Frankenmandering: The Confluence and Symbiosis of Opinion Dynamics and Gerrymandering

Blue Sky Ideas Track

Sahil Agarwal
Carleton University
Ottawa, Canada
SahilAgarwal@cmail.carleton.ca

Svetlana Obraztsova
Carleton University
Ottawa, Canada
svetlana.obraztsova@carleton.ca

Zinovi Rabinovich
Carleton University
Ottawa, Canada
zinovi.rabinovich@carleton.ca

Alan Tsang
Carleton University
Ottawa, Canada
alan.tsang@carleton.ca

ABSTRACT

We present a novel model at the intersection of two prominent sub-fields: Opinion Dynamics and Gerrymandering. While both sub-fields are well studied separately, in the real world, gerrymandering and opinion dynamics combine and compound their effects. Yet, their interaction has been unexplored in Computer Science, including the AAMAS community. In this paper, we present our Frankenmandering framework, which initiates the study of this key gerrymandering-opinion-dynamics intersection. We show how Frankenmandering can occur, its potential to capture significant – both positive and negative – effects on society, and discuss the vast volume of new research questions that arise from this framework.

KEYWORDS

Gerrymandering; Opinion Dynamics; Election Control; Market Creation

ACM Reference Format:
Sahil Agarwal, Svetlana Obraztsova, Zinovi Rabinovich, and Alan Tsang. 2026. Frankenmandering: The Confluence and Symbiosis of Opinion Dynamics and Gerrymandering: Blue Sky Ideas Track. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 6 pages. https://doi.org/10.65109/QBUW2742

1 INTRODUCTION

Singapore, a tropical gem of sun and social engineering, manages its public housing via the Housing and Development Board (HDB). HDB allocates new housing units according to a lottery. When your number is drawn you have a choice of several locations, so you can match your particular life-style preferences. But there is a catch: diversity constraints. Ethnic group quotas are imposed on every housing block, calibrated to reflect the composition of the society at large. Surprisingly, it does not dramatically impact social perceived welfare [8]. In fact, this forced **partitioning** of the population, accompanied with an **electoral process** of representatives into Residents' **Networks**, contributes to the overall country-wide social cohesion. But why? Isn't forced partitioning (a.k.a., redistricting) in voting, especially with geographic constraints (a.k.a., gerrymandering), detrimental and leads to polarisation [3, 30, 37, 53]?

In this paper, we point to the interaction of **redistricting** and the underlying **social network** of citizens as the unsung heroes of Singaporean success. Partitioning is not judged on the basis of a single election, but rather on its long-term effects on opinions and perceived wellbeing. That is, the goal of the partitioning is to reshape society itself. In fact, Frankenmandering, as we call this confluence of two great subfields of research (gerrymandering and opinion dynamics in networks), deserves a place of its own in multi-agent research.

1.1 Gerrymandering

Gerrymandering is a purposeful construction of voting districts to tilt an election's outcome towards a particular party. Classically, districts must obey geographical constraints (e.g., connectivity); though the problem remains interesting even without these restrictions (e.g., in [11, 18, 51]) and with rather unpleasant consequences of preference distortion [31] and misrepresentation [7]. These geographic constraints are commonly expressed as a graph, where nodes represent voters (or voter groups) and edges connect nodes that can be joined in the same district. Though generally such gerrymandering is difficult [22, 35, 48], recent works show that for planar graphs, which describe most polities quite well, the bag of results is mixed (polynomial time solvable for some graphs [27]). Beyond defining "neighbours", geographical constraints can be applied to the very process of drawing districts. For example, [13] discusses the effects of imposing a compactness constraint on districts, while [32] studies gerrymandering of polling station location, though the problem remains NP. Nonetheless, this does not prevent the creation of reasonably well performing, practical solutions [47, 61].

Now, given that the original goal of gerrymandering was to skew the election results, one may ask whether it would be ethical to develop practical means of gerrymandering. The answer lies, of course, on the flip side of this technology. Developing flexible, practical partitioning methods is intimately related to the fair division of resources. E.g., creating equitable school districts [55].

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/QBUW2742

In fact, quite a few "fair" redistricting approaches have been proposed (e.g. [12, 21, 42, 43, 60]). With proper research, it may even be possible to turn gerrymandering from a shadowy political battle into an open one, with socially beneficial outcomes [10]. These positive aspects of partitioning spur research on generating feasible redistricting plans [14, 16, 17, 39, 52, 56]. So that there is no lack of means for identifying negative (gerrymandering) instances and producing positive partitioning plans.

1.2 Opinion Dynamics in Networks

The earliest work in opinion dynamics is perhaps the model of John French Jr. [33] where he introduced the idea that opinions can be represented as numerical values that change over time based on social network interactions. Morris DeGroot [24] subsequently formalized the notion that an agent's opinions can be updated as a function (specifically, a convex combination) of the opinions of her social network neighbours.

In the near century that followed, opinion dynamics has blossomed as a multi-discipline field of research (see [64] for a survey). Of particular interest to us is a branch of models where influence between agents is modulated by their similarity. Krause and Hegselmann [36, 40] introduced the bounded confidence model where agents ignored any opinions differed from its own by more than a value determined by the agent's confidence value. Since then, this model has been variously modified to replace the sharp cut-off with an exponential drop-off [23], and even within the multiagent space by modelling trust values between agents [62].

As it intersects with elections, opinion dynamics has been studied from two main perspectives: forcing changes in network topology [5, 41, 58], or placing "influencers" in the network [4, 50]. The former can be deployed by social network management, while the latter is an open market. Fortunately, neither is too easy [19], though possible [6, 20, 59, 65].

1.3 Where Frankenmandering is borne

Returning to our HDB example, we recognize features of gerrymandering *and* opinion dynamics. In fact, we would be able to give some *common* reasoning as to why it works to produce social cohesion. Like in gerrymandering, HDB rules generate partitions that are microcosms of the island nation. Thus, local in-neighbourhood interests are aligned with the global, island-wide politics. There is also the element of opinion dynamics there: neighbours within the same block are far more likely to interact and influence each other, than two cultural enclaves each living in a separate block. This pressure, maintained year after year, slowly moulds a common citizenship culture. At least that would be a common sense reasoning. But is there any actual science to it?

First, enforced deliberation within a voting group is beneficial for social welfare and representation [54]. Moreover, that deliberation can be critical for enacting change to a status quo [29]. So that if neighbours in a block can resolve their differences at the local level, it is likely to replicate to other blocks as well, and then to the society at large. In other words, partitioning, while it was *not optimised* for this purpose, had a side-effect of creating a localised social network along which opinions flowed and converged.

Second, because HDB was first created to resettle people from kampongs (village compounds), it was enacting *inverse* gerrymandering [46, 63], essentially relocating existing population opinions into already established "districts". This is, again, similar (if not equivalent) to applying opinion dynamics over a graph, where nodes represent population centres (HDB blocks) and edges represent the original links between people that they have inherited from their kampongs.

So, in principle, some science, in support of why and how HDB succeeds, does exist. In other words, Frankenmandering, the intersection between gerrymandering and opinion dynamics, is real – both in the sense of the "real world" and the "real science". However, it was never studied as such. No explicit model attempts to formalise the interaction, no findings to speak of, no theory to ground it. Let us, then, define the first formal Frankenmandering model.

2 BARE BONES OF FRANKENMANDERING

First, let's give some formal notation to all the elements gerrymandering and opinion dynamics have. Then, we will put them into a Frankenmandering interaction model; and, finally, give some positive examples to show that it works.

2.1 Model of Interaction

Given a set of n voters V, each characterised by an opinion $c_v \in \mathbb{R}^m$. We use bold-face $\mathbf{c} \in (\mathbb{R}^m)^n$ to denote the vector of all opinions of all voters, and use the functional form $c_v(l)$ to refer to the l'th coordinate of the opinion 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 $\mathbf{c}|_D = \{c_v\}_{v \in D} \in (\mathbb{R}^m)^{|D|}$.

Voters are organised into two graphs: undirected graph $G^{geo} = (V, E^{geo})$ of geographic constraints, and directed graph $G^{soc} = (V, E^{soc})$ of social influence links. Thus $(u, v) \in E^{geo}$ means two voters $u, v \in V$ can potentially belong to the same district, while $(u, v) \in E^{soc}$ means that voter v is influenced by voter $u \in V$. In addition, let $\mathbb{L}[G] : (\mathbb{R}^m)^n \longrightarrow (\mathbb{R}^m)^n$ be a graph-parameterised opinion dynamics function that dictates how opinions change following the influences encoded by the graph G.

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 each time period t:

(1) Districts, $\mathbf{D}^t = \{D^t_j \subseteq V\}_{j=1}^K$, are selected, so that $\bigcup_{j=1}^{K} D^t_i = V$, and $\forall i, j\ D^t_i \cap D^t_j = \emptyset$, and the subgraph $(D, E^{geo}|_D) \subseteq G^{geo}$ satisfies some constraints (e.g., connectivity);
(2) Local elections are run, producing district representatives $r^t_j = \mathbb{F}\left(\mathbf{c}^t|_{D^t_j}\right) \in D^t_j$, $j \in \{1, ..., K\}$;
(3) Representatives become local "influencers"; i.e., graphs, $\{H^t_j = (D^t_j, E^t_j)\}_{j=1}^K$, are built so that $(r^t_j, v) \in E^t_j$ for all $v \in D^t_j$;
(4) Social influence is then exercised on all opinions by $\mathbb{L}^t = \mathbb{L}[G \cup \bigcup_{j=1}^{K} H^t_j]$, so that : $\mathbf{c}^t = \mathbb{L}^t(\mathbf{c}^{t-1})$

Now, define a goal function $J : \mathbf{c}^t \mapsto \mathbb{R}_+$, that evaluates the opinions (e.g., proximity to an "ideal" opinion). Then, a bare bones

Frankenmandering can be defined as the following optimisation:

$$\min_{\mathbf{D}^1,\dots,\mathbf{D}^T} \quad \sum_{t=0}^{t=T} J(\mathbf{c}^t)$$

$$s.t.$$

$$\forall j \in [1:K]\ \forall t \in [1:T]$$

$$r_j^t = \mathbb{F}\left(\mathbf{c}^t|_{D_j^t}\right) \qquad H_j^t = (D_j^t, E_j^t, w_j^t)$$

$$\mathbb{L}^t = \mathbb{L}[G \cup \bigcup_{j=1}^{K} H_j^t] \qquad \mathbf{c}^t = \mathbb{L}^t(\mathbf{c}^{t-1})$$

2.2 "Positive" Examples

Now, let us see what an instance of the above model can look like, and what strategic choice over districts $\mathbf{D}^t$ can do to voter opinions.

In both examples, we will use an opinion dynamic based on *discrepancy reponse functions* (DRFs). These functions determine how a voter changes its opinion relatively to the opinon of its neighbour(s). E.g., if the opinions are close enough, the voter may consider altering its opinions to be closer to those of its neighbour. However, if the opinions are too far, then the voter may seek to move away from the opinion of its neighbour (a "backfire" effect, also used in [38, 44]). We will assume that opinions are single dimensional and, thus, use the representative selection function that returns the median (w.r.t. opinions in the district) voter.

Frankenmandering Inchworm. First, let's show 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 uses the planar adjacency of voters, G^{geo}, as depicted in Figure 1. This corresponds to a realistic geography, where districts are contiguous. At the same time, we assume an empty social network graph G^{soc} ("a city of strangers").

We begin with the following opinion profile 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 (opinion dynamics function, $\mathcal{L}$), depicted for a single voter-influencer pair (u, v) with opinions c_v, c_u:

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

Frankenmandering then seeks to select a sequence of districts d^t such that we shift the entire opinion profile by exactly +1; i.e. we produce the target opinion profile $\{1, 1, 1, 2, 3, 4, 5, 6, 6, 6\}$. The intuition behind the re-districting solution proceeds in 2 phases: The first phase uses the backfire effect to "push" the most positive voters away from the main body while using the median voter still to anchor the main body. 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 illustrates a sequence of arrow diagrams that depicts the selection of districts that causes all opinions to shift exactly +1. In

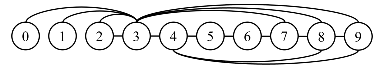

Figure 1: "Geography" Graph of Voters in Fig 2

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.

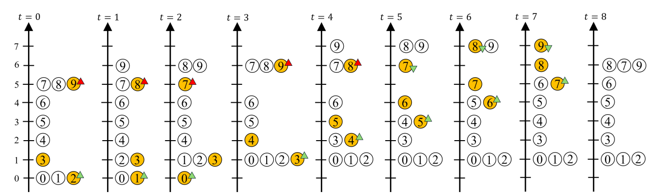

Figure 2: Gerrymandering Sequence to Shift All Voters

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 2 voters and the last voter. The social network will be a line graph, with each voter connected to (up to) two nearest peers. Finally, the opinion dynamics are guided by the following update rule:

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

Note that this opinion dynamics has "indifference" and "irrelevance" plateaus, where the opinions are too similar (indifference) or too far (irrelevance) to cause any change.

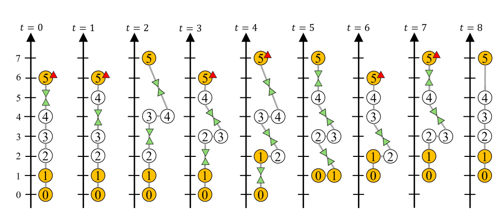

Figure 3: A Fixed Gerrymandered District to Shift All Voters

The sequence of arrow diagrams in Figure 3 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^*. As before, we use small arrows to depict the influence from the representative and social influence between voters; these happen to be exactly the red and green arrows, respectively. 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.

3 DISCUSSION AND FUTURE WORKS

Thus far, we have presented the reader with a simple, "bare bones" model of Frankenmandering and demonstrated that it can be far more insidious than gerrymandering or opinion dynamics manipulation alone. The combined power of the two elements is great. But do our examples explain the motivational example of HDB's success? Is Frankenmandering good for anything else? The answer to these questions is the very reason why Frankenmandering should become a popular research topic in the coming years.

3.1 What is it good for?

Let us start from the bad. In our examples, Frankenmandering was used to drive opinions in a specific direction. From a gerrymandering point of view this is ingenious – not only will a particular party win the parliament once, but eventually, will do so in perpetuity. But the redistricting itself was not optimised for winning an election, so state-of-the-art gerrymandering detection utilities are unlikely to detect any wrongdoing. Furthermore, there is no guarantee that current districting laws will not inadvertently create just such a tragic shift in society. This makes Frankenmandering research critical for modern democracies employing district based elections, while their citizens are addicted to social media.

On the positive side, teachers often "aim" for the median students, which can induce a drift of test scores toward the median as well. So what if we create school districts based on the Frankenmandering (rather than gerrymandering) principle? As our examples show, it promises to continually raise everyone's test scores over time.

Of course, we will not leave this discussion of good and bad without an "ugly". Targeted advertising artificially partitions networked groups of people or, rather, exploits it. For example, signing up to a group about race cars inevitably exposes one to the group's "median culture", which is often not about cars. This is Frankenmandering in the wild. Identifying it, would give advertising companies a different kind of "targeting" capability. Injecting adverts in a Frankenmander'ed fashion, does not just "find" a relevant audience, but can "create" a relevant audience as well.

3.2 What's missing outright?

There is a tremendous range of possible extensions to the "bare bones" Frankenmandering model we have presented. For example, even though it is extremely generic, our current model does not allow representatives to form a secondary social network. Think of any parliament and it becomes rather obvious that connections between politicians develop differently from the rest of the society.

We also have not modelled transient sub-populations. E.g., education programs for K-12 commonly group students by age, but this grouping unifies different people every year. If we are to apply Frankenmandering to fair school districting, we must extend the model to take this into account.

There are also additional influence network elements that are missing, like the significance/weight of the social link between voters. Non-opinion voter features (like geo-location, or party affiliation) would also be an extension of the bare bones Frankenmandering. For one, there are strong links between the local environment (geo-locale) and the needs and desires (hence, opinions) of voters. Thus, spatial opinion distribution modelling, combined with social graphs, is a stronger ground to strategic districting.

Finally, bare bones Frankenmandering does not include a combined redistricting and direct voter influence, like in, e.g., [25, 26].

3.3 What do our examples not show?

Our examples are but a limited first look at Frankenmandering. An amuse-bouche. Further work is need to produce a reliable Frankenmandering toolkit.

For one, we based our examples on the median-voter representative. But, as a framework, Frankenmandering does not require that particular choice. Instead, for example, we can ask community volunteers to step forward and people would vote for them. Or, if voters have party affiliation features, primaries are run and then candidates step forward and are being elected via your favourite voting rule. Even more interesting, what if the voting process itself follows a social structure, like in liquid democracy? Unlike the controls that directly impact voters (e.g., [1, 2, 9]), Frankenmandering would indirectly control the delegation process.

Second, our examples presume the Frankenmandering goal to be a shift in opinions. But can also be about opinion distribution. E.g., we may desire to ensure that the society does not polarise, so Frankenmandering is applied to ensure opinion diversity is preserved. More generally, our examples did not give insight into how the convergence rate of opinions under the social dynamic limits Frankenmandering scope, and vice versa. Furthermore, we know that opinion dynamics can natively produce structural changes in population opinions (e.g., [15, 49, 57]). Whether Frankenmandering can enhance or subdue these native abilities is an open question.

Our examples also presume that opinion dynamics update rate and redistricting rate are constant and equal. Which is unlikely in the real world, and thus forms another research direction to investigate.

Most obviously, of course, these were just examples. Engineered to awe. But do general Frankenmandering solutions exist? What is their complexity? Can we at least do something that would work "in practice"? Our research group toys with Markov chains, graph neural networks, and reinforcement learning to create Frankenmandering solvers. Reviewers of this paper have also suggested comparing Frankenmandering to "Majority Illusion" [28, 34, 45] for potential solution method transfer or, at least, a range of additional analytical questions to be investigated.

Frankenmandering has generated new ideas even during the review stage, and we call on the community to join us in its vigorous pursuit.

ACKNOWLEDGMENTS

This research was in part supported by the CU Development Grant – NSERC, "Gerrymandering as opinion dynamics control vector".

REFERENCES

[1] Shiri Alouf-Heffetz, Tanmay Inamdar, Pallavi Jain, Nimrod Talmon, and Yash More Hiren. 2024. Controlling Delegations in Liquid Democracy. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems, AAMAS*. 2624–2632.

[2] Shiri Alouf-Heffetz, Łukasz Janeczko, Grzegorz Lisowski, and Georgios Papasotiropoulos. 2025. The Cost Perspective of Liquid Democracy: Feasibility and Control. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 39. 13538–13546.

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

[4] Vincenzo Auletta, Francesco Carbone, and Diodato Ferraioli. 2023. Election Manipulation in Social Networks with Single-Peaked Agents. In *International Conference of the Italian Association for Artificial Intelligence*. Springer, 467–480.

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

[6] Vincenzo Auletta, Diodato Ferraioli, and Carmine Viscito. 2023. Election Manipulation on Social Networks with Abstention. In *Multi-Agent Systems*. 435–444.

[7] Yoram Bachrach, Omer Lev, Yoad Lewenberg, and Yair Zick. 2016. Misrepresentation in District Voting.. In *IJCAI*. 81–87.

[8] Nawal Benabbou, Mithun Chakraborty, Xuan-Vinh Ho, Jakub Sliwinski, and Yair Zick. 2018. Diversity constraints in public housing allocation. In *17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2018)*.

[9] Matthias Bentert, Niclas Boehmer, Maciej Rymar, and Henri Tannenberg. 2022. Who won? Winner determination and robustness in liquid democracy. *arXiv preprint arXiv:2205.05482* (2022).

[10] Felix J Bierbrauer and Mattias Polborn. 2022. *Competitive fair redistricting*. Technical Report. Working Paper. Vanderbilt University and University of Cologne.

[11] Anna Blokhina and Irina Samoilova. 2024. Electoral Control by Partition Voters Using the Schulze Method. In *2024 17th International Conference on Management of Large-Scale System Development (MLSD)*. 1–5.

[12] Niclas Boehmer, Tomohiro Koana, and Rolf Niedermeier. 2023. A refined complexity analysis of fair districting over graphs. *Autonomous Agents and Multi-Agent Systems* 37, 1 (2023), 13.

[13] Allan Borodin, Omer Lev, Nisarg Shah, and Tyrone Strangway. 2022. Little House (Seat) on the Prairie: Compactness, Gerrymandering, and Population Distribution.. In *AAMAS*, Vol. 2022. 154–162.

[14] Laurent Bouton, Garance Genicot, Micael Castanheira, and Allison L. Stashko. 2023. *Pack-Crack-Pack: Gerrymandering with Differential Turnout*. NBER Working Papers 31442. National Bureau of Economic Research, Inc. https://doi.org/None

[15] Markus Brill, Edith Elkind, Ulle Endriss, and Umberto Grandi. 2016. Pairwise diffusion of preference rankings in social networks. In *International Joint Conference on Artificial Intelligence (IJCAI 2016)*. 130–136.

[16] Sarah Cannon, Daryl DeFord, and Moon Duchin. 2024. Repetition effects in a Sequential Monte Carlo sampler. arXiv:2409.19017 [math.PR] https://arxiv.org/abs/2409.19017

[17] Sarah Cannon, Moon Duchin, Dana Randall, and Parker Rule. 2022. Spanning tree methods for sampling graph partitions. *arXiv preprint arXiv:2210.01401* (2022).

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

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

[20] Haipeng Chen, Bryan Wilder, Wei Qiu, Bo An, Eric Rice, and Milind Tambe. 2023. Complex Contagion Influence Maximization: A Reinforcement Learning Approach. In *IJCAI*. 5531–5540.

[21] Gabriel Chuang, Oussama Hanguir, and Clifford Stein. 2024. Drawing Competitive Districts in Redistricting. In *5th Symposium on Foundations of Responsible Computing, FORC 2024, Harvard University, Cambridge, MA, USA, June 12-14, 2024 (LIPIcs, Vol. 295)*, Guy N. Rothblum (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 7:1–7:22. https://doi.org/10.4230/LIPICS.FORC.2024.7

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

[23] Guillaume Deffuant, Frédéric Amblard, and Gérard Weisbuch. 2004. Modelling group opinion shift to extreme: the smooth bounded confidence model. *arXiv preprint cond-mat/0410199* (2004).

[24] Morris H DeGroot. 1974. Reaching a consensus. *Journal of the American Statistical association* 69, 345 (1974), 118–121.

[25] Sanyukta Deshpande, Ian G Ludden, and Sheldon H Jacobson. 2025. Bias in the ballot: how votemandering exploits gerrymandering and campaign strategies. *Annals of Operations Research* (2025), 1–48.

[26] Palash Dey. 2023. Priced Gerrymandering. *Theor. Comput. Sci.* 972 (2023), 114080.

[27] Jack Dippel, Max Dupré la Tour, April Niu, Sanjukta Roy, and Adrian Vetta. 2024. Gerrymandering Planar Graphs. In *AAMAS*. 463–471.

[28] Jack Dippel, Max Dupré la Tour, April Niu, Sanjukta Roy, and Adrian Vetta. 2025. Eliminating majority illusion is easy. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 39. 13763–13770.

[29] Edith Elkind, Davide Grossi, Ehud Shapiro, and Nimrod Talmon. 2024. United for change: deliberative coalition formation to change the status quo. *Social Choice and Welfare* 63, 3 (2024), 717–746.

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

[31] Aris Filos-Ratsikas and Alexandros A Voudouris. 2024. Revisiting the distortion of distributed voting. *Theory of Computing Systems* 68, 5 (2024), 1138–1159.

[32] Zack Fitzsimmons and Omer Lev. 2020. Selecting Voting Locations for Fun and Profit. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, Christian Bessiere (Ed.). ijcai.org, 224–230. https://doi.org/10.24963/IJCAI.2020/32

[33] John RP French Jr. 1956. A formal theory of social power. *Psychological review* 63, 3 (1956), 181.

[34] Umberto Grandi, Lawqueen Kanesh, Grzegorz Lisowski, MS Ramanujan, and Paolo Turrini. 2025. A Complexity-Theoretic Analysis of Majority Illusion in Social Networks. *Journal of Artificial Intelligence Research* 83 (2025).

[35] Sushmita Gupta, Pallavi Jain, Fahad Panolan, Sanjukta Roy, and Saket Saurabh. 2021. Gerrymandering on graphs: computational complexity and parameterized algorithms. In *International Symposium on Algorithmic Game Theory*. Springer, 140–155.

[36] Rainer Hegselmann and Ulrich Krause. 2002. Opinion Dynamics and Bounded Confidence Models, Analysis, and Simulation. *Journal of Artifical Societies and Social Simulation (JASSS)* 5, 3 (2002).

[37] Samuel Issacharoff. 2002. Gerrymandering and political cartels. *Harv. L. Rev.* 116 (2002), 593.

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

[39] Myung Jin Kim. 2018. Multiobjective spanning tree based optimization model to political redistricting. *Spatial Information Research* 26, 3 (2018), 317–325.

[40] Ulrich Krause. 1997. Soziale dynamiken mit vielen interakteuren. eine problemskizze. *Modellierung und Simulation von Dynamiken mit vielen interagierenden Akteuren* 3751, 2 (1997).

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

[42] Zeph Landau, Oneil Reid, and Ilona Yershov. 2009. A fair division solution to the problem of redistricting. *Social Choice and Welfare* 32, 3 (2009), 479–492.

[43] Zeph Landau and Francis Edward Su. 2014. Fair division and redistricting. *The Mathematics of Decisions, Elections, and Games* 625 (2014), 17.

[44] Jacopo Lenti, Fabrizio Silvestri, and Gianmarco De Francisci Morales. 2024. Variational Inference of Parameters in Opinion Dynamics Models. arXiv:2403.05358 [cs.CY] https://arxiv.org/abs/2403.05358

[45] Kristina Lerman, Xiaoran Yan, and Xin-Zeng Wu. 2016. The" majority illusion" in social networks. *PloS one* 11, 2 (2016), e0147617.

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

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

[48] Yoad Lewenberg, Omer Lev, and Jeffrey S Rosenschein. 2017. Divide and conquer: Using geographic manipulation to win district-based elections. In *Proceedings of the 16th conference on autonomous agents and multiagent systems*. 624–632.

[49] Christian List, Robert C. Luskin, James S. Fishkin, and Iain McLean. 2013. Deliberation, Single-Peakedness, and the Possibility of Meaningful Democracy: Evidence from Deliberative Polls. *The Journal of Politics* 75, 1 (2013), 30–95.

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

[51] Cynthia Maushagen and Jörg Rothe. 2017. Complexity of control by partition of voters and of voter groups in veto and other scoring protocols. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*. 615–623.

[52] Cory McCartan and Kosuke Imai. 2023. Sequential Monte Carlo for sampling balanced and compact redistricting plans. *The Annals of Applied Statistics* 17, 4 (2023), 3300–3323.

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

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

[55] Ariel Procaccia, Isaac Robinson, and Jamie Tucker-Foltz. 2024. School redistricting: Wiping unfairness off the map. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2704–2724.

[56] Ariel D Procaccia and Jamie Tucker-Foltz. 2022. Compact redistricting plans have many spanning trees. In *Proceedings of the 2022 annual ACM-SIAM symposium on discrete algorithms (SODA)*. SIAM, 3754–3771.

[57] Soroush Rafiee Rad and Olivier Roy. 2021. Deliberation, single-peakedness, and coherent aggregation. *American Political Science Review* 115, 2 (2021), 629–648.

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

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

[60] Ana-Andreea Stoica, Abhijnan Chakraborty, Palash Dey, and Krishna P. Gummadi. 2020. Minimizing Margin of Victory for Fair Political and Educational Districting. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '20, Auckland, New Zealand, May 9-13, 2020*, Amal El Fallah Seghrouchni, Gita Sukthankar, Bo An, and Neil Yorke-Smith (Eds.). International Foundation for Autonomous Agents and Multiagent Systems, 1305–1313. https://doi.org/10.5555/3398761.3398912

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

[62] Alan Tsang and Kate Larson. 2014. Opinion dynamics of skeptical agents. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*. 277–284.

[63] Rene Van Bevern, Robert Bredereck, Jiehua Chen, Vincent Froese, Rolf Niedermeier, and Gerhard J Woeginger. 2015. Network-based vertex dissolution. *SIAM Journal on Discrete Mathematics* 29, 2 (2015), 888–914.

[64] Haoxiang Xia, Huili Wang, and Zhaoguo Xuan. 2011. Opinion dynamics: A multidisciplinary review and perspective on future research. *International Journal of Knowledge and Systems Science (IJKSS)* 2, 4 (2011), 72–91.

[65] Xiaotian Zhou and Zhongzhi Zhang. 2021. Maximizing Influence of Leaders in Social Networks. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 2400–2408.