All Politics is Local: Redistricting via Local Fairness*

Shao-Heng Ko† Erin Taylor† Pankaj K. Agarwal Kamesh Munagala

Department of Computer Science, Duke University

Email: {sk684,ect15,pankaj,kamesh}@cs.duke.edu

September 30, 2022

Abstract

In this paper, we propose to use the concept of *local fairness* for auditing and ranking redistricting plans. Given a redistricting plan, a *deviating group* is a population-balanced contiguous region in which a majority of individuals are of the same interest and in the minority of their respective districts in the given redistricting plan; such a set of individuals have a justified complaint with how the redistricting plan was drawn. A redistricting plan with no deviating groups is called *locally fair*. We show that the problem of auditing a given plan for local fairness is NP-complete. We present an MCMC approach for auditing as well as ranking redistricting plans. We also present a dynamic programming based algorithm for the auditing problem that we use to demonstrate the efficacy of our MCMC approach. Using these tools, we test local fairness on real-world election data, showing that it is indeed possible to find plans that are almost or exactly locally fair. Further, we show that such plans can be generated while sacrificing very little in terms of compactness and existing fairness measures such as competitiveness of the districts or seat shares.

1 Overview

Redistricting in the United States is the process of partitioning a state into districts, each of which elects one representative to the Congress, for the most part, via simple majority voting. As of April 2022, or exactly one year after the US Census Bureau released the results of the 2020 decennial census, 41 out of the 50 states have finished redrawing the congressional redistricting plans for the next decade [28]. This process has triggered numerous debates and litigation along the way. Much of this debate centers on whether the plans are *gerrymandered* so that one of the two parties gets more representatives. Given its high-stakes impact and mathematical richness, there has been persistent interest in tackling redistricting as an algorithmic question since the early 1960s [16].

Within the research community, there is debate around what a “desirable” redistricting plan should mathematically be. Indeed, it is commonly agreed that “desirable” plans should at the minimum produce population balanced, contiguous, and compact districts [1]. Beyond this basic agreement, there is still debate on richer notions of desirability, particularly notions related to the “fairness” of a plan. This has motivated a long line of recent work [15, 16, 23] as well as software tools [5, 28] on *auditing* a given redistricting plan against fairness concepts. Some of these concepts have since been adopted in Wisconsin’s and Michigan’s redistricting efforts [12]. It should be noted that under most notions of desirability proposed in literature, the problem of redistricting is computationally hard [26], leading to the study of heuristic approaches that we outline later.

Global versus Local Fairness. Zooming into fairness criteria, most extant notions of fairness focus on the *global outcomes* of the redistricting plans, e.g., whether the *seat shares* proportionally represent the demographics [39], or how competitive the districts drawn are [15]. However, it is argued in [6] that global metrics do not always distinguish between *natural* gerrymandering – when the distribution of voters unavoidably prohibits certain globally fair outcomes – and *artificial* gerrymandering – when the plans are manipulated to favor a demographic group. This issue is typically addressed via *statistical tests* [16]: An MCMC method is used to generate an ensemble of population balanced, contiguous, and compact plans, and the global fairness score in question is computed for each of these plans, yielding a

*This work will appear at NeurIPS 2022.

†Equal contributions.

histogram of scores. The plan in question is deemed "fair" if its global fairness score is not an outlier in this histogram.

Furthermore, global fairness, such as proportional seat shares, are desirable and statistically testable, these seats may not represent the *local* concerns that a group of voters may have. For instance, imagine the blue party cares about rising sea levels and climate change, while the red party does not. In North Carolina, if we give blue voters on the eastern coast one seat, that representative may advocate to mitigate the impacts of climate change to the coastal residents on the state or federal level. On the other hand, a better seat share may entail making all districts near the coast red, while making the districts in the western mountains blue. However, the latter set of representatives may not advocate for issues impacting the coastal residents, since it is not of local concern to the mountains. This motivates the need for local fairness as a separate fairness measure, capturing at some level the saying "all politics is local".

Borrowing the notion of core from cooperative game theory, the work of [6] defines local fairness notion as follows: Given a redistricting plan, a voter is unsatisfied if the majority demographic in her district does not match her own demographic. A redistricting plan is locally fair if no group of unsatisfied voters could *deviate* and draw a different district such that this group of unsatisfied voters has a majority in the new district.

As in the scenario above, the advantage of such a local notion of fairness is that it captures *justified complaints* of groups of voters, as has happened in recent court judgements [2, 3]. It also provides a way of assessing enacted plans without resorting to statistical tests, making it more human interpretable and *explainable.*

Research Questions. The notion of local fairness is appealing; however, the analysis and results in [6] are theoretical and apply only to a simplified one-dimensional model. In this paper, we seek to develop algorithms to audit plans for local fairness, and subsequently systematically study this concept on real-world electoral data. In particular, we aim to study the following questions:

- Given a redistricting plan, can we efficiently test (or audit) whether the plan is locally fair?
- Are locally fair plans achievable in real redistricting tasks? If not, can we quantify how far a given plan is from being locally fair?
- Is local fairness empirically compatible with other existing global fairness concepts?

1.1 Related Work

Redistricting as Optimization. We first focus on the task of drawing plans, or computational redistricting. The idea of using computational tools in redistricting dates back to the 1960s [24, 37]. Since then, an extensive line of work (see [8] for a comprehensive survey) cast the redistricting task as an optimization problem, in which the input contains only spatial location of individuals, but not their political affiliations. The objective and constraints capture the population balance, contiguity, and compactness criteria of the districts. This problem is computationally intractable in the worst case [14], and multiple algorithmic approaches have been proposed, including Voronoi diagrams [20, 27], local search [25], simulated-annealing and hill climbing [7], and spatial evolutionary algorithms [30]. On the flip side, it is argued in [11, 40] that such "neutral" districting plans – as outputs of algorithms without political inputs – may contain unintentional biases, as well as unexpected outcomes such as "natural gerrymandering" [9, 19], i.e., the geographic distributions of voters naturally lead to disproportionate seat shares. Therefore, fairness objectives such as partisan representativeness are typically incorporated into the redistricting problem as objectives; however, these additional requirements only add to the computational difficulty of the problem [26].

Ensemble Approaches to Redistricting. Instead of optimizing and finding a single best redistricting plan, another line of work focuses on generating large *ensembles* of districting plans, with the hope that some of these plans will be fair. These methods include Flood Fill [13, 31], Column Generation [21], and the widely adopted Markov Chain Monte Carlo (MCMC) approach [18, 29, 36]. The latter approach *samples* from the space of feasible plans with a bias towards "desirable" or fairness properties. For instance, it is shown in [33] that the widely used ReCom MCMC method [16] provably biases towards compact plans. The work of [17] proposes a method for choosing one representative plan from such an ensemble based on defining distances between plans.

Auditing and Combating Gerrymandering. A somewhat different question from constructing a desirable plan is the question of *auditing* a given plan for desirability and fairness. As mentioned before, ensemble based approaches provide a natural, statistical way of auditing [22, 23]: The properties of the plan that is currently enacted is compared against the histogram of the corresponding property on the ensemble; if the plan is a statistical outlier, then it is considered more "gerrymandered" and hence less desirable. The recent work of [29] instead uses plans in the ensemble as comparators to identify packing and cracking in districting plans. On the non-statistical side, numerous approaches to auditing have also been proposed via appropriate desirability scores. These are either scores based on compactness of the plan (such as the Reock [34] and Polsby-Popper [32] scores), or scores based on partisan outcomes generated by the plan (such as the efficiency gap [35], mean-median gap [38], partisan symmetry [39], and the GEO metric [10]), or scores based on competitiveness of the plan [15]. Many of these measures are used in publicly available tools [4, 5]. Finally, there is a recent line of work that attempts to eliminate gerrymandering by completely revamping the winner-takes-all, single-member district mechanism into a multiwinner election [19].

1.2 Our Contribution

In this paper, we take the conventional view of redistricting as partitioning a planar graph on precincts into population-balanced, contiguous, and (in a heuristic sense) compact regions. We naturally extend the local fairness concept proposed in [6] to this task.

We first focus on the question of *auditing* a given plan for local fairness, that is, the non-existence of a population-balanced contiguous region in which a majority of voters are of the same party and unhappy in the given plan. We show that this problem is computationally intractable in the worst case. Our first contribution is two heuristics for the auditing problem. Our first approach, that is scalable and practical, extends existing ensemble-based methods in a novel way: We assume the districts in the ensemble are the only districts to which voters can deviate, and given a plan to be audited, we test each of these districts as a potential deviation on that plan. Our second approach drills deeper into plans where the ensemble based method finds no deviating group; indeed, if the method found a deviating group, the plan was already deemed not locally fair. On the former set, we generate several random spanning trees, and devise a polynomial time dynamic programming algorithm that audits each tree for local fairness. If any of these audits finds a deviating group, the original plan was not locally fair. The dynamic program is not as efficient as the ensemble-based method; however, we provide empirical evidence that the ensemble method suffices to deem a plan locally fair, and the dynamic program typically does not find additional deviating groups. Finally, for redistricting plans that are not locally fair, we propose a measure that quantifies the unfairness of the plans by the portion of population with a justified complaint.

As our second contribution, we empirically study the notion of local fairness on real data on recent elections in the US. We generate plans using the (by now) standard ReCom [16] ensemble method, and audit each plan for local fairness using the ensemble method, thereby producing an ordering of the plans via our unfairness measure. We empirically show that applying the criterion of local fairness prunes the space of candidate plans considerably, while still returning a set of potential candidates. Most global and statistical notions of fairness fail to do such pruning, since they are endogenously defined relative to the order statistics on the ensemble. We further show that not only is local fairness *achievable* on real redistricting tasks, but it is also compatible with extant global fairness properties. Indeed, when we compare locally fair plans and those with many deviating groups, the former tend to be just as compact, have comparable seat share outcomes, and sacrifice only a small amount of competitiveness. Thus local fairness can be used as an additional fairness criterion in conjunction with a global fairness criterion. We also investigate robustness of the local fairness concept, and show that fair redistricting plans remain consistent across different elections used. We finally show visualizations of fair and unfair plans; in particular showing that the visualization of deviating groups and likelihood of precincts being unhappy makes the fairness notion explainable.

Taken together, our results demonstrate local fairness as an effective *pruning criterion* for candidate redistricting plans while sacrificing little in other desired properties. We also note that in practice, there could be other considerations when choosing the "best" plan even among many locally-fair plans; we leave the question of choosing these considerations to policy makers.

2 Open Questions

Several open questions arise from our work. In terms of algorithm design, it is an open question of whether there is an approximation algorithm for either the auditing or generation problem, and whether such an algorithm could take into account compactness. It would also be interesting to extend our methods to capture additional real-world criteria used in redistricting, such as a penalty for splitting up counties, or a requirement for a majority-minority district. In particular, can fair plans be locally modified so that they remain fair and such real-world criteria are satisfied?

Finally, our exploration of robustness of local fairness to voter turnout is preliminary, since it compares the outcomes between two election data in one state. It would be interesting to extend our work to a stochastic setting, where each individual in the population has a "likely voter" score (or probability to vote), and we need high confidence in the non-existence of a deviating group.

References

[1] Reynolds v. Sims. `https://supreme.justia.com/cases/federal/us/377/533/`, 1964.

[2] Cooper v. Harris. `https://supreme.justia.com/cases/federal/us/581/15-1262/`, 2017.

[3] Rucho v. Common Cause. `https://supreme.justia.com/cases/federal/us/588/18-422/`, 2019.

[4] GEO metric: A metric to detect gerrymandering which uses both Geography and Election Outcomes. `https://www.the-geometric.com/`, 2022.

[5] Princeton gerrymandering project. `https://gerrymander.princeton.edu/`, 2022.

[6] P. K. Agarwal, S. Ko, K. Munagala, and E. Taylor. Locally fair partitioning. In *AAAI*, pages 4752–4759. AAAI, 2022.

[7] M. Altman and M. P. McDonald. BARD: Better automated redistricting. *Journal of Statistical Software*, 42:1–28, 2011.

[8] A. Becker and J. Solomon. Redistricting algorithms. In M. Duchin and O. Walch, editors, *Political Geometry: Rethinking Redistricting in the US with Math, Law, and Everything In Between*, chapter 16, pages 303–340. Springer, 2022.

[9] A. Borodin, O. Lev, N. Shah, and T. Strangway. Big city vs. the great outdoors: voter distribution and how it affects gerrymandering. In *IJCAI*, pages 98–104, 2018.

[10] M. Campisi, T. Ratliff, S. Somersille, and E. Veomett. Geography and election outcome metric: An introduction. *Election Law Journal: Rules, Politics, and Policy*, 2022.

[11] J. Chen, J. Rodden, et al. Unintentional gerrymandering: Political geography and electoral bias in legislatures. *Quarterly Journal of Political Science*, 8(3):239–269, 2013.

[12] M. Chen. Tufts research lab aids states with redistricting process. `https://tuftsdaily.com/news/2021/04/06/tufts-research-lab-aids-states-with-redistricting-process/`, 2021.

[13] C. Cirincione, T. A. Darling, and T. G. O'Rourke. Assessing South Carolina's 1990s congressional districting. *Political Geography*, 19(2):189–211, 2000.

[14] V. Cohen-Addad, P. N. Klein, D. Marx, A. Wheeler, and C. Wolfram. On the computational tractability of a geographic clustering problem arising in redistricting. In *FORC*, pages 3:1–3:18, 2021.

[15] D. DeFord, M. Duchin, and J. Solomon. A computational approach to measuring vote elasticity and competitiveness. *Statistics and Public Policy*, 7(1):69–86, 2020.

[16] D. DeFord, M. Duchin, and J. Solomon. Recombination: A family of markov chains for redistricting. *Harvard Data Science Review*, 3(1), 3 2021. `https://hdsr.mitpress.mit.edu/pub/1ds8ptxu`.

[17] S. A. Esmaeili, H. Grape, and B. Brubach. Centralized fairness for redistricting. *arXiv preprint arXiv:2203.00872*, 2022.

[18] B. Fifield, M. Higgins, K. Imai, and A. Tarr. Automated redistricting simulation using Markov chain Monte Carlo. *Journal of Computational and Graphical Statistics*, 29(4):715–728, 2020.

[19] N. Garg, W. Gurnee, D. Rothschild, and D. B. Shmoys. Combatting gerrymandering with social choice: The design of multi-member districts. In *EC*, pages 560–561. ACM, 2022.

[20] P. Gawrychowski, H. Kaplan, S. Mozes, M. Sharir, and O. Weimann. Voronoi diagrams on planar graphs, and computing the diameter in deterministic $\tilde{O}(n^{5/3})$ time. *SIAM Journal on Computing*, 50(2):509–554, 2021.

[21] W. Gurnee and D. B. Shmoys. Fairmandering: A column generation heuristic for fairness-optimized political districting. In *SIAM ACDA*, pages 88–99, 2021.

[22] G. Herschlag, H. S. Kang, J. Luo, C. V. Graves, S. Bangia, R. Ravier, and J. C. Mattingly. Quantifying gerrymandering in North Carolina. *Statistics and Public Policy*, 7(1):30–38, 2020.

[23] G. Herschlag, R. Ravier, and J. C. Mattingly. Evaluating partisan gerrymandering in Wisconsin. *arXiv preprint arXiv:1709.01596*, 2017.

[24] S. W. Hess, J. Weaver, H. Siegfeldt, J. Whelan, and P. Zitlau. Nonpartisan political redistricting by computer. *Operations Research*, 13(6):998–1006, 1965.

[25] D. M. King, S. H. Jacobson, and E. C. Sewell. Efficient geo-graph contiguity and hole algorithms for geographic zoning and dynamic plane graph partitioning. *Mathematical Programming*, 149(1):425–457, 2015.

[26] R. Kueng, D. G. Mixon, and S. Villar. Fair redistricting is hard. *Theoretical Computer Science*, 791:28–35, 2019.

[27] H. A. Levin and S. A. Friedler. Automated congressional redistricting. *ACM Journal of Experimental Algorithmics*, 24(1):1.10:1–1.10:24, 2019.

[28] J. Levitt. All about redistricting. `https://redistricting.lls.edu/`, 2022.

[29] J. Lin, C. Chen, M. Chmielewski, S. Zaman, and B. Fain. Auditing for gerrymandering by identifying disenfranchised individuals. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 1125–1135, 2022.

[30] Y. Y. Liu, W. K. T. Cho, and S. Wang. PEAR: A massively parallel evolutionary computation approach for political redistricting optimization and analysis. *Swarm and Evolutionary Computation*, 30:78–92, 2016.

[31] D. B. Magleby and D. B. Mosesson. A new approach for developing neutral redistricting plans. *Political Analysis*, 26(2):147–167, 2018.

[32] D. D. Polsby and R. D. Popper. The third criterion: Compactness as a procedural safeguard against partisan gerrymandering. *Yale Law & Policy Review*, 9:301, 1991.

[33] A. D. Procaccia and J. Tucker-Foltz. Compact redistricting plans have many spanning trees. In *ACM-SIAM SODA*, pages 3754–3771, 2022.

[34] E. C. Reock. A note: Measuring compactness as a requirement of legislative apportionment. *Midwest Journal of Political Science*, 5(1):70–74, 1961.

[35] N. O. Stephanopoulos and E. M. McGhee. Partisan gerrymandering and the efficiency gap. *University of Chicago Law Review*, 82:831, 2015.

[36] W. K. Tam Cho and Y. Y. Liu. Toward a talismanic redistricting tool: A computational method for identifying extreme redistricting plans. *Election Law Journal*, 15(4):351–366, 2016.

[37] W. Vickrey. On the prevention of gerrymandering. *Political Science Quarterly*, 76(1):105–110, 1961.

[38] S. S.-H. Wang. Three tests for practical evaluation of partisan gerrymandering. *Stanford Law Review*, 68:1263, 2016.

[39] G. S. Warrington. Quantifying gerrymandering using the vote distribution. *Election Law Journal*, 17(1):39–57, 2018.

[40] A. Wheeler and P. N. Klein. The impact of highly compact algorithmic redistricting on the rural-versus-urban balance. In *SIGSPATIAL/GIS*, pages 397–400, 2020.