Redistricting via Local Fairness

Anonymous Author(s)
Affiliation
Address
email

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 Introduction

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 [29]. 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 [15].

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 [14, 15, 23] as well as software tools [5, 29] 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 [11]. It should be noted that under most notions of desirability proposed in literature, the problem of redistricting is computationally hard [27], 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 [40], or how competitive the districts drawn are [14]. 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 –

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

when the plans are manipulated to favor a demographic group. This issue is typically addressed via *statistical tests* [15]: An MCMC method is used to generate an ensemble of population balanced, contiguous, and compact plans, and the fairness score in question is computed for each of these plans, yielding a histogram of scores. The plan in question is deemed "fair" if its score is not an outlier in this histogram.

The usage of statistical tests to assess plans makes it hard to interpret such global notions from the perspective of an individual voter or block of voters. In contrast, a *local* notion of fairness looks at complaints voters may have with the plan at hand. 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.

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, 4]. It also provides a way of accessing enacted plans without resorting to statistical tests, making it more human interpretable and *explainable*.

Though this notion of local fairness is appealing, 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 in this paper:

- 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, 38]. 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 [13], and multiple algorithmic approaches have been proposed, including Voronoi diagrams [20, 28], local search [26], simulated-annealing and hill climbing [7], and spatial evolutionary algorithms [30]. On the flip side, it is argued in [10, 41] 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 [27].

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 [12, 31], Column Generation [21], and the widely adopted Markov Chain Monte Carlo (MCMC) approach [18, 37]. The latter approach *samples* from the space of feasible plans with a bias towards "desirable" or fairness properties. For instance, it is shown in [34] that the widely used ReCom MCMC method [15] provably biases towards compact plans. We provide a more in-depth description of ReCom in Section 3. Finally, 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. 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 [35] and Polsby-Popper [33] scores), or scores based on partisan outcomes generated by the plan (such as the efficiency gap [36], mean-median gap [39], or partisan symmetry [40]), or scores based on competitiveness of the plan [14]. Many of these measures are used in publicly available tools [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 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. 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 plans while sacrificing little in other desired properties.

2 Model and Preliminaries

In keeping with recent literature [13, 15, 16, 21], the input to the redistricting problem is a planar connected graph $G = (V, E)$ where each vertex $v \in V$ represents an indivisible geographic unit (a precinct or a census block),[1] and an edge is placed between two vertices if they are geographically adjacent. Going forward, we refer to each $v \in V$ as a *precinct* and G as the *precinct graph*.

Redistricting Plans. For each precinct $v \in V$, let $\rho(v) > 0$ denote its population and let $\tau(v) \in [0, \rho(v)]$ denote the number of voters in v. We let $\gamma(v) \in [0, 1]$ and $\beta(v) = 1 - \gamma(v)$ denote the fraction of $\tau(v)$ who vote *red* and *blue*, respectively. Note that it is assumed each individual voter is

[1] Typically, precincts are not split by redistricting plans [15].

exactly one of the two colors. For an arbitrary subset of precincts $W \subseteq V$, set $\rho(W) \coloneqq \sum_{v \in W} \rho(v)$, $\tau(W) \coloneqq \sum_{v \in W} \tau(v)$, $\gamma(W) \coloneqq \frac{\sum_{v \in W}(\gamma(v) \cdot \tau(v))}{\tau(W)}$, and $\beta(W) \coloneqq 1 - \gamma(W)$.

Definition 1 (k-redistricting plan)**.** *A* k-redistricting plan *of* $G = (V, E)$ *is a partition of* V *into* k *pairwise-disjoint subsets (or* districts*)* $D_1, D_2, \ldots, D_k \subseteq V$*. Each district assumes the color of the majority of its voters. For a redistricting plan* Π*, let* $\mathcal{B}_\Pi$ *(resp.* $\mathcal{R}_\Pi$*) denote the set of precincts in blue (resp. red) districts in* Π*.*

In the following, we will fix an error parameter $\varepsilon > 0$, and the desired number of partitions k. Note that the average population per district is $\frac{\rho(V)}{k}$. We say a district $D \subseteq V$ is *ε-feasible* if: (1) D induces a connected subgraph, and (2) the population of D is ε away from average, *i.e.*, $(1-\varepsilon) \cdot \frac{\rho(V)}{k} \leq \rho(D) \leq (1+\varepsilon) \cdot \frac{\rho(V)}{k}$. A redistricting plan Π is *ε-feasible* if each district $D_i \in \Pi$ is ε-feasible.

We note that this definition is consistent with the general practice in the U.S, where the total population of a precinct v is determined not by eligible voters $\tau(v)$, but by the total population $\rho(v)$ based on census information. Since ε and k will be fixed throughout, we drop the prefixes and refer to k-redistritcting plans and ε-feasible districts as just *redistricting plans* and *feasible* districts, repsectively.

Local Fairness. We extend the notion of *local fairness* proposed by [6] to the graph-based redistricting problem. We say that a feasible district $W \subseteq V$ is *red-majority* (*red* in short) if $\gamma(W) \geq \beta(W)$, and *blue-majority* (blue) otherwise. We call this majority color as the *color* of W. Given a redistricting plan Π, any voter whose color agrees with the color of its assigned district in Π is deemed *happy* with respect to Π, and the remaining voters are *unhappy*.

Definition 2 (c-locally fair)**.** *Given a feasible redistricting plan* Π *of* G *and a constant* $c \in [1/2, 1]$*, a feasible district* $W \subseteq V$ *is a* red c-deviating group *with respect to* Π *if* W *is red and at least a* c*-fraction of its voters are unhappy red voters, or formally,* $\sum_{v \in W \cap \mathcal{B}_\Pi} \gamma(v) \cdot \tau(v) > c \cdot \tau(W)$*. A* blue c-deviating group *is defined analogously. We call a feasible redistricting plan* Π *of* G c-locally fair *if there are no red or blue* c*-deviating groups with respect to* Π*.*

When $c = 1/2$, only a simple majority of voters in a deviating group must be unhappy. In this special case, we omit the prefix c by referring to red deviating groups, blue deviating groups, and locally fair redistricting plans. Throughout, we refer to locally fair redistricting plans as fair plans.

We are thus interested in the following two problems.

LF AUDITING problem. Given a feasible redistricting plan Π, decide whether Π is c-locally fair.

LFP GENERATION problem. Find a feasible redistricting plan Π of G such that Π is c-locally fair, or report that none exists.

For redistricting plans Π that are not locally fair, we quantify its degree of unfairness as follows. Consider all deviating groups of Π, and define its unfairness score as the relative size of the union of all voters that were unhappy in some deviating group. Formally, let $W^*_{\mathrm{B}}(\Pi) \coloneqq \bigcup \{W \subseteq V \mid W \text{ is a blue deviating group of } \Pi\}$ denote the set of all precincts that lie in some blue deviating group of Π. Similarly, define $W^*_{\mathrm{R}}(\Pi)$. Then the unfairness score of Π is defined as:

$$\mathrm{unf}(\Pi) \coloneqq \frac{\beta(W^*_{\mathrm{B}}(\Pi)) \cdot \tau(W^*_{\mathrm{B}}(\Pi)) + \gamma(W^*_{\mathrm{R}}(\Pi)) \cdot \tau(W^*_{\mathrm{R}}(\Pi))}{\tau(V)}.$$

Note that $\mathrm{unf}(\Pi) \in [0, 1]$, and is equal to zero if Π is locally fair.[2]

Compactness. In addition to requiring that districts be contiguous and population balanced, many redistricting models also require that the districts be *compact*. However, in contrast to the former two criteria, there is no universally agreed measure of compactness [8, 21]. For example, the Princeton Redistricting Report Cards[3] uses Reock [35] and Polsby-Popper [3, 33] scores, both of which are

[2] We note that another intuitive measure of unfairness considers the population of all precincts that lie in some deviating group, rather than the number of unhappy voters. However, the same set of unhappy voters can form multiple distinct deviating groups by incorporating different sets of voters who were already happy. Thus, the definition of unf score is more meaningful in terms of unhappy voters.

[3] `https://gerrymander.princeton.edu/`

derived from the area and perimeter of the geographic districts drawn. It also uses the number of counties split into multiple districts. In the discrete model of precinct graphs we consider, one common measure of compactness is the number of cut edges formed by the plan, i.e., the number of edges whose endpoints lie in different districts of Π [13, 15]. Though we do not consider compactness in the generation and audit problems, the algorithms we use are biased towards compact plans, as we empirically demonstrate.

Organization. In Section 3, we present two algorithms for the LF AUDITING and LFP GENERATION problems. We then describe the experimental setup and empirical results in Section 4. Additional methodological and experimental results are presented in the Supplementary Material.

3 Efficiently Auditing Local Fairness

In Appendix A, we prove both the LF AUDITING and LFP GENERATION problem are NP-complete, thereby necessitating heuristic or approximately optimal approaches.

Our main contribution in this section is efficient algorithms for LF AUDITING. The methods trade off computational efficiency for accuracy in determining local fairness, and we provide empirical evidence that the more computationally efficient of these methods suffices to deem plans as locally fair. For LFP GENERATION, we simply generate an ensemble of feasible redistricting plans using existing MCMC approaches, and run the LF AUDITING algorithm to find the unfairness score $\text{unf}(\Pi)$ for each generated plan Π, thereby ranking the plans in the ensemble in terms of fairness.

3.1 Ensemble-based Auditing

The computationally more efficient approach makes use of ensembles in a novel way as follows:

1. **Generation.** Generate an ensemble of t redistricting plans $\mathcal{E} = \{\Pi_1, \Pi_2, \ldots, \Pi_t\}$.
2. **Districts to test.** Let $\Delta = \{D \subseteq V \mid \exists j \text{ s.t. } D \in \Pi_j \in \mathcal{E}\}$ be the collection of districts in the plans in the ensemble $\mathcal{E}$.
3. **Audit.** Treat Δ as the candidate set of deviating groups. For each plan $\Pi \in \mathcal{E}$, compute if each $D \in \Delta$ is a c-deviating group (either red or blue). This step yields, for each $\Pi \in \mathcal{E}$, the set $\tilde{W}^*_{\text{B}}(\Pi) \subseteq V$ of precincts in blue deviating groups $D \in \Delta$ (which we use as an approximation of $W^*_{\text{B}}(\Pi)$), and similarly $\tilde{W}^*_{\text{R}}(\Pi)$. We then use $\tilde{W}^*_{\text{B}}(\Pi)$ and $\tilde{W}^*_{\text{R}}(\Pi)$ to compute the score the unfairness score $\text{unf}(\Pi)$ for each $\Pi \in \mathcal{E}$.

For the first step of an ensemble of plans $\mathcal{E}$, we use the ReCom algorithm [15]. This algorithm is a Markov Chain whose state space is the full set of (ε-)feasible (k-)partitions of G. In each step, it randomly combines two (or more in its general form) districts of the current redistricting plan, generates a random spanning tree of the subgraph induced by the combined districts, and re-partitions the subgraph into two parts by cutting edges in the spanning tree so that the new resulting districts remain population balanced. The random spanning tree step *biases* the Markov chain towards plans with compact districts if compactness is measured by number of cut edges [34].

In the remainder of the paper, we describe our auditing framework when ReCom is used as the ensemble generation technique. We note however that we can use any other generation method that produces compact plans, for instance, iterative-merging flood-fill algorithms [10].

3.2 Auditing via Dynamic Programming on Trees

The ensemble-based LF AUDITING method makes one-sided error – if it finds a deviating group, then the plan is not fair, while the absence of a deviating group from Δ only means the plan is *likely* fair. Indeed, in our experiments, there are often multiple redistricting plans that the ensemble-based audit deems likely fair, calling for a more systematic method for further analyzing the candidate plans which are potentially fair. This motivates the approach we now describe.

We will first show that the LF AUDITING problem is efficiently solvable if G is a tree. We will then exploit this fact and use this as a heuristic to solve LF AUDITING on a general planar graph G, given partition Π, as follows:

1. **Random spanning trees.** First generate a collection of random spanning trees $\mathcal{T}$ of G.
2. **Audit each tree.** For each tree $T \in \mathcal{T}$, decide if there exists a feasible deviating group with respect to Π such that this district is a connected subtree of T. Note that such a district will also be a connected subgraph in G.

3.2.1 Dynamic Program for Trees

We now show how to solve the LF AUDITING problem efficiently on trees. We check for the existence of blue deviating groups for $c = 1/2$; the case for red groups and larger c follows similarly. We take as input the precinct graph $G = (V, E)$, a spanning tree T of G, and a redistricting plan Π of G. Let $\sigma = \rho(V)/k$, so that the range for the size of a district is $[(1-\varepsilon)\sigma, (1+\varepsilon)\sigma]$. For each precinct v, let $D(v)$ denote the district that contains v in Π. Let the number of unhappy blue individuals in precinct v be denoted $\text{uhp}(v) = \beta(v) \cdot \tau(v)$ if $D(v) \in \mathcal{R}_\Pi$, and zero otherwise. The task is to decide whether there is a connected subtree $W \subset T$ such that $\sum_{v \in W} \rho(v) \in [(1-\varepsilon)\sigma, (1+\varepsilon)\sigma]$ and $\sum_{v \in W} \text{uhp}(v) > \frac{1}{2} \cdot \sum_{v \in W} \tau(v)$.

For each $v \in V$, let T_v be the subtree of T rooted at v. For each $v \in V$ and $i, p \leq (1+\varepsilon)\sigma$, let $A[v, i, p]$ denote the maximum number of unhappy blue voters in a subtree $W \subseteq T_v$ such that $v \in W$, $\tau(W) \leq i$, and $\rho(W) = p$. To compute $A[v, i, p]$, we take the best subset of children of v such that unhappy blue voters are maximized subject to population and voter constraints.

If v is a leaf precinct in G, then $A[v, i, p] = \text{uhp}(v)$ if $i \geq \tau(v)$ and $p = \rho(v)$; otherwise we have $A[v, i, p] = 0$.

Otherwise, let $U = \text{children}(v) = \{u_1, \ldots, u_{\deg(v)}\}$ denote the set of v's children precincts in T. Let $i' = i - \tau(v)$ and $p' = p - \rho(v)$. We have

$$A[v, i, p] = \text{uhp}(v) + \max_{\{i_j, p_j\}} \left\{ \sum_{j=1}^{\deg(v)} A[u_j, i_j, p_j] \mid \left(\sum_{j=1}^{\deg(v)} i_j \leq i' \right) \wedge \left(\sum_{j=1}^{\deg(v)} p_j = p' \right) \right\}. \quad (1)$$

To compute the second term in Equation (1) efficiently, let $B_{v,i,p}[j, x, y]$ denote the maximum number of total unhappy blue voters in a union of subtrees rooted at $\{u_1, \ldots, u_j\}$ with total voter count at most x and total population level y. The second term of the RHS in Equation (1) is $B_{v,i,p}[\deg(v), i', p']$. We then have $B_{v,i,p}[1, x, y] = A[u_1, x, y]$ for all x, y, and for all $j \geq 2$ we have

$$B_{v,i,p}[j, x, y] = \max_{x' \in [0,x], y' \in [0,y]} \{B_{v,i,p}[j-1, x-x', y-y'] + A[u_j, x', y']\}. \quad (2)$$

The algorithm thus proceeds by computing all $A[v, i, p]$ in an outer loop in increasing order of i and p and in bottom-up order of v. Each computation of $A[v, i, p]$ proceeds with an inner loop that computes $B_{v,i,p}[j, x, y]$ in increasing order of j, x, and y. For each v, after computing all $A[v, i, p]$, we check if there is any $p \in [(1-\varepsilon)\sigma, (1+\varepsilon)\sigma]$ such that $A[v, i, p] > \frac{i}{2}$. If so, there exists a blue deviating group rooted at v with population p, total voter count at most i, and total number of unhappy blue voters $A[v, i, p]$.

Time complexity. The algorithm computes $O(|V| \cdot \sigma^2)$ values of $A[v, i, p]$, each computing $O(\Delta(T) \cdot \sigma^2)$ values of $B_{v,i,p}[j, x, y]$, where $\Delta(T)$ is the maximum degree of a precinct in T. Computing each $B_{v,i,p}[j, x, y]$ requires $O(\sigma^2)$-time to loop through all values of x' and y'. Hence the overall time complexity is $O(|V| \cdot \sigma^6 \cdot \Delta(T))$.

3.2.2 Ensemble-based Auditing is Approximately Sufficient

The dynamic programming method is less desirable because of two significant limitations. The first is its running time, which is prohibitively high for practical use. The second limitation is the introduction of a different type of error from the ensemble-based method. Note that as with the ensemble-based method, the non-existence of a deviating group only provides high confidence (but not absolute certainty) about the plan being locally fair. At the same time, there is a more subtle type of error the DP can make in the other direction: even if it finds a deviating group on the tree, this may not be a *compact* deviating group in the original graph. Therefore, the algorithm can output "not locally fair" when the only deviating groups are non-compact. We therefore need a final step where we check the deviating group to see that it is feasible and compact on the original graph.

In Appendix B.1, we first show approaches to speed up the running time of the dynamic program to make it practical. In Appendix B.2, under a slightly different data model for which the dynamic program runs efficiently, we empirically show that it fails to find compact c'-deviating groups in plans that the ensemble-based audit method deems c-locally fair, for c' slightly larger than c. On the other hand, it easily finds such deviating groups in plans that the ensemble-based method deems unfair. In summary, the dynamic programming approach obviates its own use in practice, since we provide strong evidence that the ensemble-based auditing method suffices in order to deem redistricting plans as locally fair, assuming the strength c of the deviating group is relaxed slightly.

4 Experiments

In our experiments, we will attempt to answer the following questions. First, is local fairness achievable on real redistricting tasks? Second, is it compatible with extant measures of global fairness? Finally, is the notion robust if the underlying data changes? Given the previous discussion, our experiments in this section focus exclusively on the ensemble based method. In this section, we give results for the ensemble-based auditing. In Appendix B we discuss the experiments for auditing via dynamic programming.

4.1 Datasets and Methods

All data used in our experiments is obtained from the MGGG States open repository [32]. We obtain shapefiles and precinct graphs for North Carolina (NC), Pennsylvania (PA), Texas (TX), and Wisconsin (WI). These states are chosen because their elections typically tend to be competitive, leading to greater chance for plans to be locally unfair. We set $\rho(v)$ to the 2010 census total population in each precinct v. The default election we use is the 2016 presidential election, while the 2012 presidential election is also used in the robustness tests. For each precinct v, we collect the number of total votes for the Republican party r_v and the total vote amount for the Democrat party b_v in the 2016 presidential election from [32]. We set $\gamma(v) = r_v/(r_v + b_v)$, $\beta(v) = 1 - r_v$, and $\tau(v) = r_v + b_v$, so $\tau(v)$ is the total number of red (Republican) and blue (Democrat) voters.

Using the ReCom algorithm [15], we generate an ensemble of redistricting plans for each state.[4] Each ensemble consists of $1,000$ redistricting plans, each being the outcome of an independent $10,000$-step Markov chain (with default population balance parameter $\varepsilon = 0.02$) seeded with a recent congressional electoral plan of the state. We set k to be the number of congressional districts in the 2016 election in each state. We then obtain the collection Δ of candidate districts for each state by taking the union of the districts in each plan in the ensemble. The properties of input graphs of states and their corresponding ensembles are summarized in Appendix C.

State	$c = .5$	$c = .51$	$c = .52$	$c = .55$
NC	4.9%	8.7%	13.8%	63.9%
TX	2.8%	5.4%	8.5%	26.0%
PA	11.9%	28.6%	40.7%	95.2%
WI	76.8%	83.6%	87.1%	98.8%

Table 1: Plans without c-deviating groups

4.2 Locally Fair Plans: Counts and Visualizations

We first use the ensemble-based audit method to audit each plan in the ensemble against the set Δ of districts from the entire ensemble. For each $c \in \{0.5, 0.51, 0.52, 0.55\}$, we count the number of plans in the ensemble without c-deviating groups in Δ. We present the results in Table 1. For all values of c, there exists plans in the ensemble that admit no c-deviating groups in Δ and thus are identified fair by the ensemble-based algorithm. Clearly, a larger value of c implies more plans are identified as locally fair. With $c = 1/2$, very few plans are identified as fair in three of the four states.

In Figures 1a and 1c, we present examples of fair plans (with no 0.5-deviating groups) in NC and PA respectively, while in Figures 1b and 1d, we present "unfair" plans with many 0.5-deviating groups

[4] Note that ReCom generates plans without taking into account electoral data.

in these states. We show one of the blue deviating groups with black outline. In these figures, the districts are coded by its color and the extent of partisanship: districts with a larger value of γ (resp. β) are colored in darker red (resp. blue). We show visualizations for other states in Appendix E.

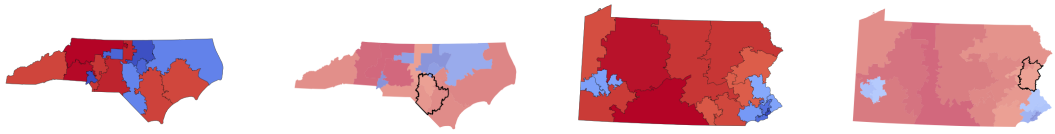

(a) NC Fair Plan (b) NC Plan with blue dev. group (c) PA Fair Plan (d) PA Plan with blue dev. group

Figure 1: North Carolina and Pennsylvania plans without and with deviating groups

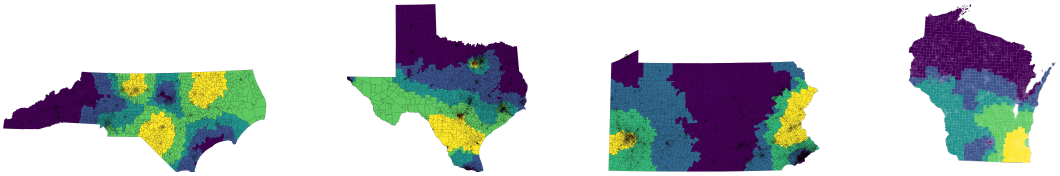

(a) NC (b) TX (c) PA (d) WI

Figure 2: Precincts which are contained in more deviating groups have lighter color.

We note that a large fraction of plans are locally fair in some states like WI. To understand this, for each state, we plot (in Figure 2) a heat-map of the number of $1/2$-deviating groups (of either color) that contain a precinct, with yellow representing larger counts. We observe that precincts with highest counts are located either in an urban area or in proximity of one. This phenomenon is consistent with the perceived correlation between voter distribution and type of residence [9, 41]. In states with multiple dense urban areas (NC and TX), there is sufficient flexibility in the redistricting process to "crack" a highly-concentrated urban demographic into multiple districts. In this case, the urban area may form a deviating group resulting in high numbers of unfair plans. On the other hand, for states like WI without such flexibility (fewer urban areas, or urban areas close to state boundaries), the numbers of unfair plans remain low.

We note that our visualizations – in particular, the deviating groups in unfair plans, as well as the heatmap of likelihood of unhappiness of a precinct – make the local fairness concept *explainable*.

4.3 Compatibility with Extant Fairness and Compactness Notions

The ensemble-based auditing approach can be viewed as a pruning method that identifies a subset of plans as locally fair. We now ask: Are the fair plans better or worse on extant global fairness and compactness criteria than average plans in the ensemble? Towards this end, for $c = 1/2$, we rank the plans in the ensemble by the unfairness score $\text{unf}(\Pi)$. We compare properties of the top 5% plans in the ranking (which are most locally fair) against the entire ensemble.

Seat share outcomes. For each plan, we compute *Minor%*, the percentage of seat shares claimed by the color with a minority of total voters. The resulting distribution is shown in Figure 3a. Note that the top 5% locally fair plans do just as well as or better than the entire ensemble on this metric, sometimes achieving lower variance in seat shares.

Average partisanship. For each district, we define its *partisanship* as the percentage of the total number of voters of the majority color. We then define the average partisanship of Π to be the average of partisanship values over its districts (ignoring small differences in population). Therefore, a low average partisanship (towards 50%) implies better competitiveness of the districts drawn. In Figure 3b, we compare the average partisanship among the two subsets of plans as well as a recent real-world congressional redistricting plan. The locally fair plans produce slightly more partisan districts, since this reduces the number of unhappy voters, making finding a deviating district harder. However, the (roughly) 2-3% of shift in average partisanship is small and comparable to other uncertainties (e.g., voter turnouts or year-to-year election result gaps). Furthermore, the average partisanship of the fair

plans is significantly smaller than that of the real-world redistricting plan for three states, showing that local fairness remains compatible with reasonably small partisanship.

Average Compactness. Another measure of quality of districting plans is compactness – a non-compact district not only makes less geographic sense, but is also more likely to have been gerrymandered to favor one party over another. We compute the Polsby-Popper scores [33] of districts as a measurement of compactness. For a district $D \in \Pi$, the Polsby-Popper score is defined as $4\pi A(D)/P(D)^2$, where $A(D)$ and $P(D)$ are the area and the perimeter of D, respectively. Therefore, a higher value implies a more compact district. The average compactness of Π is the average of Polsby-Popper score over its districts. In Figure 3c, we show that the average compactness of the locally fair plans remains comparable to that of the entire ensemble. Taken together, our results show that local fairness is compatible with fair seat share and compactness, while sacrificing only a small amount on partisanship of the districts.

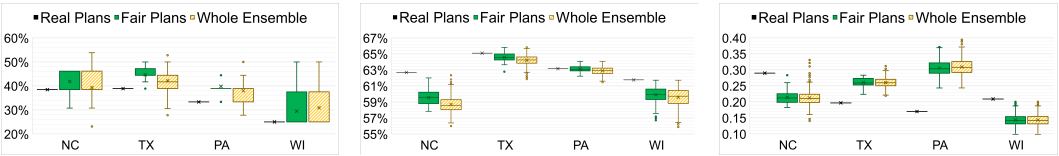

(a) Seat shares (Minor%) (b) Average partisanship (c) Average compactness

Figure 3: Distribution of fairness and compactness metrics among subsets of generated plans

4.4 Additional Results

We first compute the local fairness of the plans actually in use. Using the unf ranks, the actually implemented plan falls in the 84th percentile (fairer than 84% of the plans) in NC, 22nd percentile in TX, and 29th percentile in PA. While the NC plan is (somewhat surprisingly) above average in local fairness, it is also very partisan. In fact, as shown in Figure 3b, it is more partisan than *every* plan in the ensemble, while it is possible to find fair plans that do not make this compromise.

In Appendix D, we finally show that the locally fair notion is also robust to voting patterns in the sense that there is a large overlap between fair plans identified by auditing the ensemble using the voter data from the 2012 and the 2016 Presidential elections for the states of NC, PA, and TX.

5 Conclusions

In summary, in contrast to extant global notions of fairness in redistricting such as seat distribution or district competitiveness, the new notion of local fairness is an explainable notion based on happiness and justified complaints of populations in compact geographic regions. Our experiments show that in practice local fairness is not only possible to achieve, but choosing locally fair plans also does not come at the expense of other global fairness properties properties.

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 LF AUDITING or LFP 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 c-deviating group. We leave this as a challenging open question.

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] North Carolina v. Covington. https://supreme.justia.com/cases/federal/us/581/16-1023/, 2018.

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

[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*, 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] 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.

[11] 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.

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

[13] 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.

[14] 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.

[15] 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.

[16] M. Duchin and B. E. Tenner. Discrete geometry for electoral geography. *arXiv preprint arXiv:1808.05860*, 2018.

[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. *arXiv preprint arXiv:2107.07083*, 2021.

[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. S. Hochbaum and A. Pathria. Node-optimal connected k-subgraphs. *University of California, Berkeley*, 1994.

[26] 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.

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

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

[29] J. Levitt. All about redistricting. `https://redistricting.lls.edu/`, 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] MGGG. MGGG states. `https://github.com/mggg-states`, 2022.

[33] 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.

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

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

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

[37] 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.

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

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

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

[41] 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.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes] See Supplementary Material
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A] Datasets are publicly available.
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [No]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]