Frankenmandering: Repeated Social Graph Gerrymandering

Sahil Agarwal, Svetlana Obraztsova , Zinovi Rabinovich and Alan Tsang
Affiliation: Carleton University, Ottawa
Email: SahilAgarwal@cmail.carleton.ca,
{svetlana.obraztsova,zinovi.rabinovich,alan.tsang}@carleton.ca

Abstract

Opinion influence by means of manipulating a social graph structure, as well as election manipulation by means of artificial grouping (gerrymandering), are well-known fields of application and investigation in computational social choice and machine learning. However, while studied separately, in real life, both of these manipulations (social graph alteration and gerrymandering) occur simultaneously. In this paper, we offer the first model of such a simultaneous process, which takes the form of repeated gerrymandering with an underlying social graph for opinion diffusion. We term this process "Frankenmandering", and provide the first steps in its analysis: examples of principal feasibility and the impact of the underlying social graph.

1 Introduction

Influencing the spread of opinion in a social network has been studied from two main perspectives: forcing changes in network topology [35, 24, 4], or placing "influencers" in the network [5, 29]. Former can be deployed by social network management, while the latter is an open market. Fortunately, neither is too easy [11], though possible [39, 6, 36, 12].

However, in district-based election schemes, political parties have found another dubious tactic: Gerrymandering – altering the grouping of voters to skew the representative's election process. It has been studied in practical political terms [2, 37, 27] and more theoretical terms [15, 9, 26]. Strangely, Gerrymandering remains (as a task) a single-shot process, a way to exploit the existing set of opinions, rather than means to both exploit and *form* them. This is in spite of the existence of some studies in political science (e.g., [30]) of the long-term effects of gerrymandering. Doubly strange, since the question of grouping people for forced interaction, targetting their opinions in an elections context, has been raised (e.g., [31]).

In this paper, we bridge the gap and show that gerrymandering can be effectively used as an opinion *formation* tool with *long-term* electoral effects. Specifically, we link gerrymandering with the formation (or modification) of social opinion networks, and thus use it as a *long-term* opinion control method. We term this new model of repeated gerrymandering process over social networks **"Frankenmandering"**.

2 Model of Interaction

Given a set of n voters V, each characterised by a position in some "physical" space $p_v \in \mathbb{R}^d$ and an opinion $c_v \in \mathbb{R}^m$. The physical location, p_v, of voters is used to restrict their grouping into districts. We will use bold-face $\mathbf{c} \in (\mathbb{R}^m)^n$ to denote the vector of all opinions of all voters, and use functional form $c_v(l)$ to refer to the l'th coordinate of the preference vector c_v of voter $v \in V$. In addition, for a subset $D \subset V$, we will denote the sub-vector of opinions within the subset by $\mathbf{c}|_D = \{c_v\}_{v \in D} \in (\mathbb{R}^m)^{|D|}$. Let there also be a directed (social) graph $G = (V, E, w)$ with $w : E \to \mathbb{R}$ being an edge weight function. In addition, let $\mathbb{L}[G] : (\mathbb{R}^m)^n \to (\mathbb{R}^m)^n$ be a graph-parameterised opinion dynamics function, and

$d : (\mathbb{R}^m)^n \to \mathbb{R}$ a social opinion evaluation function. Finally, let there be a representative selection function $\mathbb{F}$, so that for any $\widetilde{V} \subseteq V$, $\mathbb{F}\left(\mathbf{c}|_{\widetilde{V}}\right) \in \widetilde{V}$.

Set $\mathbf{c}^0$ to be the initial opinion of all voters, and consider the following process iterated for every time period t:

1. Districts, $\mathbf{D}^t = \{D_j^t \subseteq V\}_{j=1}^K$, are being selected, so that $\bigcup_{j=1}^{K} D_i^t = V$, and $\forall i, j$ holds $D_i^t \cap D_j^t = \emptyset$;
2. Local elections are run, producing district representatives $r_j^t = \mathbb{F}\left(\mathbf{c^t}|_{D_j^t}\right) \in D_j^t, j \in \{1, ..., K\}$;
3. Representatives become local "influencers". I.e., Graphs, $\{H_j^t = (D_j^t, E_j^t, w_j^t)\}_{j=1}^K$, are built so that $(r_j^t, v) \in E_j^t$ for all $v \in D_j^t$, and $w_j^t : E_j^t \to \mathbb{R}$ reflect the influence of the chosen representative on the "constituents" of its district;
4. Social influence is then exercised on all opinions via $\mathbb{L}^t = \mathbb{L}[G \cup \bigcup_{j=1}^{K} H_j^t]$, so that : $\mathbf{c}^t = \mathbb{L}^t(\mathbf{c}^{t-1})$

2.1 District Geometry and Initial Opinion Distribution

In our initial studies, we will assume that voters are placed on a regular (2D) grid, though they do not necessarily *form* that grid. This initial simplification pursues several goals.

First, such a positioning will simulate well the real-life geographical maps used in applied gerrymandering. Such maps are commonly reduced to planar graphs, and we will have the benefit of prior work on algorithmic gerrymandering as experimentation baselines (e.g., [15, 14]). Second, grids and planar graphs are a convenient medium for the application of spatial distribution models, such as Gaussian and Dirichlet Processes [16, 19, 33], to capture and track the distribution of opinions in geographical space and over the social network.

Second, speaking of the "geography" and voter's "physical" locations $p_v \in \mathbb{R}^d$, we must note that it is common practice to restrict district formation. Specifically, it is required that there are contiguous non-overlapping regions of the location space $A_j \subset \mathbb{R}^d$, $j \in \{1, ..., K\}$, so that $D_j \subset A_j$. Further restrictions come in the form of an *exclusion* area $A^{\emptyset} \subset \mathbb{R}^d$, and the requirement that no district-forming area A_j intersects with $A^{\emptyset}$. The exclusion area dictates geographic limitations, such as a lake or a river, that a district cannot cross/bridge.

The importance of limiting the choice of possible districts cannot be understated, as it is a distinguishing feature of Gerrymandering from Partition Control of elections (see, e.g., [7, 10], and the tremendous volume of citations therein). Both Gerrymandering and Partition Control strive to exploit the splitting of voters into (similarly sized) sub-groups. Partition Control even takes a deeper look at multi-stage and on-line voting processes, but it is Gerrymandering that explicitly focuses on the (geographic/topological) limitations of possible voter groupings.

In practice, following a classical robotics trick, geographic limitations tend to be resolved by "tessellating" the *permissible* space of locations, and then forming districts by collecting connected "shards". This can further be reduced to a complementary graph of the tessellation, where each shard is represented by the graph's node and each edge represents a side shared by two shards. This naturally leads to redistricting being reduced to graph separation, as in, e.g. [14, 15]. In our feasibility study in Section 4 we adopt this view of redistricting as well, making sure that the "geography" graph is planar.

However, in out formal model we maintain the "physical" locations, as they may be correlated with opinions and, thus, exploited for Gerrymandering with partial opinion information.

2.2 Local Representative Election

The grand formalism of The Model does not limit the manner in which a district representative is chosen, but merely suggests that it exists and it is uniform for all districts. However, to begin our analysis of The Model we focus on metric/spatial social choice. It is a long-standing theoretical and practical modelling technique [17, 32, 34]. More recent research even provides some bounds and assurances as to the quality of the outcome (e.g., [3, 1], with the former explicitly targetting district-based voting) and computational feasibility under uncertainty (e.g., [20]).More specifically, we assume a "median" representative is elected:

$$r_j = \arg\min_{l^* \in D_j} \sum_{\lambda \in D_j} \|c_l - c_\lambda\|$$

2.3 Opinion Dynamics

There are multiple opinion dynamics we can target, starting from the (weighted) Ising model and to aggregate dynamics with backlash/backfire [22, 25]. The latter is of particular interest here.

During an interaction of two voters $u, v \in V$, connected by an edge $e = (uv) \in E$, the opinion c_v will undergo a change depending on the relative proximity of c_u:

$$c_v^{t+1} = \begin{cases} c_v^t + \mu^+(c_u^t - c_v^t) & \|c_v - c_u\| < \epsilon^+ \\ c_v^t - \mu^-(c_u^t - c_v^t) & \|c_v - c_u\| \geq \epsilon^- \\ c_v^t & otherwise \end{cases}$$

where $\epsilon^- >> \epsilon^+$ are, respectively, the "backfire" and "confirmation" thresholds, and μ^+, μ^- are the corresponding degrees of sensitivity. Sensitivity arguments can be a function of the general edge weight $w(e)$ as described by the social network graph $G = (V, E, w)$. These effects are aggregated (and the influence coefficients are possibly normalised) across all edges.

Now, in [22, 25], from which we borrow the "absorption/backlash" response functions, the opinions c_v are scalar values in $[0, 1]$ or $[-1, 1]$. However, in our model optinions are multi-dimensional, so that $c_v \in R^m$ with $m > 1$ and not necessarily have finite range. Essentially, we see them as multi-issue opinions, interpreting each dimension of the opinion vector c_v as a (logit) stance on that issue.

The simplest way to expand the "absorption/backlash" response to multi-issue opinions would be to apply the dynamic function above per opinion coordinate. In addition, sensitivity coefficients can be correlated with the overall (normalised) distance of opinions, so that overall close opinions will depress the backfire sensitivity, while the overall distant opinions will depress the confirmation coefficients. In fact, borrowing some inspiration from [13], the dynamic may be written as:

$$c_v^{t+1} = c_v^t + \mu_{uv}^t(c_u^t - c_v^t),$$

where $\mu^t \propto w_{uv}\langle c_u^t, c_v^t\rangle$, so that the update coefficient is proportional to: a) the strength/magnitude of the individual opinions; b) their alignment ($\cos$ of the angle between the opinion vectors, or another distance); and c) the strength of the influence, w_{uv} of voter u on voter v. However, in this paper, we will be using an even more generic form of neighbourhood influence aggregation, formally described as follows.

Let us define for a $c \in \mathbb{R}^m$ a unit vector with the same direction $\vec{\mathbb{1}}(c) = \frac{1}{\|c\|}c$, extending this function so that $\vec{\mathbb{1}}(\vec{0}) = \vec{0}$. Let $\mu : \mathbb{R} \to \mathbb{R}$ be the *discrepancy response function (DRF)*, which we presume to be bounded so that there are $\mu^- < \mu^+ \in \mathbb{R}$ so that $\mu^- \leq \mu(x) \leq \mu^+$. Now, given a (directed) social graph $G = (V, E, w)$, let us denote the *influencing* and *following* neighbours of a voter $v \in V$ by

$N^{\downarrow}(v) = \{u \in V | (uv) \in E\}$ and, respectively $N^{\uparrow}(v) = \{u \in V | (vu) \in E\}$ Then, opinion dynamics over a (directed) social graph $G = (V, E, w)$ are expressed by:

$$c_v^{t+1} = c_v^t + \frac{1}{Z} \sum_{u \in N^{\downarrow}(v)} \mu\left(\|(c_u^t - c_v^t)\|\right) w_{uv} \vec{\mathbb{1}}(c_u^t - c_v^t),$$

where Z is a normalisation factor that controls sensitivity to the cardinality of a voter's neighbourhood. So, for example, $Z \propto \sum_{u \in N^{\downarrow}(v)} w_{uv}$ normalises the relative importance of neighbours without regard to their number, while $Z \propto \sum_{u \in N^{\downarrow}(v)} w_{uv}\mu\left(\|(c_u^t - c_v^t)\|\right)$ also normalises the response to the discrepancy in the neighbours opinion.

DRF is what dictates whether a voter is attracted or repulsed by the opinion of its neighbour(s). We will be using DRF with two distinct phases of *similarity bias* (or "assimilation"), wherein somewhat different but close opinions are attractive, and *backfire*, wherein opinions that are too distinct to a voter will be repulsive. Furthermore, we will allow for three neutral ranges: "indifference" (wherein voters deem the difference in opinion inconsequential), "ambivalence" (wherein voters choose a positive or negative reaction) and "irrelevance" (where the opinion of the other voter is so distinct that it is disregarded). Figure 1 depicts the overall shape of the DRFs that we will employ.

Figure 1: Discrepancy Response Function used in Section 4.2

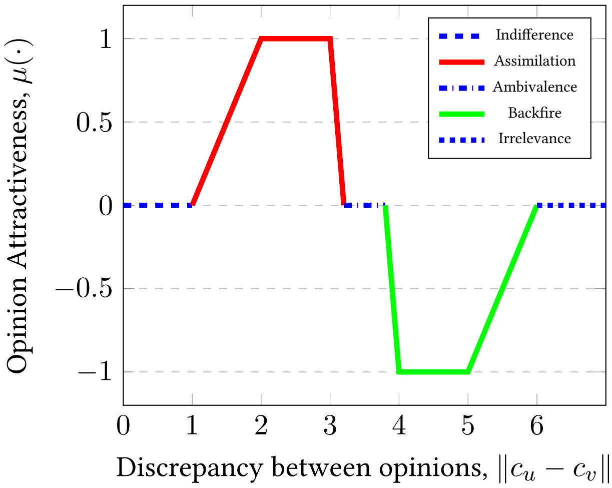

3 Creating Frankenmandering

The Model only describes the interaction between a social graph opinion dynamics and a redistricting process. But on it own, The Model lacks purpose. To develop one we expand the concept of Gerrymandering [21, 18]– a *single instance* redistricting aimed at ensuring that representatives hold a specific type of opinion.

3.1 Gerrymandering

To begin, let us recast the classical Gerrymandering problem in terms and formalism of our Model. In this case, the opinions represent a preference order over m candidates, and the representative is there to support the aggregate preference order of its district. The winner of an election is determined by aggregating the preferences of the representatives, and Gerrymandering targets this final (single-shot) outcome.

E.g., the following procedure may be followed. Let a (plurality) ballot projection $\beta : \mathbb{R}^m \to \mathbb{R}^m$ such that if $b_v = \beta(c_v)$ for some $v \in V$ then b_v is a one-hot vector and $b_v(l^*) = 1$ implies $l^* \in \arg\max_{i\in[1:m]} c_v(i)$.

Then the representative for district D_j is defined by

$$r_j = \arg\max_{r\in D_j} \left\langle \beta(c_r), \sum_{v\in D_j} \beta(c_v) \right\rangle$$

That is the ballot of the representative corresponds to choosing the candidate with the highest number of plurality votes in the district.

A constructive Gerrymandering algorithm then accepts a desired candidate $l^* \in [1 : m]$, and seeks re-districting, $\{D_j\}_{j=1}^K$, that makes l^* a winner. I.e., $l^* = \arg\max_{i\in[1:m]} b(i)$, where $b = \sum_{j=1}^{K} b_r$.

A destructive Gerrymandering algorithm seeks a re-districting so that $b(l^*) < \max_{i\in[1:m]} b(i)$, preventing l^* from winning the election.

The softer versions of the above also exist, where redistricting seeks to maximise (or minimise) the number of votes received by the idealised candidate.

3.2 Gerrymandering a Dynamic Voter Base

Noticeably, Gerrymandering aims at a static voter base, i.e., it is a single shot operation whose after-effects on the voter's opinions are not taken into account. In fact, we did not find any work on *repeatedly* gerrymandering a population of voters, though some works (e.g., [30]) in political science do discuss whether the fact of gerrymandering influences such population metrics as polarisation.

In this work, we show that (repeated) re-districting can be used as a control mechanism of opinion dynamics, thus having long-term and, potentially, permanent effect on future elections. We term such re-districting: **Frankenmandering**.

Frankenmandering accepts as input some desired "reference"/ideal opinion c^*, which it then seeks to spread throughout the population within a given number, T, of strategic re-districtings. We define **Strong (or Populist) Frankenmandering** as minimising the total distance from actual voter opinions to the ideal one. Formally, it is captured by the following optimisation problem:

$$\begin{aligned}
\min_{\mathbf{D}^1,\ldots,\mathbf{D}^T} \quad & \sum_{v\in V} \|c^* - c_v^T\| \\
& s.t. \\
\forall j \in [1:K]\ \forall t \in [1:T] \quad & \\
& r_j^t = \mathbb{F}\left(\mathbf{c^t}|_{D_j^t}\right) \\
& H_j^t = (D_j^t, E_j^t, w_j^t) \\
& \mathbb{L}^t = \mathbb{L}[G \cup \bigcup_{j=1}^{K} H_j^t] \\
& \mathbf{c}^t = \mathbb{L}^t(\mathbf{c}^{t-1})
\end{aligned}$$

Frankenmandering is a control problem: it optimises opinion diffusion over a social network via (restricted) graph connectivity control. In its naive form above it poses a much stronger requirement on the final outcome than Gerrymandering at time $t = T$ would do. However, it does have a more Gerrymandering-like variant, which we term **Senat Frankenmandering**, where the objective function targets the representative's opinion at time $t = T$:

$$\min_{\mathbf{D}^1,\ldots,\mathbf{D}^T} \quad \sum_{j=1}^{K} \|c^* - \mathbf{c}^T(r_j^T)\|$$

4 Frankenmandering Feasibility

We would like to present two hand-crafted examples. The first example is to show that Frankenmandering, as a repeated extension of Gerrymandering, is possible. That is, there exist instances of opinion profiles where repeated redistricting leads to a consistent shift in opinions. The second is to show that a social network can amplify even a single-step Gerrymandering solution into a Frankenmandering monster. That is, even a single redistricting (maintained for a long time) can lead to persistent and consistent opinion shift.

4.1 Frankenmandering Inchworm

In this first example, we demonstrate that by carefully selecting the districts $\mathbf{D}^t$ at each step, we can cause all opinions of a population to indefinitely shift in the positive direction. Our example relies on planar graph adjacency of voters, as depicted in Figure 3. This corresponds to a realistic geography, where districts are contiguous. At the same time, we assume an empty social network graph G ("a city of strangers").

We begin with the following profile of opinions for $n = 10$ voters: $\{0, 0, 0, 1, 2, 3, 4, 5, 5, 5\}$. In each iteration, we create a single district d^t of size 3, and size-1 districts for all other voters. The representative (elected median voter) then exerts influence on the two other voters in d^t according to the following update rule, depicted for a single voter-influencer pair (u, v) with opinions c_v, c_u:

$$c_v^{t+1} = \begin{cases} c_v^t + \text{sign}(c_u^t - c_v^t) & \|c_v - c_u\| < 3 \\ c_v^t - \text{sign}(c_u^t - c_v^t) & \|c_v - c_u\| \geq 3 \end{cases}$$

Frankenmandering then seeks to select a sequence of districts d^t such that we shift the entire profile of voter opinions by exactly $+1$; i.e. we end up with the target opinion profile $\{1, 1, 1, 2, 3, 4, 5, 6, 6, 6\}$. The intuition behind the solution is that re-districting proceeds in two phases: The first phase uses the backfire effect to "push" the most positive voters away from the main body while keeping the median voter close to the least positive voters to attract them toward a central value. The second phase selects increasingly higher opinion median voters to "pull" voters toward increasingly positive values. While this shift occurs, it is important to keep a "ladder" of voters with intermediate opinion values that can be used as median voters. The movement of voters is reminiscent of the locomotion of the inchworm.

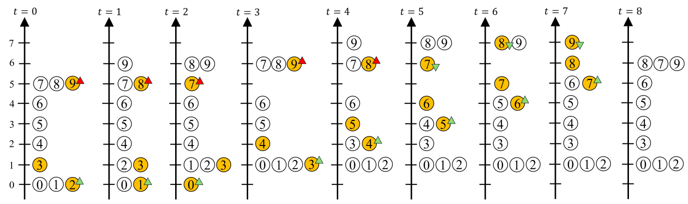

Figure 2: A Sequence of Gerrymandered Districts to Shift All Voters

Figure 2 illustrates a sequence of arrow diagrams that depicts the selection of districts that causes all opinions to shift exactly $+1$. In each arrow diagram, each voter is positioned (vertically) according to her opinion. The sequence of arrow diagrams proceeds from left to right. In each iteration t, the 3 voters shaded in orange are selected to form the district d^t. The median voter becomes the representative, whose opinion remains uninfluenced. The voters' opinions may shift, and these shifts are depicted as a small arrow: green arrows denote an attractive effect toward the representative, and red arrows denote the backfire effect pushing the voter away from the representative. These shifts are then reflected at the next arrow diagram at $t+1$. The final ($t = 8$) diagram reproduces the initial ($t = 0$) opinion profile, with all opinions shifted $+1$. The same pattern can then be repeated to shift the entire population's opinions indefinitely.

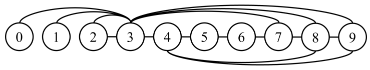

Figure 3: "Geography" Graph of Voters in Fig 2

4.2 Social Network Effect Example

In this second example, we show that we can achieve the same effect of shifting the opinions of the entire population with a single, persistent districting of voters, *if* the social influence between the voters can be leveraged. Let us define an initial opinion profile on $n = 6$ voters of $\{0, 1, 2, 3, 4, 6\}$, and a target opinion profile of $\{1, 2, 3, 4, 5, 7\}$. We will fix a district d^* to include exactly the first two voters and the last voter. The social network will be a line graph, with each voter connected to (up to) two most similar peers. Finally, the opinion dynamics have DRF as in Figure 1 or, more explicitly:

$$c_v^{t+1} = \begin{cases} c_v^t & \|c_v - c_u\| < 2 \\ c_v^t + \text{sign}(c_u^t - c_v^t) & \|c_v - c_u\| < 4 \\ c_v^t - \text{sign}(c_u^t - c_v^t) & \|c_v - c_u\| < 6 \\ c_v^t & \|c_v - c_u\| \geq 6 \end{cases}$$

Note the explicit use of the "indifference" and "irrelevance" plateaus of the DRF in this example.

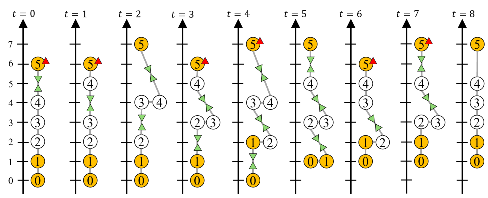

Figure 4: A Fixed Gerrymandered District to Shift All Voters

The sequence of arrow diagrams in Figure 4 shows how this is accomplished. As in Figure 2, the vertical positions in each arrow diagram denote the opinions of the voters. The social network among the voters is depicted as a gray line graph. Only voters 0, 1, and 5 form the fixed district d. Influence from the representative is depicted by small arrows as before (in this case, these are exactly the red arrows). Social influence between voters through the social network are drawn directly on the network (in this case, these are exactly the green arrows). By $t = 8$, the same opinion profile is regenerated, with all opinions shifted $+1$ from $t = 0$. With no further intervention, this system will shift opinions of all voters indefinitely in the positive direction.

5 Discussion and Future Work

In this paper, we present a novel interaction model of repeated redistricting with an underlying social network. The redistricting forces some nodes to become temporary opinion influencers in the social network, wherein the opinions undergo a non-trivial dynamic with "assimilation" and "backfire" properties. We term the strategic use of such redistricting **Frankenmandering**, and formulate an opinion control problem that corresponds to it. We show the feasibility of Frankenmanding by providing two key examples, where the opinions essentially "inchworm" in the desired direction.

We are building two algorithmic solutions to Frankenmandering. First, we are developing a reinforcement learning based solution. In spite of ostensibly high complexity, deep learning solutions of influencer placing in an opinion network do exist, e.g. [12, 8, 38]. Our current designs follow these and other approaches capable of some reasoning over graph embeddings [23, 28]. Second, we are considering an ad-hoc algorithm based on the "inchworm" solution for more general graphs by reordering the "geography" and the social graphs to align with the order implied by the geodesic distance of voter opinions from the ideal, target opinion c^*.

We submit our model and preliminary findings to the community, seeking comments, critique, and cooperation to facilitate our progress.

References

[1] Ben Abramowitz, Elliot Anshelevich, and Wennan Zhu. Awareness of voter passion greatly improves the distortion of metric social choice. In *Proceedings of the 15th International Conference on Web and Internet Economics (WINE)*, pages 3–16, 2019.

[2] Micah Altman and Michael McDonald. Redistricting and polarization. *American Gridlock: e Sources, Character, and Impact of Political Polarization*, pages 103–131, 2015.

[3] Elliot Anshelevich, Aris Filos-Ratsikas, and Alexandros A Voudouris. The distortion of distributed metric social choice. *Artificial intelligence*, 308:103713, 2022.

[4] Vincenzo Auletta, Diodato Ferraioli, and Vincenzo Savarese. Manipulating an election in social networks through link addition. *Journal of Ambient Intelligence and Humanized Computing*, 11: 4073–4088, 2020.

[5] Vincenzo Auletta, Francesco Carbone, and Diodato Ferraioli. Election manipulation in social networks with single-peaked agents. In *International Conference of the Italian Association for Artificial Intelligence*, pages 467–480. Springer, 2023.

[6] Vincenzo Auletta, Diodato Ferraioli, and Carmine Viscito. Election manipulation on social networks with abstention. In *Multi-Agent Systems*, pages 435–444, 2023.

[7] Anna Blokhina and Irina Samoilova. Electoral control by partition voters using the schulze method. In *2024 17th International Conference on Management of Large-Scale System Development (MLSD)*, pages 1–5, 2024. doi: 10.1109/MLSD61779.2024.10739529.

[8] Vivek S Borkar and Alexandre Reiffers-Masson. Opinion shaping in social networks using reinforcement learning. *IEEE Transactions on Control of Network Systems*, 9(3):1305–1316, 2021.

[9] Allan Borodin, Omer Lev, Nisarg Shah, and Tyrone Strangway. Little house (seat) on the prairie: Compactness, gerrymandering, and population distribution. In *AAMAS*, volume 2022, pages 154–162, 2022.

[10] Benjamin Carleton, Michael C Chavrimootoo, Lane A Hemaspaandra, David E Narváez, Conor Taliancich, and Henry B Welles. Separating and collapsing electoral control types. *Journal of Artificial Intelligence Research*, 81:71–116, 2024.

[11] Matteo Castiglioni, Diodato Ferraioli, Nicola Gatti, and Giulia Landriani. Election manipulation on social networks: seeding, edge removal, edge addition. *Journal of Artificial Intelligence Research*, 71:1049–1090, 2021.

[12] Haipeng Chen, Bryan Wilder, Wei Qiu, Bo An, Eric Rice, and Milind Tambe. Complex contagion influence maximization: A reinforcement learning approach. In *IJCAI*, pages 5531–5540, 2023.

[13] Xi Chen, Panayiotis Tsaparas, Jefrey Lijffijt, and Tijl De Bie. Opinion dynamics with backfire effect and biased assimilation. *PloS one*, 16(9):e0256922, 2021.

[14] Amittai Cohen-Zemach, Yoad Lewenberg, and Jeffrey S Rosenschein. Gerrymandering over graphs. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pages 274–282, 2018.

[15] Jack Dippel, Max Dupré la Tour, April Niu, Sanjukta Roy, and Adrian Vetta. Gerrymandering planar graphs. In *AAMAS*, page 463–471, 2024. ISBN 9798400704864.

[16] Jason A Duan, Michele Guindani, and Alan E Gelfand. Generalized spatial dirichlet process models. *Biometrika*, 94(4):809–825, 2007.

[17] James M Enelow and Melvin J Hinich. *The spatial theory of voting: An introduction*. CUP Archive, 1984.

[18] Robert S Erikson. Malapportionment, gerrymandering, and party fortunes in congressional elections. *American Political Science Review*, 66(4):1234–1245, 1972.

[19] Alan E Gelfand, Athanasios Kottas, and Steven N MacEachern. Bayesian nonparametric spatial modeling with dirichlet process mixing. *Journal of the American Statistical Association*, 100(471): 1021–1035, 2005.

[20] Aviram Imber, Jonas Israel, Markus Brill, Hadas Shachnai, and Benny Kimelfeld. Spatial voting with incomplete voter information. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 9790–9797, 2024.

[21] Samuel Issacharoff. Gerrymandering and political cartels. *Harv. L. Rev.*, 116:593, 2002.

[22] Wander Jager and Frédéric Amblard. Uniformity, bipolarization and pluriformity captured as generic stylized behavior with an agent-based simulation model of attitude change. *Computational & Mathematical Organization Theory*, 10:295–303, 2005.

[23] Wei Ju, Zheng Fang, Yiyang Gu, Zequn Liu, Qingqing Long, Ziyue Qiao, Yifang Qin, Jianhao Shen, Fang Sun, Zhiping Xiao, et al. A comprehensive survey on deep graph representation learning. *Neural Networks*, page 106207, 2024.

[24] A.-R. Lagos and G.P. Papavassilopoulos. Network topology design to influence the effects of manipulative behaviors in a social choice procedure. *Journal of the Franklin Institute*, 359(7): 3046–3070, 2022.

[25] Jacopo Lenti, Fabrizio Silvestri, and Gianmarco De Francisci Morales. Variational inference of parameters in opinion dynamics models, 2024. URL https://arxiv.org/abs/2403.05358.

[26] Omer Lev and Yoad Lewenberg. "reverse gerrymandering": Manipulation in multi-group decision making. In *AAAI*, pages 2069–2076, 2019.

[27] Harry A Levin and Sorelle A Friedler. Automated congressional redistricting. *Journal of Experimental Algorithmics (JEA)*, 24:1–24, 2019.

[28] Chen Ling, Junji Jiang, Junxiang Wang, My T Thai, Renhao Xue, James Song, Meikang Qiu, and Liang Zhao. Deep graph representation learning and optimization for influence maximization. In *International conference on machine learning*, pages 21350–21361. PMLR, 2023.

[29] Xiaoxue Liu, Shohei Kato, Fenghui Ren, Guoxin Su, Minjie Zhang, and Wen Gu. Information gerrymandering in elections. In *Knowledge Management and Acquisition for Intelligent Systems*, pages 83–97. Springer, 2023.

[30] Nolan McCarty, Keith T Poole, and Howard Rosenthal. Does gerrymandering cause polarization? *American Journal of Political Science*, 53(3):666–680, 2009.

[31] Kanav Mehra, Nanda Kishore Sreenivas, and Kate Larson. Deliberation and voting in approval-based multi-winner elections. In *AAMAS*, pages 77–93, 2023.

[32] Samuel Merrill and Bernard Grofman. *A unified theory of voting: Directional and proximity spatial models*. Cambridge University Press, 1999.

[33] Carl Edward Rasmussen and Christopher K. I. Williams. *The structure and dynamics of networks*. MIT Press, 2006.

[34] Norman Schofield. *The spatial model of politics.* Routledge, 2007.

[35] Sigal Sina, Noam Hazon, Avinatan Hassidim, and Sarit Kraus. Adapting the social network to affect elections. In *AAMAS*, pages 705–713, 2015.

[36] Alexander J Stewart, Mohsen Mosleh, Marina Diakonova, Antonio A Arechar, David G Rand, and Joshua B Plotkin. Information gerrymandering and undemocratic decisions. *Nature*, 573(7772): 117–121, 2019.

[37] Rahul Swamy, Douglas M King, Ian G Ludden, Kiera W Dobbs, and Sheldon H Jacobson. A practical optimization framework for political redistricting: A case study in arizona. *Socio-Economic Planning Sciences*, 92:101836, 2024.

[38] Mohamed N Zareer and Rastko R Selmic. Maximizing opinion polarization using double deep q-learning on social networks. *IEEE Access*, 2025.

[39] Xiaotian Zhou and Zhongzhi Zhang. Maximizing influence of leaders in social networks. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, page 2400–2408, 2021.