Algorithmic Redistricting and Black Representation

ZACHARY SCHUTZMAN

Abstract. In the United States, the careful crafting of electoral districts has been a powerful tool for politicians to limit groups' political power or exclude them from representation entirely, most prominently to the detriment of political parties and racial and ethnic minority communities. Beginning in the 1960s, experts began proposing algorithmic solutions to the redistricting problem, where a 'neutral' computer program could draw 'fair' districts free of human influence. Despite the traction that these proposals achieved both in academic and popular discourse, little work has been done to understand the extents to which algorithmically drawn districts do or do not comport with notions of fairness and equity in this space. In this work, we perform such an analysis, running several proposed algorithms to generate districts in Alabama and Michigan. We observe that in both of these states, all four algorithms generate plans which provide fewer districts where Black voters would be expected to decide the outcome of the election, relative to both the proportion of Black people in their populaces as well as to the number of Black opportunity districts in the plans actually enacted by their legislatures. We conclude with some discussion about the role of algorithms in redistricting moving forward, and how these tools might be used to enhance, rather than restrict, the ability for various communities to achieve elected representation.

1 INTRODUCTION

The practice of *gerrymandering*, politicians intentionally manipulating the boundaries of electoral districts to help allies or hurt rivals trying to achieve political representation is a practice as old as the United States itself. Over the centuries, abuse of the redistricting process has been used to help or hurt political parties, racial and ethnic groups, and individual representatives in their quest for political representation in legislative bodies from the U.S. Congress to local municipal and county governments. Beginning in the 1960s, advocates for 'fair' redistricting have proposed the use of computer algorithms to generate districting plans. In 1961, the economist William Vickrey proposed a framework in which an algorithm, given only the most basic data needed to draw districts which are geographically connected and balanced in population, could use opaque processes and randomness to separate the construction of districting plans from human influences entirely [18]. Almost immediately after, researchers and scientists designed theoretical and practical algorithms to achieve this, following a general framework of building algorithms to draw districts which maximized some notion of 'compactness', a family of measures which describe the geometric regularity of a district, subject to the constraints of geographic connectedness and balanced population. Before the 1960s, states had an enormous amount of leeway in how they designed and implemented districting plans, particularly state legislative districts, and through the 1940s, the Supreme Court affirmed this right [8]. Many states neglected to redistrict at all in the first half of the twentieth century, even as residents moved from rural regions to urban areas. According to Rep. Morris Udall, at the time the Supreme Court intervened in the early 1960s to demand that congressional and state legislative districts contained nearly equal numbers of people, there were some state legislative chambers where the largest district had hundreds or even thousands of times more voters than the smallest one [16].

It was in this environment of malapportionment that Vickrey and others made their arguments. The foremost abuse in the process was allowing districts which were nowhere close to being equal in population and so centering the design of algorithms around that as a hard constraint. However, despite these being very real and salient issues at the time, population imbalance is not a central concern nowadays, and the discussion of gerrymandering and unfairness in redistricting now revolves around racial and partisan inequities, preserving political units like municipalities and counties, and issues of incumbency. Despite this, the classical framework of algorithms which optimize for compactness subject to population balance and geographic connectedness still persists.

Here, we take a narrow perspective of the impacts of algorithmic redistricting on these issues and consider the extent to which various algorithms proposed by researchers and experts do or do not draw plans which contain districts that allow Black voters the *opportunity to elect candidates-of-choice*, or *opportunity districts*, in two states with significant but geographically very different Black populations: Alabama and Michigan. Advocates for algorithmic redistricting argue that these so-called 'neutral' algorithms will generate districting plans which are 'fair' with respect to these other criteria, arguing that because the algorithm cannot see any additional data, for example the partisan composition of a region or the geographic distribution of people of different races, then the algorithm cannot generate a plan which is unfair on a partisan or racial dimension. This is, however, unsupported by any mathematical or experimental evidence and we understand very clearly from other domains that simply omitting a variable from some algorithm or model does not guarantee that the output will be uncorrelated with that variable. Indeed, our results show that the algorithm-generated districting plans clearly and consistently contain fewer Black opportunity districts than either the enacted plan currently in use or what a standard of proportionality might demand.

Related Work. Despite sixty years of work designing and implementing redistricting algorithms, little attention has been paid to examining and comparing the districting plans the algorithms and generate along dimensions of social factors like how well they draw districts which afford various communities the ability to elect their candidate-of-choice. For an overview of redistricting algorithms, see [3] and the introduction of [11]. For some broader discussion about the trade-offs between Vickrey's notion of fairness as arising from procedural neutrality and the idea of evaluating the fairness of a districting plan based on the slate of candidates expected to be elected from it, see [15].

In this work, we compare the outputs of four algorithms selected broadly from the domain of algorithmic redistricting. We discuss these algorithms in more detail in Section 2.2. None of the sources of these algorithms carefully consider the impacts on any community of the generated districts. [14] asserts at the outset that the best districting plan is the one which maximizes compactness and intentionally ignores any of the underlying social factors. [7] and [12] do not address these factors at all. The authors of [7] mentions the dichotomy between fairness arising from the neutrality of process and definitions of fairness being applied to the outcome but does not discuss it further and does not perform any evaluation of the plans output by their algorithm; they primarily approach the problem as one of computational geometry. The authors of [11] offers some discussion of notions of fairness, but they also do not evaluate the hypothetical outcomes of their districting plans. Their discussion is largely about how to balance other ideas, like preserving communities or municipalities and how to adjust the input data to force the algorithm to take those as constraints.

The closest related work is [5] which considers the legal implications of using algorithms to provide a 'neutral baseline' of representation along partisan and racial axes. They find that doing so would broadly decrease the ability of Black and Hispanic/Latino communities to exert influence in state legislatures. While our findings are entirely consistent with theirs, we ask a different question and use a different methodology. They compare the use of a single algorithm on a breadth of states while we examine the outputs of several algorithms on only two states. Additionally, their concern is primarily with the average behavior of their algorithm, e.g. the expected number of districts where Black voters can elect a candidate-of-choice, while we examine more closely the plans our selection of algorithms draw, e.g. the extent to which the algorithms concentrate the power of Black voters in few districts or dilute it across many. The authors of a response to this article highlight some methodological concerns and run additional experiments which offer some robustness to the former's findings [9].

1.1 Our Contributions

In this work, we make the first comparison of the expected impact of redistricting algorithms on the ability for Black voters to elect candidates-of-choice to the legislature. In initiating this study, we begin with a fairly narrow scope, considering hypothetical state senate districts in Alabama and Michigan, examining the extent to which four algorithms selected from both scientific literature and open source software do or do not draw districts in which Black voters potentially have sufficient political power to decide the outcome of the election. In doing this, we show empirically that on both states, all four of these algorithms generate fewer Black opportunity districts than the enacted plan (as of 2021) provides and what a standard of proportionality would demand. We additionally discuss some of the patterns that emerge in these outcomes despite the differences in the specifications of the algorithms as well as the kind of districts they draw. We conclude by revisiting the motivation for using algorithms to draw electoral districts and, informed by the results and analysis in Section 3, offer some ideas for where these computational tools may be used in a way that empowers marginalized communities in the redistricting process, rather than enforcing and legitimizing detrimental outcomes.

2 DATA & ALGORITHMS

2.1 Data

In order to run the algorithms themselves, we need relatively little data. Each algorithm draws connected, population-balanced districts composed of *census blocks*. For this, we can use data provided by the U.S. Census Bureau, which provides files containing both the geography of the census blocks as well as demographic data, for each block [17]. This demographic data includes the total population of each block, which is necessary to run the redistricting algorithms as well as the Black population and Black population over the age of 18, which are necessary for the subsequent evaluation of the generated districts. The block geography and demographics come from the 2010 decennial census data which is used as the official redistricting data until it is superceded by the 2020 census data, which as of this writing has not been released.

In this paper, we consider two states, Alabama and Michigan, and use the algorithms to draw state senate districts for each state. According to the 2010 data, Alabama has a population around 4.73 million and is 27 percent Black. The state is afforded 35 state senate districts with approximately 136,000 people per district. Figure 1a shows the geographic distribution of Black people in Alabama. Birmingham and Montgomery in the center of the state and Mobile in the southwest all majority-Black cities of approximately 200,000 people. The wide strip across the center of the state, sometimes called the "Black Belt" contains Montgomery as well as many smaller predominantly Black cities and a large rural Black population. In aggregate, this region is over 50 percent Black.

The 2010 population of Michigan is approximately 9.93 million and is 14 percent Black. The state has 38 state senate districts, each containing around 260,000 people. Figure 1b shows the geographic distribution of Black people in Michigan. In stark contrast to Alabama, the Black population is overwhelmingly concentrated in and around Detroit in the southeast part of the state. Wayne County, which contains Detroit, has a total population of roughly 1.8 million, 33 percent of which is Black. There is a sizeable Black community in the city of Flint in the center of the state, but the population is relatively small compared to that of Detroit, where approximately 80,000 of the 420,000 residents of Genesee County are Black. The remainder of the state is overwhelmingly white.

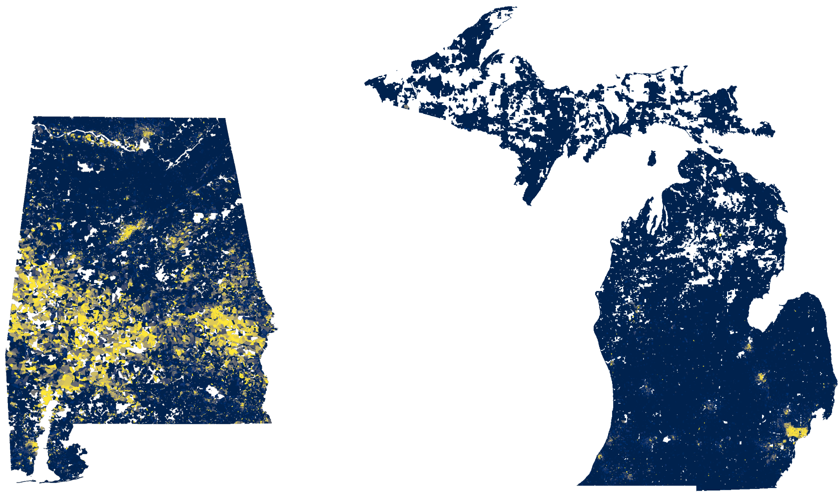

(a) Black population distribution in Alabama

(b) Black population distribution in Michigan

Fig. 1. The Black population distributions for Alabama and Michigan. The lighter, yellow regions have more Black residents.

We choose to look at state senate rather than the more widely discussed U.S. Congressional districts in order to potentially generate a richer set of outcomes. For instance, Alabama is afforded seven members in the U.S. House of Representatives. Its current plan contains one Black opportunity district and experts and stakeholders have argued and sued for the state to redraw the plan to include a second. The algorithms will typically find zero or one, with the rare possibility of two, opportunity districts at this scale. In contrast, the enacted state senate plan contains eight Black opportunity districts, and drawing as many as 11 may be possible, so working at this scale allows for a wider range of possible outcomes and a more careful inspection and analysis of the behavior of the algorithms.

We finally describe precisely what a 'Black opportunity district' is. More broadly, a district provides some group or community the *opportunity to elect a candidate-of-choice* if the group in question votes cohesively, and therefore 'candidate-of-choice' is a relevant concept, and possesses sufficient political power to ensure that candidate wins an election. Throughout the U.S., including in Alabama and Michigan, Black voters act overwhelmingly cohesively, where upwards of 90 percent of Black voters support Democratic candidates in competitive elections [6]. However, the proportion of Black voters a district must contain in order to properly constitute an 'opportunity district' is not the same in every state or locality. In Alabama, while there is a small portion of white voters who might reliably support the Black-preferred candidate, white voters in Alabama overwhelmingly support Republicans. Experts estimated this level of white support for Black-preferred candidates in the center of the state to be roughly 17 percent. Therefore, while a district need not be strictly majority Black in order to reliably elect the Black voters' candidate-of-choice, it needs to be fairly close to that, where perhaps 45 percent might be sufficient. More precisely, we would want a district where 95 percent of the Black voting age population together with 17 percent of the white population constitute a clear majority in the district in order to safely consider it a 'Black opportunity district'.

Michigan is very different from Alabama in that there is a much larger proportion of white Democrats. Estimates based on the 2016 U.S. Presidential election and the 2018 U.S. Senate election in the state place the proportion of white Democrats around 45 percent. For this reason, if Black voters' candidate-of-choice wins the Democratic primary in a reliably Democratic district, we can expect white Democrats to support this candidate in the general election and it is therefore the strength of Black voters in the *primary* that determines whether a district is or is not a Black opportunity district. In Michigan, we use election data [2, 13] to estimate the partisan lean of a district and we consider one to be a 'Black opportunity district' if it both reliably will elect a Democratic candidate in the general election *and* Black voters constitute a majority of the estimated Democratic primary electorate.

2.2 Algorithms

We now briefly describe the four algorithms considered. At a high level, all of these algorithms are designed to take as input a collection of geographic units, such as census blocks, each equipped with a total population. They then solve a geometric optimization problem to draw districts which are connected, nearly equal in population, and as compact as possible, where the definition of compactness is either explicitly or implicitly encoded in the objective of the optimization problem.

Annealing. The first algorithm, which we call *Annealing*, was designed by a software engineer and its source code is available publicly online [14]. It is well-discussed in public discourse, having featured in a *Washington Post* article [10] and used in *FiveThirtyEight*'s *Atlas of Redistricting* project [4]. The algorithm works by choosing random district centers and assigning each census block to its closest center. Districts which are underpopulated incorporate nearby blocks from districts which are overpopulated, then the centers are recomputed. This annealing process of grabbing blocks repeats until the districts are population-balanced. The entire process, starting from new random centers, is repeated numerous times and the plan which generates the most compact districts is returned as the final output.

Arcs. We call the second algorithm *Arcs* due to its unique characteristic of drawing districts bounded by circular arcs. This algorithm appears in a recent paper [11] and while it draws inspiration and lessons from several existing algorithms, the resulting districts are highly unlike those of any other algorithm. To draw k districts algorithm works by selecting a corner of the bounding box of the state and drawing the circular arc centered there which splits the state into two pieces which support $\lceil \frac{k}{2} \rceil$ and $\lfloor \frac{k}{2} \rfloor$ districts. Then the algorithm is run recursively on each half, selecting a corner of the bounding box for each. The sequence of bounding box corners which is ultimately selected is the one which maximizes the compactness of the districting plan.

Voronoi. Algorithms using Voronoi diagrams and related procedures to generate districting plans have been proposed since at least as early as the mid-2000s and we consider the most recent iteration in this line of work [7] which uses a generalization of Voronoi diagrams called *power diagrams* to partition a state into districts. The algorithm to draw k districts works as follows. First, k random points $c_1, c_2, \ldots, c_k$ are chosen and a circle with radius $r_1, r_2, \ldots, r_k$ is associated to each point. Each census block is some distance from the boundary of every one of these circles, and we assign a block to the district corresponding to its nearest circle. The algorithm then iteratively recomputes the centers c_i to be the geographic center of the blocks assigned to it as well as adjusts the radii r_i to increase or decrease the population of each district. Once both the centers and radii converge to some local minimum, the algorithm halts and returns a districting plan.

Tree. The final algorithm we consider is a highly randomized one which appears in the software package `gerrychain` [12], a suite of algorithms used to generate large ensembles of districting

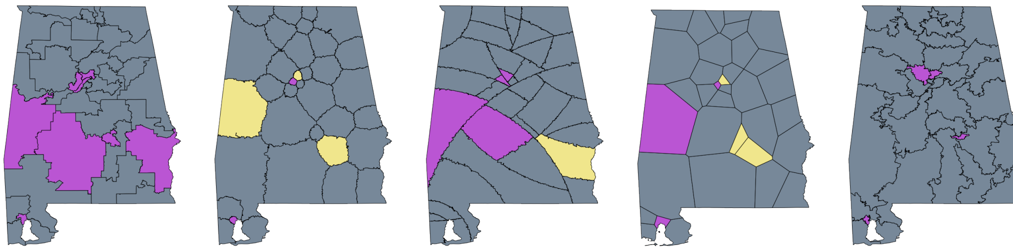

(a) Enacted (b) Annealing (c) Arcs (d) Voronoi (e) Tree

Fig. 2. The enacted state senate districts in Alabama and the four algorithmically generated plans. Pink districts have at least 55 percent expected support for the Black-preferred candidate; yellow districts have between 50 and 55 percent expected support.

plans, against which to compare a proposed or enacted plan. The algorithm works by recursively drawing *spanning trees* in the underlying adjacency graph of the census blocks. To construct this graph, represent each census block as a vertex and an edge between two vertices exists if and only if the two corresponding census blocks share a geographic boundary. The algorithm then draws a random spanning tree of this graph and chooses an edge to delete, splitting the single tree into two such that the total population in one of these trees is the appropriate size for a single district. This is then frozen as a district, the spanning trees are forgotten, and a new tree is drawn over the remainder of the state. Once all districts have been drawn, the algorithm terminates and returns a completed plan.

3 RESULTS

3.1 Alabama

Here we present the enacted and algorithmically generated plans for Alabama, depicted in Section 3.1. We can see that visually all of these plans are quite distinct, reflecting the differences in specification between the algorithms. Using the benchmark of 95 percent of Black voting age population plus 17 percent of white voting age population exceeding 50 percent, we can examine how many Black opportunity districts appear in each plan, the geographic location of these districts, and how reliably they could be expected to elect Black voters' candidate-of-choice.

Beginning with the enacted plan (Figure 2a, Figure 3a), we have eight districts where the expected support for Black voters' preferred candidate exceeds 50 percent. Three of these districts are anchored in Birmingham, one each in Mobile and Montgomery, and three across the more rural central region of the state. We can plot the expected support for the Black-preferred candidate in these eight districts as well as the next four districts with the highest Black population and observe a clear demarcation between these eight and the next four (as well as the remaining 23 not shown), where the eight district has over 55 percent expected support for the Black-preferred candidate while the ninth has this value well below 40 percent.

In the Annealing plan (Figure 2b, Figure 3b), we see a stark difference. This plan contains only two districts with expected support for the Black-preferred candidate over 55 percent and an additional three where the expected support falls between 50 and 55 percent. These five districts include two in Birmingham, one in Mobile, one on the periphery of Montgomery, and one large rural one in the west, and it is straightforward to see how this plan 'misses' several of the opportunity districts in the enacted plan. Plotting the expected support from the 12 districts with the highest Black

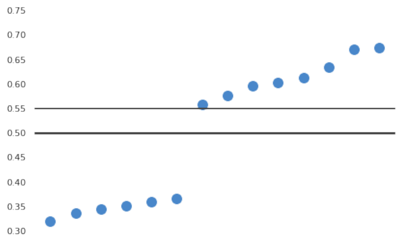

(a) Enacted

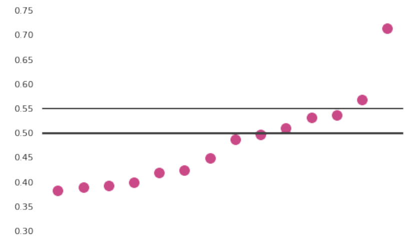

(b) Annealing

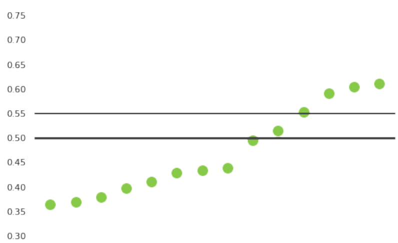

(c) Arcs

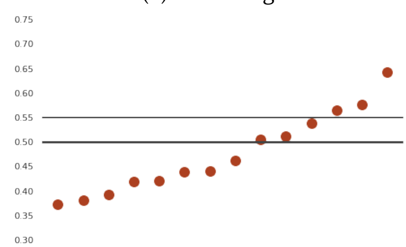

(d) Voronoi

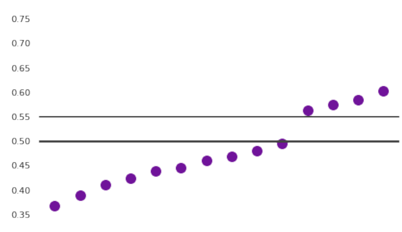

(e) Tree

Fig. 3. Plots of the expected support for the Black-preferred candidate in the 14 districts with the highest Black population in Alabama. The horizontal axes are the sorted districts and the vertical axis is the expected fraction of votes the Black-preferred candidate wins.

population, we can see the dramatic gap between the eight and ninth districts in the enacted plan does not appear in the Annealing plan. Rather, the there is a gentle decline in expected support for the Black-preferred candidate from the second district onward. Relative to the enacted plan, the high-support districts underperform and the low-support districts overperform, together suggesting this plan is very effective at diluting Black voting strength.

For the Arcs plan (Figure 2c, Figure 3c), we see something similar. This algorithm again draws five potential opportunity districts with four expected to fall above 55 percent support for the Black-preferred candidate and one in the intermediate 50 to 55 percent range. This algorithm finds different opportunity districts from the Annealing ones. Here we have again two around Birmingham and one rural one in the west, but also two additional large rural districts in the middle of the state. Once again we can examine the plot of expected support for the Black-preferred candidate by district and observe the steady decline in expected support as in the Annealing plan rather than the abrupt leap as in the enacted plan, once again suggesting that the algorithmically generated plan is diluting Black voting strength. The Voronoi plan (Figure 2d, Figure 3d) performs similarly to the Annealing plan, but with an extra district in the 50 to 55 percent range around Montgomery and slightly higher expected support for the Black-preferred candidate in the rural

(a) Enacted (b) Annealing (c) Arcs (d) Voronoi (e) Tree

Fig. 4. The enacted state senate districts in Michigan and the four algorithmically generated plans. Pink districts have at least 55 percent expected support for the Black-preferred candidate in the Democratic primary; yellow districts have between 50 and 55 percent expected support.

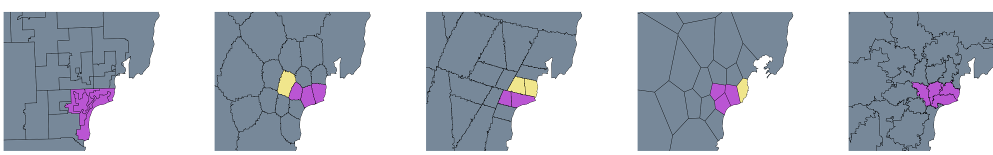

(a) Enacted (b) Annealing (c) Arcs (d) Voronoi (e) Tree

Fig. 5. Insets of the Detroit region for the plans shown in Section 3.2.

district in the west. Once again, the plot of expected support indicates dilution of Black voting power.

The Tree algorithm draws a plan (Figure 2e, Figure 3e) with four districts above the 55 percent expected support level and none in the 50 to 55 percent range, which is a departure from the previous outputs. It also does not draw any rural opportunity districts, finding two in Birmingham and one each in Montgomery and Mobile. Plotting the expected support level shows, however, that the effects of this plan aren't too different from the other three algorithmic ones. We once again see a steady decline in expected support with a small gap between the fourth and fifth districts. However, we still see districts falling just below the 50 percent line, indicating that Black voters in these districts might be narrowly shut out of political power, which is one of the consequences of the kind of vote dilution present in these algorithmically generated plans.

3.2 Michigan

We repeat the same plan generation and analysis for Michigan. Again, the four algorithmic plans are highly visually distinct both from each other as well as from the enacted plan which was designed to follow the rectangular county boundaries of the state. Michigan does not have the same level of racially polarized voting as Alabama in the general electorate, so the determination of whether or not a district might be a Black opportunity district will be different. In Michigan, approximately 90 percent of Black voters support the Democratic Party candidate in recent statewide general elections. On the other hand, about 45 percent of white voters do as well. This means that in a reliably Democratic district, if Black voters' preferred candidate wins the Democratic primary, that person is highly likely to win the general election as well, regardless of any racially polarized voting patterns in the primary. Therefore, in assessing whether or not a district provides Black voters the opportunity to elect a candidate-of-choice, the district must be *both* a reliably Democratic district *and* one in which Black voters are expected to comprise a majority of the Democratic

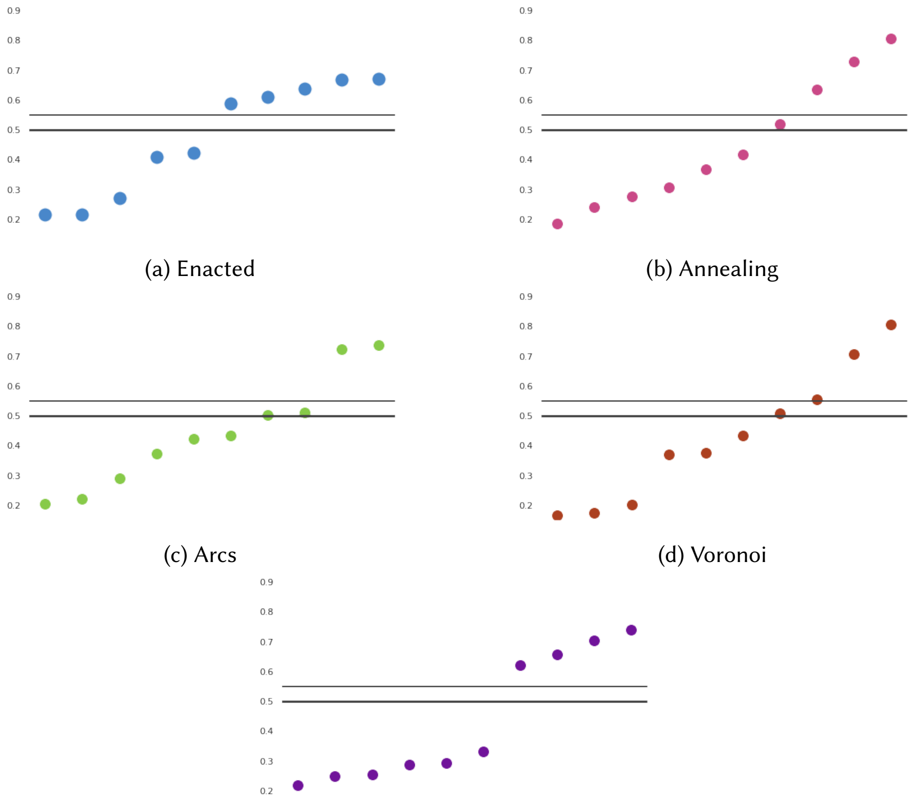

(a) Enacted (b) Annealing

(c) Arcs (d) Voronoi

(e) Tree

Fig. 6. Plots of the expected support in the Democratic primary for the Black-preferred candidate in the 10 Democratic districts with the highest Black population in Michigan. The horizontal axes are the sorted districts and the vertical axis is the expected fraction of votes the Black-preferred candidate wins.

primary electorate. To perform these estimates we use precinct-level election data from the 2016 U.S. Presidential and 2018 U.S. Senate elections in Michigan.[2, 13] Evaluating against these two elections gives slightly different quantitative estimates for Black turnout but qualitatively the results and conclusions are extremely similar. In this section, we present results with respect to the Presidential election and include the corresponding figures for the Senate election in Appendix A.

The enacted plan (Figure 4a, Figure 5a, Figure 6a) contains five districts which meet the criteria of being considered Black opportunity districts. All five of these are anchored in the city of Detroit in the southeast part of the state, where the majority of Michigan's Black residents live. Furthermore, all five are overwhelmingly Democratic and all have elected a Black state senator in recent years, although only three of the five are majority Black. Being that Michigan is approximately 12 percent Black, a standard of proportionality would ask for four or five Black opportunity districts. The four algorithmically generated plans underrepresent Black voters relative to both the outcomes of the enacted plans and the standard of proportionality. Unlike all of the potential opportunity districts the algorithms find in Michigan are in and around Detroit, in contrast to Alabama where we were looking for opportunity districts around cities as well as in rural areas. The Annealing (Figure 4b, Figure 5b, Figure 6b) and Voronoi (Figure 4d, Figure 5d, Figure 6d) plans contain three clear opportunity districts and one more where Black voters barely constitute a majority of the expected Democratic primary electorate; the Arcs plan (Figure 4c, Figure 5c, Figure 6c) contains two clear opportunity districts and two more with a bare majority; the Tree plan (Figure 4e, Figure 5e, Figure 6e) comes the closest to the enacted plan with four clear Black opportunity districts.

Similar to Alabama, some of the plots in Section 3.2 show a steady decline in Black voting strength when stepping through the districts from right to left, and we see a similar effect of the algorithms drawing high-concentration districts where the number of Black voters dramatically exceeds the number that would be necessary to simply expect a safe victory in that district and low-concentration districts where the number of Black voters is clearly insufficient to elect a candidate-of-choice. Compared to the enacted plan as well as the proportionality benchmark, this results in fewer Black opportunity districts. However, some contain a 'gap' where there is a significant decrease in Black voting strength between one district and the next. These plans still, however, underrepresent Black voters. In the enacted plan, we see a gap between the fifth and sixth districts where the first five districts are clearly Black opportunity districts and the sixth is not. In the Arcs plan (Figure 6c), the gap occurs after only the second district and in the Tree plan (Figure 6e), the gap occurs after the fourth district. The algorithmically drawn plans in Michigan also all have at least two districts where Black voters comprise at least 70 percent of the Democratic electorate, which differs from the plans in Alabama where only one plan had a single such district. This 70 percent level is clearly much higher than is necessary to call these 'opportunity districts'. By overconcentrating Black voters in these districts, the algorithm is reducing their ability to exert influence in neighboring ones, therefore reducing the number of opportunity districts in the plan overall. This action of concentrating some voting bloc into very few districts to reduce their influence elsewhere is called 'packing-and-cracking', and, when done intentionally, is a key strategy in partisan and racial gerrymandering.

4 DISCUSSION & FUTURE DIRECTIONS

4.1 Discussion

In the broader discourse of algorithmic redistricting, much as in other subject area domains, there are people suggesting that computer algorithms totally replace the human decisionmaking components in the process much in the way that Vickrey proposed. The results shown here that drawing districts in this way could have severe detrimental effects on the political influence of already marginalized groups and proposals to use algorithms to draw districts should be treated with skepticism and scrutiny in the face of the assertion that because they are unaware of some social or demographic factor such as race that they are therefore unable to create outcomes which are unfair with respect to that feature.

Another way some propose to use algorithms in redistricting is as a way of informing human decisionmakers about what is 'possible' or 'typical'. In the relatively new area of *ensemble analysis*, an algorithm is used to draft hundreds or thousands of districting plans and then an analyst can evaluate this sample of plans and ask questions like 'What is the average number of districts won by Republicans across all the plans?' The article by Chen and Stephanopoulos [5] and response by Duchin and Spencer [9] take this approach and find, similarly to our results, that such a procedure may generate plans that, even on average and in aggregate, underrepresent minority racial and ethnic groups relative to both the enacted plans and a standard of proportionality, similar to our results in this work. The concern is therefore not algorithms *replacing* human decisionmaking and analysis but rather algorithms *legitimizing* unfair or discriminatory decisionmaking. One could imagine a legislature intent on suppressing Black political strength appealing to "A neutral algorithm drew 10,000,000 plans and none of them had more than one majority Black district" as a way to justify and legitimize enacting a plan which underrepresents Black voters. Indeed, Judge Easterbrook [1] advocated for this kind of analysis to be used as the baseline in legal contexts and these analyses seek to understand the ramifications of adopting this framework.

Looking forward to the future of algorithms and redistricting, it is not the case that algorithms and computerized map-making are totally useless and harmful to the process of redistricting, rather the harm stems from human decisionmakers intentionally or unintentionally giving too much deference to their outputs. In principle, a more deliberate approach to the human-computer interface in this domain could be used to make the redistricting process more transparent, fair, and accessible. Computers are very bad at inferring human values; as we saw in this work, a race-unaware algorithm should not be expected to draw districts which achieve any sort of notion of fairness or equity with respect to race. To further compound this, humans are very bad at fully articulating all the features and facets of a districting plan we find desirable or undesirable, especially in a way that is precise and mathematical enough to be encoded in a computer algorithm. However, computers are very good at solving mathematical problems and searching for plans or districts which meet particular criteria. Rather than using algorithms to generate plans to enact, legislators, stakeholders, and the public could use algorithmic redistricting tools to explore and understand trade-offs and the frontiers of what is and is not possible for a districting plan to accomplish in a particular jurisdiction. One could ask an algorithm whether it is possible to draw a state senate plan in Alabama with ten Black opportunity districts, look at the resulting plan, and ask another question based on relevant community feedback, such as whether or not it is possible to draw ten opportunity districts *and* keep the city of Tuscaloosa entirely within one district. In this way, algorithms become a tool as part of an iterated discussion about social and political values rather than the arbiters of fairness.

We briefly conclude by highlighting some future areas of computational research aimed at better understanding the results in this work and making the findings more robust. First, extending this kind of analysis to more than two groups would be extremely valuable in extending these findings to jurisdictions with greater racial and ethnic diversity. That the algorithms underrepresent Black voters in both Alabama and Michigan despite these two states having extremely different racial geography is interesting and warrants more precise study. By just relabelling the racial data in the experiments in this work, one could concoct districting plans in which there are a number of Black opportunity districts equal to the proportion of Black people in the state or even overrepresent Black voters relative to this proportion, but this leads to questions about how much ‘population swapping’ must be done to achieve this as well as whether or not any real jurisdictions match the racial geography of these hypothetical versions of Alabama and Michigan. Finally, one could extend this work by exploring other variants of ‘neutral’ algorithms. We worked with purely geometric formulations here, but other principles of redistricting, such as minimizing the extent to which counties and municipalities are split, to be ‘neutral’ criteria as well, and extending this analysis to incorporate those features may either provide robustness to our findings or highlight a salient difference between the two approaches.

REFERENCES

[1] 2008. Gonzalez v. City of Aurora, Illinois. , 594 pages.

[2] 2018. https://www.michigan.gov/sos/

[3] Amariah Becker and Justin Solomon. 2020. Redistricting Algorithms. *arXiv preprint arXiv:2011.09504* (2020).

[4] Aaron Bycoffe, Ella Koeze, David Wasserman, and Julia Wolfe. 2018. The Atlas Of Redistricting. https://projects.fivethirtyeight.com/redistricting-maps/

[5] Jowei Chen and Nicholas Stephanopoulos. 2020. The Race-Blind Future of Voting Rights. *Yale Law Journal, Forthcoming, Harvard Public Law Working Paper* 20-22 (2020).

[6] Chestnut v Merrill. 2020. *United States District Court For The Northern District Of Alabama Southern Division.*

[7] Vincent Cohen-Addad, Philip N Klein, and Neal E Young. 2018. Balanced centroidal power diagrams for redistricting. In *Proceedings of the 26th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. ACM, 389–396.

[8] Colgrove v Green. 1946. *Supreme Court of the United States*.

[9] Moon Duchin and Doug Spencer. 2021. Models, race, and the law. In *Yale Law Journal Forum*, Vol. 130.

[10] Christopher Ingraham. 2014. This computer programmer solved gerrymandering in his spare time. https://www.washingtonpost.com/news/wonk/wp/2014/06/03/this-computer-programmer-solved-gerrymandering-in-his-spare-time/?noredirect=on&utm_term=.47fccb34f63d

[11] Harry Levin and Sorelle Friedler. 2019. Automated Congressional Redistricting. *ACM Journal of Experimental Algorithms* (2019).

[12] Metric Geometry and Gerrymandering Group. 2019. mggg/gerrychain: v0.2.12. https://github.com/mggg/gerrychain

[13] Metric Geometry and Gerrymandering Group and Ruth Buck. 2019. mggg-states. https://github.com/mggg-states

[14] Brian Olson. 2013. Impartial Automatic Redistricting. https://bdistricting.com/2010/

[15] Zachary Schutzman. 2020. Trade-Offs in Fair Redistricting. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*. 159–165.

[16] Moriss K Udall. 1964. Reapportionment – I: 'One Man, One Vote' ... That's all she Wrote! https://speccoll.library.arizona.edu/online-exhibits/items/show/1759

[17] U.S. Census Bureau. 2011. 2010 Census. U.S. Department of Commerce.

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

A PLOTS FOR MICHIGAN USING SENATE DATA

In this section, we reproduce the figures from Section 3.2 using the 2018 U.S. Senate election data to make inferences about election outcomes in the algorithm-drawn districts. In 2016, Hillary Clinton, a Democrat lost the state of Michigan to Donald Trump, a Republican by less than a quarter of a percent. In 2018, the Democratic incumbent senator Debbie Stabenow defeated Republican challenger John James by about 6.5 percentage points. Broadly, Stabenow had a weaker performance than Clinton in urban Detroit but a stronger performance in the Detroit suburbs and rural parts of the state. Because of this geographical difference, and particularly because this goes hand-in-hand with an increase in white Democrats with respect to this election data, the evaluation of whether or not Black voters constitute a majority of the Democratic electorate in a hypothetical district may be different when using the 2018 Senate election data versus the 2016 Presidential election data.

The district boundaries themselves are identical, since the algorithms do not have access to the political data used to perform these analyses, but which districts are opportunity districts or not might differ. In particular, two districts in the Arcs plan and one in the Voronoi had expected support for the Black-preferred candidate between 50 and 55 percent with respect to the 2016 Presidential election results but fell below 50 percent with respect to the 2018 U.S. Senate election results, due to the minor difference in the candidates' performance.

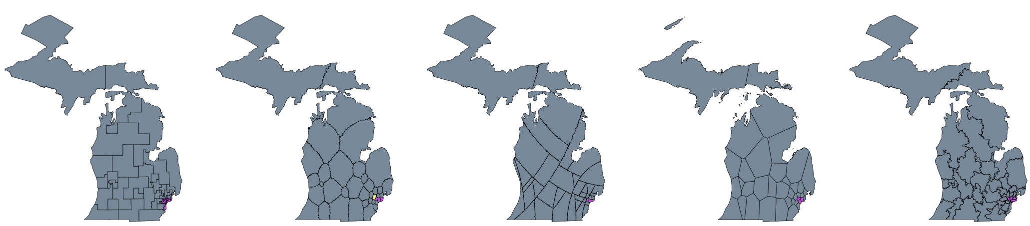

(a) Enacted (b) Annealing (c) Arcs (d) Voronoi (e) Tree

Fig. 7. The enacted state senate districts in Michigan and the four algorithmically generated plans. Pink districts have at least 55 percent expected support for the Black-preferred candidate in the Democratic primary; yellow districts have between 50 and 55 percent expected support.

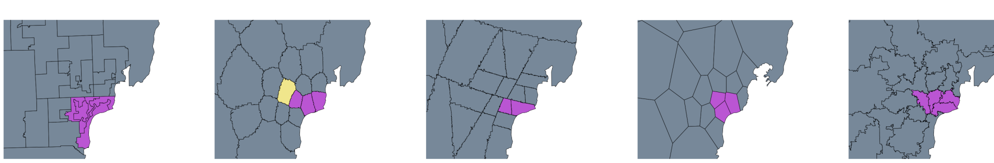

(a) Enacted (b) Annealing (c) Arcs (d) Voronoi (e) Tree

Fig. 8. Insets of the Detroit region for the plans shown in Appendix A.

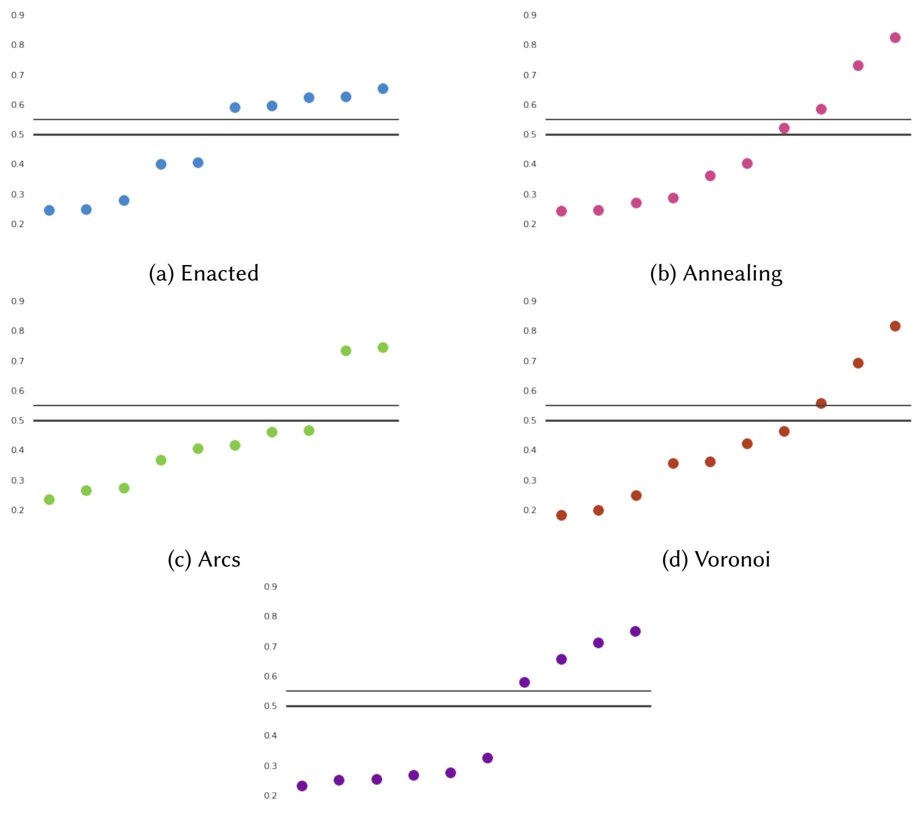

(a) Enacted (b) Annealing

(c) Arcs (d) Voronoi

(e) Tree

Fig. 9. Plots of the expected support in the Democratic primary for the Black-preferred candidate in the 10 Democratic districts with the highest Black population in Michigan. The horizontal axes are the sorted districts and the vertical axis is the expected fraction of votes the Black-preferred candidate wins. The inference is made using 2018 U.S. Senate election data.