Opinion Diffusion on Society Graphs Based on Approval Ballots

Extended Abstract

Jayakrishnan Madathil
University of Glasgow
Glasgow, United Kingdom
jayakrishnan.madathil@glasgow.ac.uk

Neeldhara Misra
Indian Institute of Technology, Gandhinagar
Gandhinagar, India
neeldhara.m@iitgn.ac.in

Yash More
Indian Institute of Technology, Gandhinagar
Gandhinagar, India
yash.mh@iitgn.ac.in

ABSTRACT

A society graph, as considered by [Faliszewski et al., IJCAI 2018], is a graph corresponding to an election instance where every possible ranking is a node, and the weight of such a node is given by the number of voters whose vote correspond to the said ranking. A natural diffusion process on this graph is defined, and an immediate question that emerges is whether there is a diffusion path that leads to a particular candidate winning according to a certain voting rule—this turns out to be NP-complete.

In this contribution, we consider the setting when votes are approval ballots, as opposed to rankings—and we consider both the possible and necessary winner problems. We demonstrate that it is possible to efficiently determine if a candidate is a possible winner (i.e, if there exists a diffusion path along which a given candidate wins the election) if the underlying society graph is a star (i.e, tree of diameter at most two), while the problem is NP-complete for trees of diameter d for $d > 2$. Analogously, we show that it is possible to efficiently determine if a candidate is a necessary winner (i.e, a winner for every possible diffusion path) if the underlying society graph is a star, while the problem is coNP-complete for trees of diameter d for $d > 2$. We also show the following results on structured graphs for the possible winner problem: the problem is strongly NP-complete on a disjoint union of paths, and on trees of constant diameter. We also report preliminary experiments from an ILP-based implementation.

KEYWORDS

opinion diffusion; integer linear programs; plurality; possible winner; necessary winner; NP-hardness; graphs

ACM Reference Format:

Jayakrishnan Madathil, Neeldhara Misra, and Yash More. 2024. Opinion Diffusion on Society Graphs Based on Approval Ballots: Extended Abstract. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 3 pages.

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

1 INTRODUCTION

We consider the interplay of opinion diffusion and elections outcomes in the context of elections held using approval ballots. An election is most naturally modeled as a profile of voter opinions over a set of candidates. With increasing communication between voters (e.g, peer-to-peer discussions on social media platforms) and between candidates and voters (e.g, coordinated campaigns), the final profile representing votes that are submitted is typically not consistent with the original opinions of the voters—these opinions have likely evolved over discourse related to the election process, and we would like to understand how this evolution impacts election outcomes.

For instance, it is conceivable that because of the way voters influence each others' opinions, the winner with respect to the original voting profile may be quite different from the winner with respect to the profile that manifests after the convergence of some appropriate diffusion process [8] that models the evolution of individual opinions. This prompts natural questions in the spirit of election "control" [3, 6, 9], where one is interested in knowing if the diffusion process can be influenced to engineer a specific outcome, either constructive (where the goal is to make a specific candidate win), or destructive (where the goal is to ensure that a particular candidate does not win).

A natural way to understand how voter opinions evolve based on how they are influenced by other voters is to associate a graph with the set of voters—typically a voter is adjacent, in this graph, to other voters who she is most likely to be influenced by. One view might be that a voter is likely to be influenced by her close and trusted friends. This leads the "social network" model, where a voter is adjacent to people she knows [1, 2]. However, one might argue that not all friends hold the same influence over an individual, and it would typically be challenging to solicit accurate data about how people influence each other through interpersonal relationships.

The other approach is to say that people are influenced by other people who hold similar *views*, even if they do not know them socially. Indeed, it is common for people to subscribe to certain ideologies based on targeted campaigns by vocal supporters of said ideologies, even if they have no direct association with the campaigners. This is naturally modeled using what are called "society graphs," where we have nodes corresponding to "worldviews" (to be more specific, every possible ballot is represented by a vertex), and all voters who subscribe to a particular worldview (as represented by a ballot) are effectively mapped to the node representing said worldview [4, 7]. We then have that a pair of nodes are adjacent if the ballots are similar. One of the original motivations for this model is that the edges represent the fact that voters from various clusters are very likely to interact with each other due to the similarity of their views. For instance, like-minded voters tend to follow similar social media personalities and because of the nature of discovery algorithms on social media platforms, they see "more of the same," leading to interactions between the voters over these common sources of information.

This model is particularly useful in the context of large-scale elections, where information about friendships are other relationships capturing scope of influence may be scarce, and it would be reasonable to assume that people who have similar opinions are most likely to influence each other provided the numbers are appropriately substantial. It is also a useful view from the perspective of campaigners, who naturally organize their outreach around groups of like-minded people, whether or not they are socially connected.

Related Work. The notion of society graphs was introduced by Knop et al. [7]. Our work is inspired by problems studied by Faliszewski et al. [4, 5], where similar problems are studied in the context of society graphs based on rankings. In a society graph based on rankings, every ranking of candidates is a vertex, and there is an edge between a pair of vertices corresponding to rankings p and q if it is possible to transform p into q with a single swap of two consecutive candidates. They obtain algorithms that are FPT in the number of candidates for the problems of determining possible winners and the feasibility of certain kinds of bribery, that work for all ILP-expressible voting rules.

2 OUR CONTRIBUTION

Our main contribution is to extend the society graph model to the setting of approval ballots.

Our Model. We consider elections based on approval ballots. An election consists of a set $\mathcal{C}$ of m candidates and set $\mathcal{V}$ of n voters, and for each voter $v \in \mathcal{V}$, v's vote is given by $S_v \subseteq \mathcal{C}$, the set of candidates v approves. Let $w : 2^{\mathcal{C}} \to \mathbb{N} \cup \{0\}$ be the function defined as follows: for $S \subseteq \mathcal{C}$, $w(S)$ is the number of voters v such that $S = S_v$; that is, $w(S)$ is the number of voters whose set of approved candidates is precisely S. We call $w(S)$ the weight of S, and we call w the weight function. Notice that as each candidate approves exactly one subset, we have $\Sigma_{S \subseteq \mathcal{C}} w(S) = n$, the number of voters. For a candidate $c \in \mathcal{C}$, the number of votes v has is $\Sigma_{S \subseteq \mathcal{C}; c \in S} w(S)$, and we denote this quantity by $\mathsf{sc}(c)$, the score of c. We say that a candidate c is the winner (resp. a co-winner) of the election if $\mathsf{sc}(c) > \mathsf{sc}(c')$ (resp. $\mathsf{sc}(c) \geqslant \mathsf{sc}(c')$) for every $c' \in \mathcal{C} \setminus \{c\}$. The *society graph* associated with an election is the graph G with vertex set consists of all subsets S of $\mathcal{C}$ with $w(S) > 0$, and two vertices S and T of G are adjacent if and only if $|S \triangle T| = 1$, i.e., the the sets S and T differ by exactly one candidate. Notice that an election is completely described by the society graph G and the weight function w.

Diffusion Process. We consider a diffusion process on G, which works as follows. The atomic unit of the diffusion process is an *update step*, by which we mean the following: if a vertex S has a neighbor T such that $w(T)$ is greater than the weight of all the other neighbors of S plus S's own weight, then S would gravitate towards T. Formally, for a vertex S, an update step with respect to S is performed if S has a neighbor T such that

$$w(T) > w(S) + \Sigma_{S' \in N(S) \setminus \{T\}} w(S'), \tag{1}$$

where $N(S)$ is the set of neighbors of S. And at the end of such an update step, the weight of S is updated to 0 and the weight of T is updated to $w(T) + w(S)$. We say that the society graph G is stable if no update step can be performed; that is, Equation 1 does not hold for any pair of vertices S and T of G. By a diffusion process on G, we mean a sequence of update steps that renders G stable in the end. Notice that either G is stable, or at least one diffusion process exists (and that every diffusion process has at most |G| update steps, where G is the number of vertices of G).

Computational Problems. We are interested in the problems of the following type. (a) The Possible Winner problem (resp. the Possible Co-Winner problem): Given an election specified by its society graph G and the weight function w and a preferred candidate $p \in \mathcal{C}$, is there is a diffusion process on G such that p is the winner (resp. co-winner) at the end of the diffusion process? (b) The Necessary Winner problem (resp. the Necessary Co-Winner problem): Given an election specified by its society graph G and the weight function w and a preferred candidate $p \in \mathcal{C}$, is p the winner (resp. co-winner) at the end of *every* diffusion process on G?

Our Complexity Results. We prove a number of algorithmic and hardness results, and in particular, establish a complexity dichotomy for the Possible Winner and the Necessary Co-Winner problems when the society graph is a tree. Our results are as follows.

- Possible Winner can be solved in polynomial time if the society graph is a tree of diameter two, and is NP-hard on society graphs of diameter three.
- Necessary Co-Winner can be solved in polynomial time if the society graph is a tree of diameter two, and is coNP-hard on society graphs of diameter three.
- Possible Winner is strongly NP-complete even when the society graph is a tree of constant diameter.
- Possible Winner is strongly NP-complete even when the society graph is a disjoint union of paths.

Our Experimental Results. We implemented an ILP-based approach for the Possible Winner and Possible Co-Winner problems, and we document some preliminary experiments performed using Gurobi, a widely available ILP solver. Intuitively, what we observed was indicative of the diffusion not causing too much damage—in the sense that although different pathways do lead to different candidates winning on specific instances, it is not the case that a large number of winners emerge simply from the diffusion process. We note that our ILP formulation is flexible and can account for several variations of the base model.

ACKNOWLEDGMENTS

Jayakrishnan Madathil is supported by the Engineering and Physical Sciences Research Council [EP/V032305/1], and Neeldhara Misra is supported by the Science and Engineering Research Board [ECR/2018/002967]. All authors are grateful for support from IIT Gandhinagar.

REFERENCES

[1] Robert Bredereck and Edith Elkind. 2017. Manipulating Opinion Diffusion in Social Networks.. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI*. 894–900.

[2] Markus Brill, Edith Elkind, Ulle Endriss, and Umberto Grandi. 2016. Pairwise diffusion of preference rankings in social networks. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*. 130–136.

[3] Federico Corò, Emilio Cruciani, Gianlorenzo D'Angelo, and Stefano Ponziani. 2019. Vote for me! Election control via social influence in arbitrary scoring rule voting systems. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*. 1895–1897.

[4] Piotr Faliszewski, Rica Gonen, Martin Kouteckỳ, and Nimrod Talmon. 2018. Opinion Diffusion and Campaigning on Society Graphs. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI*.

[5] Piotr Faliszewski, Rica Gonen, Martin Koutecký, and Nimrod Talmon. 2020. Opinion Diffusion and Campaigning on Society Graphs. *CoRR* abs/2010.00651 (2020). arXiv:2010.00651 https://arxiv.org/abs/2010.00651

[6] Piotr Faliszewski and Jörg Rothe. 2016. Control and Bribery in Voting. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, 146–168.

[7] Dušan Knop, Martin Kouteckỳ, and Matthias Mnich. 2018. A Unifying Framework for Manipulation Problems. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*. 256–264.

[8] Paulo Shakarian, Abhivav Bhatnagar, Ashkan Aleali, Elham Shaabani, Ruocheng Guo, et al. 2015. *Diffusion in social networks*. Springer.

[9] Bryan Wilder and Yevgeniy Vorobeychik. 2018. Controlling Elections through Social Influence. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*. 265–273.