by interacting with other individuals.

Nowadays, with the rise of online social networks, opinion formation processes have assumed a crucial role in a multitude of domains, including social media-related phenomena such as the diffusion of information through social campaigns, the formation of echo chambers, the impact of influencers and the manipulation of opinions through the diffusion of misinformation. These phenomena are having a huge impact in critical domains such as political campaigns. For example, managers of political campaigns largely use targeted messaging strategies to influence the electors, or even tools to predict electoral outcomes based on the analysis of public opinion trends observed through information campaigns. Another critical domain where phenomena related to opinion formation processes is having a huge impact is the spread of health-related information in vaccination campaigns, which could have a significant impact on the outcome of a pandemic. For example, Cascini et al. [2], made a systematic review that investigates how social media affected attitudes towards COVID-19 vaccination. Their work aims to understand the influence of social media on public health campaigns and how it can address vaccine hesitancy. Other applications of the opinion formation processes can be found in the commercial domain too. For example, Tu et al. [3], found out that marketing campaigns and polarizing content may influence network polarization differently, with polarizing content exerting a more substantial effect.

AI4CC-IPS-RCRA-SPIRIT 2024: International Workshop on Artificial Intelligence for Climate Change, Italian Workshop on Planning and Scheduling, RCRA Workshop on Experimental evaluation of algorithms for solving problems with combinatorial explosion, and SPIRIT Workshop on Strategies, Prediction, Interaction, and Reasoning in Italy. November 25-28th, 2024, Bolzano, Italy [1].
*Corresponding author.
auletta@unisa.it (V. Auletta); dferraioli@unisa.it (D. Ferraioli); grferrara@unisa.it (G. Ferrara)

In the last years social media platforms have revolutionized how individuals form and share their opinions and several studies point out the crucial role that these platforms play in shaping the flow of information and, consequently, the formation of public opinion. As a matter of fact, these platforms adopted sophisticated algorithms to enhance user engagement by curating content tailored to individual preferences. These algorithms analyze vast amounts of data, including users' past behavior, interests and interactions, to present content that aligns with their existing beliefs and preferences. While this personalization enhances user experience by making content more relevant to the user, it also creates echo chambers, where users are predominantly exposed to viewpoints that reinforce their existing opinions. This phenomenon, known as *filter bubbles*, limits the diversity of information that users encounter, making it increasingly difficult for them to consider alternative perspectives. Chitra et al. [4] discussed the impact of this phenomenon on social networks, highlighting how algorithms can create echo chambers by recommending content that aligns with users' existing beliefs, leading to increased polarization among individuals.

Furthermore, users are often overwhelmed by the constant influx of information, leading them to rely on the platform's curated feed rather than seeking out diverse sources. As a result, they are less likely to engage in meaningful discussions with those who hold differing opinions. However, these mechanisms, which have been created with the intention of enhancing the efficacy and usability of social media, are inadvertently contributing to a significant increase in the prevalence of disagreement and polarization within society that may have a severe negative impact on collective decision making processes.

Here, *disagreement* occurs when individuals hold different opinions, beliefs or perspectives on a particular topic. Negative consequences might arise from disagreement, especially if the whispered outcome of the opinion formation process was to maximize consensus towards a given opinion or position. Disagreement could be harmful to the successful outcome of a campaign, making it result in a huge disaster.

On the other hand, *polarization* refers to the increasing divergence of opinions within a population and it is another problem which could prevent a collective decision making process from succeeding. There is polarization when a society becomes divided into two sharply contrasting groups or sets of opinions. Polarization might be potentially very dangerous because individuals tend to hold more extreme positions within their respective groups and, in addition, people tend to interact primarily with like-minded individuals, reinforcing their existing beliefs.

Given the role that social media platforms have acquired in our society it is fundamental to encourage their commitment to promote social welfare. For example, they should redesign their algorithms to minimize both disagreement and polarization in order to facilitate enhanced social cohesion, increased trust among individuals and groups, and a reduced possibility of conflicts.

Our contribution. This work wants to address the issue of minimizing disagreement and/or polarization of opinions that occur as a result of the opinion formation process. The questions we want to answer in this sense are:

1. *Could we find an efficient algorithm to minimize disagreement and/or polarization independently from the structure of the underlying social network?*
2. *Could we make the algorithm work effectively independently from the used opinion formation model?*
3. *Could we make the algorithm work effectively even when the initial opinions are highly discordant/polarized?*

Furthermore, it would be beneficial to examine the issue in greater depth by considering a scenario in which the characteristics of the underlying social network are completely unknown or partially known. For instance, if the information pertaining to the weights on edges is not available, learning the weights could facilitate a deeper understanding of the strength of the connections between individuals who are involved in the opinion formation process and whose ideas have potentially changed over time. In this context, the main question we try to answer is:

4. *How accurately can we infer network structures (edge weights) based on observed opinion dynamics?*

The first three questions were answered by considering a greedy algorithm selecting some seeds to manipulate, in order to minimize a given metric of interest, before applying a certain opinion formation model. Unfortunately, this algorithm is not executable in a reasonable amount of time for medium or large networks, so a more efficient and scalable heuristics approximating the greedy one was proposed. The latter, which will be referred to as LocalGreedy, outperforms the greedy one, since with an approximation ratio of 95.2%, it is almost 169 times faster than it on the average. With respect to the setting of the network with unknown weights, the performance in minimizing the error in the learned weights depends on the specific opinion formation model. In general, the algorithm has an average error value on the learned weights that is between 0.1 and 0.3 at the end of the learning process.

Related Works. An opinion formation model is a mathematical framework designed to simulate the evolution of opinions within a population over time. It aims to reproduce the dynamics of opinion changes based on various factors, including social interactions, media influence, and individual characteristics. The study of opinion formation models has a long history: Degroot proposed a first such model [5] to lead a group or a committee of people to reach a consensus on a certain subject, developing an agreement on a single final opinion.

Another model of great importance is the Friedkin-Johnsen model [6, 7], which is a variation of the Degroot [5] one, but it takes also into account the effects of individual beliefs in the opinion formation process. In both these works, an agent is influenced by all its contacts, regardless of their opinions.

Several other models weakened this assumption by assuming that one agent opinion is influenced only by those contacts with opinions close to the one currently expressed by the agent [8, 9]. A weighted version of the latter was also proposed [10], where the interactions between agents are not only based on opinion proximity, but also account for the heterogeneity in social bonds and trust. In this variant the influence of each neighbor is modulated by a weight representing the strength of the social relationship, providing a more realistic portrayal of opinion dynamics in social networks.

While our work mainly focuses on these four models, several other variants and generalizations of them have been introduced in literature, by considering discrete opinions [11, 12], limited/local interactions [13, 14], dynamic settings where social relationships and internal beliefs evolve over time [15, 16, 17, 18], mixture of attraction and repulsion in opinion formation [19, 20, 21].

In this work we focus on minimizing the disagreement and/or polarization of opinions in any possible social network. These metrics were also taken into consideration by Matakos et al. [22], Musco et al. [23] and Zhu et al. [24] with the aim of minimizing them. Other works, such as Chen et al. [25], used an adversarial approach and they aimed to maximize them instead. Additional other works, such as Gionis et al. [26], want to maximize the overall opinion of the society instead. Note that most of the literature on opinion dynamics focused on different goals: minimizing the *utilitarian social cost*, defined as the sum of the agents' costs (according to the redefinition of the Friedkin-Johnsen model by Bindel et al. [7]) [7, 12, 15, 16]; the *truthfulness* of the declared opinions, by bounding how much the social pressure deviates the agents' opinions from their private beliefs [27, 28, 29]; the distance from a *consensus* [9, 18].

Manipulation problems involving social networks in the setting of opinion diffusion, adopt approaches that are essentially divided in: changing the opinion of some nodes [30, 31, 32, 33, 34]; managing (adding/removing) edges [33, 34, 35, 36]; changing the order in which opinions are updated [27, 28, 29, 33, 37, 38, 39]. We here focus mainly on the first approach. Note that this is strictly related with the *Social Influence Maximization problem* (SIM) introduced by Domingos and Richardson [40] in the context of viral marketing. Two in-depth surveys related to the problem of the seed set selection for influence maximization are Li et al. [41] and Banerjee et al. [42].

Very few works are known considering the opinion formation processes in a setting with unknown network parameters: De et al. [43] studied a setting in which opinion formation processes are run according a Degroot model on social networks with weights not known a priori, and the main objective is to learn the parameters of the underlying model; Wai et al. [44] propose an online optimization algorithm to actively learn trust parameters in social networks using the Degroot model under the influence of stubborn agents. Note that the Degroot model consistently attains consensus, while our work considers also more complex models that do not necessarily converge to a consensus.

2. Definitions

The social network is represented as a weighted graph $G = (V, E, W)$, where V is the set of vertices, E is the set of edges and $\forall(u, v) \in E$, $w_{uv} \in [0, 1]$ represents the strength of the link between the node u and the node v. Each node of the network represents an agent i, which owns a real-valued initial opinion $x_i \in [0, 1]$. Agents iteratively update their own opinion according to a given opinion formation model. We will denote as $y_i^{(t)}$ the opinion expressed by agent i at timestep t, and we will consider $y_i^{(0)} = x_i$. Several continuous models were considered in this work.

2.1. Models

Degroot. In the Degroot [5] model, agents update their opinion by making a weighted average of their neighbors' opinions at the current timestep:

$$y_i^{(t)} = \sum_{j \in N(i)} w_{ij} y_j^{(t-1)}, \tag{1}$$

where $y_i^{(t)}$ is the opinion of node i at time t, $y_j^{(t-1)}$ is the opinion of node j at time $t-1$, $w_{ij} \in [0, 1]$ is the weight on the edge (i, j) in G, $N(i)$ is the set of neighbors of i in G.

Friedkin-Johnsen (FJ). In the Friedkin-Johnsen model [6], for $t = 1$, the opinions (Y) of n individuals are completely determined by a set of beliefs X_1, that is

$$Y_1 = X_1 B_1,$$

where Y_1 is a $n \times 1$ vector of opinions, X_1 is a $n \times k$ matrix of scores on k beliefs, and B_1 is a $k \times 1$ vector of coefficients giving the effects of each of the beliefs; for $t > 1$, the opinions (Y_t) of the n individuals continue to be affected by a set of X_t beliefs, but are also endogenously affected by their own and others' opinions at the immediately previous instant, that is

$$Y_t = \alpha_t W_t Y_{t-1} + \beta_t X_t B_t,$$

where α_t is a scalar weight of the endogenous conditions, β_t is a scalar weight of the belief, W_t is a $n \times n$ matrix of the effects of each opinion held at time $t-1$ on the n opinions held at time t. The equilibrium opinions obtained through this process are $\mathbf{y} = (\mathbf{L} + \mathbf{I})^{-1}\mathbf{x}$, where $\mathbf{L}$ is the Laplacian matrix of the graph, $\mathbf{I}$ is the identity matrix and $\mathbf{Q} = (\mathbf{L} + \mathbf{I})^{-1}$ is known as *fundamental matrix*. However, for our purposes, we are interested in considering the updating rule provided in the Bindel et al.'s formulation [7] of the FJ model:

$$y_i^{(t)} = \frac{x_i + \sum_{j \in N(i)} w_{ij} y_j^{(t-1)}}{1 + \sum_{j \in N(i)} w_{ij} y_j^{(t-1)}}, \tag{2}$$

where x_i is the initial opinion of agent i and the other parameters are the same defined above in the Degroot updating rule.

Deffuant. In the Deffuant model [8], at each timestep t, two agents are randomly selected and if their opinions are such that $|y_i^{(t-1)} - y_j^{(t-1)}| < d$, they update to:

$$y_i^{(t)} = y_i^{(t-1)} + \eta \cdot (y_j^{(t-1)} - y_i^{(t-1)}), \tag{3}$$

$$y_j^{(t)} = y_j^{(t-1)} + \eta \cdot (y_i^{(t-1)} - y_j^{(t-1)}) \tag{4}$$

where η is the convergence parameter, whose value may vary between 0 and 0.5.

Hegselmann-Krause (HK). In the HK model [9], agents update their opinion based on an average of the opinions of their neighbors whose opinion at the previous timestep did not differ more than a confidence bound ε_i from their own. The updating rule is:

$$y_i^{(t)} = |I(i, y^{(t-1)})|^{-1} \sum_{j \in I(i, y^{(t-1)})} y_j^{(t-1)}, \tag{5}$$

where $I(i, y^{(t-1)}) = \{1 \le j \le n : |y_i^{(t-1)} - y_j^{(t-1)}| < \varepsilon_i\}$, and ε_i is the confidence bound of agent i.

Weighted Hegselmann-Krause (WHK). In this work, we also consider a weighted variant of the HK model. One notable approach, introduced by Toccaceli et al. [10], incorporates weights to capture the heterogeneity of interaction frequencies and social bonds between agents. Their model introduces non-linear updates based on the sign of the agent's opinion, where the influence of neighboring opinions is adjusted differently depending on whether the opinion is positive or negative.

In contrast, we consider a simplified version of the model in [10], where all opinions are taken in the $[0, 1]$ range uniformly, without distinguishing between positive and negative opinions. Specifically, in our formulation, the influence of each neighboring agent is directly proportional to the weight assigned to the interaction, and the opinion update is based purely on the weighted average of neighbors' opinions. The update rule is:

$$y_i^{(t)} = \frac{\sum_{j \in I(i, y^{(t-1)})} w_{ij} y_j^{(t-1)}}{\sum_{j \in I(i, y^{(t-1)})} w_{ij}}$$

where $I(i, y^{(t-1)}) = \{1 \le j \le n : |y_i^{(t-1)} - y_j^{(t-1)}| < \varepsilon_i\}$, ε_i is the confidence bound of agent i and the weights w_{ij} satisfy $w_{ij} \ge 0$. In practice, we see that the behavior of WHK is almost identical to that of HK.

2.2. Metrics

To evaluate the quality of the outcome of the opinion formation dynamics, we consider the following metrics, which were previously defined by Musco et al. [23].

Disagreement. The *disagreement* metric is defined as the sum of the squared differences in the opinions between neighbors:

$$D = \sum_{(i,j) \in E} w_{ij}(y_i - y_j)^2. \tag{6}$$

Polarization. The *polarization* metric, is defined as the variance of opinions around the mean, given by:

$$P = \sum_{i \in V} \left(y_i - \frac{\sum_{j \in V} y_j}{n} \right)^2. \tag{7}$$

Disagreement-Polarization Index. The *DP-index* metric is defined as the sum of disagreement and polarization:

$$I = D + P. \tag{8}$$

2.3. Problems

According to the above definitions, in this work we considered the following problems:

PROBLEM 1. (S-MIN-D) *Given a graph $G = (V, E)$ and an integer k, identify the set S of k nodes such that fixing the expressed opinions of the nodes in S to a target opinion 0.5 minimizes the overall disagreement $D(\mathbf{y}|S)$ (6).*

PROBLEM 2. (S-MIN-P) *Given a graph $G = (V, E)$ and an integer k, identify the set S of k nodes such that fixing the expressed opinions of the nodes in S to a target opinion 0.5, minimizes the overall polarization $P(\mathbf{y}|S)$ (7).*

PROBLEM 3. (S-MIN-DP) *Given a graph $G = (V, E)$ and an integer k, identify the set S of k nodes such that fixing the expressed opinions of the nodes in S to a target opinion 0.5, minimizes the overall DP-index $I(\mathbf{y}|S)$ (8).*

PROBLEM 4. (LEARN-W) *Given a weighted graph $G = (V, E)$ with unknown weights W and a Opinion Formation Model M executing after k seeds have been identified and manipulated to express a target opinion, learn the weights of the network.*

It is not difficult to see that all the considered problems are NP-hard.

3. Our Algorithms

3.1. LocalGreedy

In Gionis et al. [26] they consider the problem of maximizing the overall opinion of the society and model it as a problem similar to the influence maximization one faced in Kempe et al. [45]. Hence, in order to select a seed set S of nodes maximizing the overall opinion, they adapt the greedy algorithm discussed in Kempe et al. [45] for monotone and submodular optimization problems. Unfortunately neither polarization, nor disagreement are monotone and submodular with respect to the choice of the seeds. Anyway, we experimentally observe that the performance of the greedy algorithm are close to the one achieved by the optimal exhaustive search algorithm (at least for the small networks in which the latter algorithm can be run) for all the opinion formation models considered above and for all metrics of interests. Unfortunately, the computational performance of this greedy algorithm becomes prohibitive even for networks of modest size.

In order to achieve a precision comparable with the one showed by greedy, but a greater efficiency in the execution, a local version of the greedy algorithm has been designed. This further approximation is based on the idea that, in order to estimate the impact of a potential seed in the opinion formation process, it is not necessary to simulate the process on the whole network, which could potentially have a huge number of nodes and edges, but it is sufficient to execute it in ”its neighborhood”. In particular, the algorithm allows to locally run the simulation specifying:

- the *depth d*, which is the number of levels of nodes of the network which have to be considered starting from a given node,
- the *convergence threshold* ε, which implicitly decides the number of iterations that are necessary for the model to stop.

It is straightforward to observe that the depth must be set according to the size of the network (e.g. if the network has thousands of nodes and edges, it is probable that a depth of 1 will give a poor approximation). Also the threshold ϵ has to be differently set with respect to the pure greedy algorithm, since it is likely that the simulation will reach convergence faster on the sub-network, rather than on the entire network. Our procedure is described in detail in Algorithm 1.

3.2. An online learning algorithm

In order to make the LEARN-W problem solvable for all the considered opinion formation models, an online learning approach will be used. In this algorithm it is necessary to define two main entities: a *Learner*, who owns a network G', which is a copy of G with estimated weights W', the opinion formation process to apply M and the metric of interest σ, and she is able to compute a seed set S, optimizing that

Algorithm 1 LocalGreedy

Input: Graph $G = (V, E)$, Budget k, Cost Function σ, Model M, Depth d
Output: Seed set S

```
S ← deque()
for i ← 1 to k do
    n̄ ← None
    Δ̄ ← −∞
    y ← Run(M,S)                              ▷ Run the opinion formation model M
    for each n ∈ V do
        if n ∉ S then
            Append(S, n)                      ▷ Temporarily add node n to the seed set S
            N ← Neighbourhood(G, n, d)        ▷ Compute the neighbourhood
            ȳ ← RunNeighbourhood(M, N, S)     ▷ Locally run the opinion formation model
            Pop(S)                            ▷ Remove node n from the seed set S
            μ ← σ(N, y)                       ▷ Compute the current local metric μ
            μ̄ ← σ(N, ȳ)                       ▷ Compute the new local metric μ̄
            Δ ← μ − μ̄                         ▷ Compute the marginal improvement Δ
            if Δ > Δ̄ then                     ▷ Check if this improvement is the best so far
                n̄ ← n                         ▷ Update the best node candidate to n
                Δ̄ ← Δ                         ▷ Update the best marginal gain to Δ
            end if
        end if
    end for
    Append(S, n̄)                              ▷ Add the best node found in this iteration to the seed set S
end for
return S
```

metric; an *Environment*, which owns the network G with the original weights W and knows the opinion formation model M which has to be applied.

The algorithm works in the following way:

1. the *Learner* applies LocalGreedy on the network G' with fake weights W' and it obtains a seed set S, which is passed to the Environment,
2. the *Environment*, with that seed set, applies the opinion formation process and obtains opinions, which are returned to the *Learner*,
3. the *Learner* executes the same operation: applies the opinion formation model with that seed set, but on the network with estimated weights W', and it obtains opinions. Based on the opinions returned from the *Environment* and the ones it computed autonomously, it updates the weights of the network.

The problem that arises naturally is how the weights have to be updated. The weights update has to be done in a way such that increasing the number of iterations, they tend to the ones of the real network owned by the Environment.

4. Experimental Results

An experimental analysis has been run on a comprehensive collection of synthetic and real-world networks to evaluate and demonstrate the functionality of the proposed framework.

The real-world networks were taken both from SNAP [46], Netzscheduler [47] and Network Repository [48] and they include well-known small datasets such as: Karate Club, Les Miserables, Dolphins, Twitter Interaction Network for the US Congress, Diseasome, Netscience, GR-QC network, US Power

Grid, Erdos-992, Bcspwr10, Ia-reality, Facebook Pages Government, Wikipedia elections, Dmela Network, Anybeat Online Social Network, Facebook Pages Company, Condmat, GPlus. The synthetic networks, instead, are generated using random graph models to simulate different network topologies. These networks are valuable for understanding how algorithms perform under controlled conditions with varying network densities and sizes.

For each of these networks, we compared the performance of Greedy with respect to the exhaustive search algorithm, and the performance and the scalability of LocalGreedy with respect to Greedy.

For sake of space, we here only discuss some results and refer the interested reader to https://gitlab.com/CS-lab1/mitigate-disagreement-polarization for a more comprehensive presentation of the results of our experimental analysis with figures and numerical details. For the Degroot process (truncated before the process reached convergence) and small networks it is possible to see that LocalGreedy is 30 times lower than Greedy on the average, and it achieves a comparable approximation in terms of polarization. Similar results for the same experiment can be obtained with the other metrics.

As a matter of fact, for the Polarization metric it is possible to see that the best trade-off between performance and scalability is to set LocalGreedy with depth 1, while for the DP-index metric, the best trade-off is achieved by LocalGreedy with depth 2. The same experiments were repeated for the other opinion formation models, and the observed behavior is similar to the one observed for the Degroot model.

In order to evaluate the scalability of LocalGreedy, we have also evaluated the size of the largest network for which the problem has been solved by the given algorithm in a certain amount of time. We observed that starting with networks with 400 nodes and a probability of inserting arcs of 0.4, Greedy's execution time exceeds one hour. The same happens for networks with 500 nodes and a 0.2 probability of inserting arcs and 600 nodes and a 0.1 probability of inserting arcs. On the other hand, in 1 hour LocalGreedy has been able to execute the same experiment for random graphs until 4000 nodes and 1600493 edges.

Furthermore, we evaluated the performance of the online learning algorithm on networks of different dimensions. The learning was executed on a time horizon of 1000 iterations and the initial opinions were generated uniformly at random, while the fake weights were initialized to have values close to 0.

The manipulation was run selecting just one seed and using $depth = 1$ for the LocalGreedy algorithm.

Accuracy in learning opinions was high, with errors decreasing rapidly and converging to near-zero levels. This ensured that the final opinions closely matched those of the original network. Weight estimation also exhibited consistent improvement, with the mean squared error steadily declining over iterations. Although larger networks presented slightly higher errors, the values remained within acceptable ranges. Across different opinion formation models, the results were generally similar. For Degroot and FJ models, both opinion and weight errors were minimal, demonstrating reliable learning across these frameworks. For the WHK model, performance depended on the choice of parameters such as confidence bounds; neutral or higher bounds produced better results. The trends were consistent across both real-world network topologies and random graphs. Smaller networks exhibited faster convergence and lower computation times, while larger networks required more resources but maintained steady accuracy in learning metrics. Parameter choices also influenced performance.

5. Conclusions and Future Research

In this work we considered the problem of selecting a set of seeds in a social network to start an information diffusion campaign and reduce polarization and disagreement in the opinions expressed by the individuals.

A potential extension of our framework may consider the direction of exploring other Opinion Formation Models, as well as other metrics. With respect to LocalGreedy, its main problem is the limit in term of scalability on significantly huge real networks, so future work could try to improve its performance with more sophisticated choices of the subset of nodes to be considered to run the model. Another possible idea could be exploiting community detection to refine the choice of the seed set, or even merge community detection and LocalGreedy to make the algorithm scaling to bigger networks. In addition, LocalGreedy is an approximation of the Kempe et al. [45] Greedy algorithm, but advanced and faster versions of it exists (CELF [49], CELF++ [50], ...), so it may be interesting to extend the LocalGreedy approach to these algorithms too. Another possible extension may be to perform both seeding and edge manipulation techniques at the same time, in order provide a better minimization of the metrics.

For the learning part, exploiting the generality of the online learning framework, it could be a good idea to find edge updates techniques requiring less execution time and/or having better quality in the real weights' approximation.

References

[1] D. Aineto, R. De Benedictis, M. Maratea, M. Mittelmann, G. Monaco, E. Scala, L. Serafini, I. Serina, F. Spegni, E. Tosello, A. Umbrico, M. Vallati (Eds.), Proceedings of the International Workshop on Artificial Intelligence for Climate Change, the Italian workshop on Planning and Scheduling, the RCRA Workshop on Experimental evaluation of algorithms for solving problems with combinatorial explosion, and the Workshop on Strategies, Prediction, Interaction, and Reasoning in Italy (AI4CC-IPS-RCRA-SPIRIT 2024), co-located with 23rd International Conference of the Italian Association for Artificial Intelligence (AIxIA 2024), CEUR Workshop Proceedings, CEUR-WS.org, 2024.

[2] F. Cascini, A. Pantovic, Y. A. Al-Ajlouni, G. Failla, V. Puleo, A. Melnyk, A. Lontano, W. Ricciardi, Social media and attitudes towards a COVID-19 vaccination: A systematic review of the literature, EClinicalMedicine 48 (2022). Publisher: Elsevier.

[3] S. Tu, S. Neumann, A Viral Marketing-Based Model For Opinion Dynamics in Online Social Networks, in: WWW '22: Proceedings of the ACM Web Conference 2022, WWW '22, Association for Computing Machinery, New York, NY, USA, 2022, pp. 1570–1578. Event-place: Virtual Event, Lyon, France.

[4] U. Chitra, C. Musco, Understanding Filter Bubbles and Polarization in Social Networks, 2019.

[5] M. H. Degroot, Reaching a Consensus, Journal of the American Statistical Association 69 (1974) 118–121. Publisher: Taylor & Francis.

[6] N. E. Friedkin, E. C. Johnsen, Social influence and opinions, The Journal of Mathematical Sociology 15 (1990) 193–206. Publisher: Routledge.

[7] D. Bindel, J. Kleinberg, S. Oren, How bad is forming your own opinion?, Games and Economic Behavior 92 (2015) 248–265.

[8] G. Deffuant, D. Neau, F. Amblard, G. Weisbuch, Mixing beliefs among interacting agents, Advances in Complex Systems 03 (2000) 87–98.

[9] R. Hegselmann, U. Krause, Opinion dynamics and bounded confidence: Models, analysis and simulation, Journal of Artificial Societies and Social Simulation 5 (2002) 1–24.

[10] C. Toccaceli, L. Milli, G. Rossetti, Opinion Dynamic Modeling of Fake News Perception, in: R. M. Benito, C. Cherifi, H. Cherifi, E. Moro, L. M. Rocha, M. Sales-Pardo (Eds.), Complex Networks & Their Applications IX, Springer International Publishing, 2021, pp. 370–381. Place: Cham.

[11] D. Ferraioli, P. W. Goldberg, C. Ventre, Decentralized dynamics for finite opinion games, Theoretical Computer Science 648 (2016) 96–115.

[12] F. Chierichetti, J. Kleinberg, S. Oren, On discrete preferences and coordination, in: Proceedings of the Fourteenth ACM Conference on Electronic Commerce, EC '13, Association for Computing Machinery, New York, NY, USA, 2018, pp. 233–250. Event-place: Philadelphia, Pennsylvania, USA.

[13] D. Fotakis, D. Palyvos-Giannas, S. Skoulakis, Opinion Dynamics with Local Interactions, in: IJCAI, 2016, pp. 279–285.

[14] D. Fotakis, V. Kandiros, V. Kontonis, S. Skoulakis, Opinion Dynamics with Limited Information, in: WINE, 2018, pp. 282–296.

[15] K. Bhawalkar, S. Gollapudi, K. Munagala, Coevolutionary Opinion Formation Games, in: STOC, 2013, pp. 41–50.
[16] V. Bilò, A. Fanelli, L. Moscardelli, Opinion Formation Games with Dynamic Social Influences, Theoretical Computer Science 746 (2018) 444–458.
[17] D. Ferraioli, C. Ventre, Social Pressure in Opinion Games, in: IJCAI, 2017, pp. 3661–3667.
[18] V. Auletta, A. Fanelli, D. Ferraioli, Consensus in Opinion Formation Processes in Fully Evolving Environments, in: AAAI, 2019, pp. 6022–6029.
[19] V. Auletta, I. Caragiannis, D. Ferraioli, C. Galdi, G. Persiano, Generalized Discrete Preference Games, in: IJCAI, 2016, pp. 53–59.
[20] E. Acar, G. Greco, M. Manna, Group Reasoning in Social Environments, in: AAMAS, 2017, pp. 1296–1304.
[21] V. Bilò, D. Ferraioli, C. Vinci, General Opinion Formation Games with Social Group Membership, in: L. D. Raedt (Ed.), Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22, 2022, pp. 88–94.
[22] A. Matakos, E. Terzi, P. Tsaparas, Measuring and moderating opinion polarization in social networks, Data Mining and Knowledge Discovery 31 (2017) 1480–1505.
[23] C. Musco, C. Musco, C. E. Tsourakakis, Minimizing Polarization and Disagreement in Social Networks, in: Proceedings of the 2018 World Wide Web Conference, WWW '18, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 2018, pp. 369–378. Event-place: Lyon, France.
[24] L. Zhu, Q. Bao, Z. Zhang, Minimizing Polarization and Disagreement in Social Networks via Link Recommendation, in: M. Ranzato, A. Beygelzimer, Y. Dauphin, P. S. Liang, J. W. Vaughan (Eds.), Advances in Neural Information Processing Systems, volume 34, Curran Associates, Inc., 2021, pp. 2072–2084.
[25] M. F. Chen, M. Z. Rácz, An Adversarial Model of Network Disruption: Maximizing Disagreement and Polarization in Social Networks, IEEE Transactions on Network Science and Engineering 9 (2022) 728–739.
[26] A. Gionis, E. Terzi, P. Tsaparas, Opinion maximization in social networks, in: Proceedings of the 2013 SIAM International Conference on Data Mining (SDM), 2013, pp. 387–395.
[27] V. Auletta, I. Caragiannis, D. Ferraioli, C. Galdi, G. Persiano, Minority Becomes Majority in Social Networks, in: WINE, 2015, pp. 74–88.
[28] V. Auletta, I. Caragiannis, D. Ferraioli, C. Galdi, G. Persiano, Information Retention in Heterogeneous Majority Dynamics, in: WINE, 2017, pp. 30–43.
[29] V. Auletta, D. Ferraioli, G. Greco, Reasoning about Consensus when Opinions Diffuse through Majority Dynamics, in: IJCAI, 2018, pp. 49–55.
[30] E. Yildiz, A. Ozdaglar, D. Acemoglu, A. Saberi, A. Scaglione, Binary Opinion Dynamics with Stubborn Agents, ACM Trans. Econ. Comput. 1 (2013). Place: New York, NY, USA Publisher: Association for Computing Machinery.
[31] J. Ghaderi, R. Srikant, Opinion dynamics in social networks with stubborn agents: Equilibrium and convergence rate, Automatica 50 (2014) 3209–3215.
[32] A. Silva, Opinion Manipulation in Social Networks, in: S. Lasaulce, T. Jimenez, E. Solan (Eds.), Network Games, Control, and Optimization, Springer International Publishing, 2017, pp. 187–198. Place: Cham.
[33] R. Bredereck, E. Elkind, Manipulating opinion diffusion in social networks, in: IJCAI International Joint Conference on Artificial Intelligence, International Joint Conferences on Artificial Intelligence, 2017, pp. 894–900.
[34] M. Castiglioni, D. Ferraioli, N. Gatti, G. Landriani, Election Manipulation on Social Networks: Seeding, Edge Removal, Edge Addition, Journal of Artificial Intelligence Research 71 (2021) 1049–1090.
[35] S. Sina, N. Hazon, A. Hassidim, S. Kraus, Adapting the social network to affect elections, in: Proceedings of the International Conference on Autonomous Agents and MultiAgent Systems (AAMAS), 2015, pp. 705–713.

[36] V. Auletta, D. Ferraioli, V. Savarese, Manipulating an election in social networks through link addition, J. Ambient Intell. Humaniz. Comput. 11 (2020) 4073–4088.

[37] V. Auletta, I. Caragiannis, D. Ferraioli, C. Galdi, G. Persiano, Robustness in discrete preference games, in: Proceedings of the International Conference on Autonomous Agents and MultiAgent Systems (AAMAS), 2017, pp. 1314–1322.

[38] V. Auletta, D. Ferraioli, V. Fionda, G. Greco, Maximizing the Spread of an Opinion when \emph-Tertium Datur Est, in: Proceedings of the International Conference on Autonomous Agents and MultiAgent Systems (AAMAS), 2019, pp. 1207–1215.

[39] V. Auletta, D. Ferraioli, G. Greco, On the Effectiveness of Social Proof Recommendations in Markets with Multiple Products, in: Proceedings of the the European Conference on Artificial Intelligence (ECAI), 2020, pp. 19–26.

[40] P. Domingos, M. Richardson, Mining the network value of customers, in: Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '01, Association for Computing Machinery, New York, NY, USA, 2001, pp. 57–66. Event-place: San Francisco, California.

[41] Y. Li, J. Fan, Y. Wang, K.-L. Tan, Influence Maximization on Social Graphs: A Survey, IEEE Transactions on Knowledge and Data Engineering 30 (2018) 1852–1872.

[42] S. Banerjee, M. Jenamani, D. K. Pratihar, A survey on influence maximization in a social network, Knowledge and Information Systems 62 (2020) 3417–3455.

[43] A. De, S. Bhattacharya, P. Bhattacharya, N. Ganguly, S. Chakrabarti, Learning a Linear Influence Model from Transient Opinion Dynamics, in: Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management, CIKM '14, Association for Computing Machinery, New York, NY, USA, 2014, pp. 401–410. Event-place: Shanghai, China.

[44] H.-T. Wai, A. Scaglione, A. Leshem, Active online learning of trusts in social networks, in: 2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2016, pp. 4139–4143.

[45] D. Kempe, J. Kleinberg, E. Tardos, Maximizing the spread of influence through a social network, in: Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '03, Association for Computing Machinery, New York, NY, USA, 2003, pp. 137–146. Event-place: Washington, D.C.

[46] J. Leskovec, A. Krevl, SNAP Datasets: Stanford Large Network Dataset Collection, 2014. URL: http://snap.stanford.edu/data.

[47] T. P. Peixoto, The netzschleuder network catalogue and repository, Zenodo10 5281 (2020).

[48] R. A. Rossi, N. K. Ahmed, The Network Data Repository with Interactive Graph Analytics and Visualization, in: AAAI, 2015. URL: https://networkrepository.com.

[49] J. Leskovec, A. Krause, C. Guestrin, C. Faloutsos, J. VanBriesen, N. Glance, Cost-effective outbreak detection in networks, in: Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '07, Association for Computing Machinery, New York, NY, USA, 2007, pp. 420–429. Event-place: San Jose, California, USA.

[50] A. Goyal, W. Lu, L. Lakshmanan, CELF++: Optimizing the greedy algorithm for influence maximization in social networks, in: Proceedings of the 20th International Conference Companion on World Wide Web, WWW 2011, volume 47-48, 2011, pp. 47–48.