The impact of highly compact algorithmic redistricting on the rural-versus-urban balance

Archer Wheeler
Brown University

Philip N. Klein
Brown University

ABSTRACT

It is commonly believed that, in congressional and state legislature elections in the United States, rural voters have an inherent political advantage over urban voters. We study this hypothesis using an idealized redistricting method, *balanced centroidal power diagrams*, that achieves essentially perfect population balance while optimizing a principled measure of compactness. We find that, using this method, the degree to which rural or urban voters have a political advantage depends on the number of districts and the population density of urban areas. Moreover, we find that the political advantage in any case tends to be dramatically less than that afforded by district plans used in the real world, including district plans drawn by presumably neutral parties such as the courts. One possible explanation is suggested by the following discovery: modifying centroidal power diagrams to prefer placing boundaries along city boundaries significantly increases the advantage rural voters have over urban voters.

ACM Reference Format:
Archer Wheeler and Philip N. Klein. 2020. The impact of highly compact algorithmic redistricting on the rural-versus-urban balance. In *Proceedings of ACM Conference (Conference'17).* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

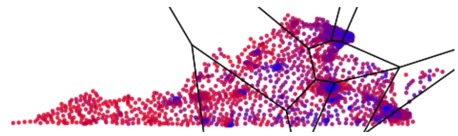

Figure 1: Algorithmic redistricting of Virginia using balanced power diagrams and populations from the 2010 census. Each dot represents the results of a precinct in the 2016 election with a color gradient corresponding to the outcome.

1 INTRODUCTION

Representatives to the U.S. House of Representatives and to many state legislative bodies are selected by winner-take-all elections across districts in states. A *district plan* for a state is a partition of the state's map into regions, called *districts*. A state's districts should be close to equal in population. Moreover, districts are expected to be *compact* and *contiguous* (notions that are not formally defined in the law). It is well known that district plans have been engineered to provide advantage to individual candidates or to parties (this is called *gerrymandering*) [3, 11, 23]. Gerrymandered districts can lead to the advantaged person or party being less responsive to voter preferences.

Voters in rural areas and voters in urban areas tend to vote for opposing parties, in the US and elsewhere [19]. It is considered well-established that geography—what parts of the map are urban and what parts are rural, and how many people live in each—has a major impact on the relative electoral success of rural voters versus urban voters. Rodden [19] has written the definitive work on the phenomenon, addressing its historical origins and its implications for the present. While he clearly acknowledges the role of gerrymandering, he convincingly argues that the rural-voter advantage is inherent in the geography—the dense packing of left-leaning voters into urban areas, and the dispersion of right-leaning voters through the larger rural areas. Rodden suggests that "a party-blind process that produces geometrically compact districts" would simply benefit the rural party, but we find the phenomenon is more nuanced.

In this paper we explore the hypothesis that rural-voter advantage is inherent in the geography, rather than being a consequence of features of specific district plans. For this exploration, we use the sort of party-blind redistricting algorithm for optimizing compactness that Rodden cautioned against. Fryer and Holden [13] state three properties that they argue any measure of compactness should satisfy, and propose a measure, RPI, that uniquely satisfies these properties. In this paper, we use a method [6] that we believe tends to find district plans that are nearly optimal with respect to RPI. We analyzed these district plans as follows: we simulated elections in a subset of U.S. states[1] based on historical electoral data, and calculated the likelihood of electoral outcomes. All results are included in Appendix A.

We find that using these compact district plans leads to elections that are significantly more competitive and exhibit less partisan advantage than existing district plans. This is true even in the case of Virginia where the existing district plan was redrawn by a court. Moreover, the partisan advantage does not consistently belong to the rural party. One possible explanation for this result is that the algorithmically generated district plans pay no attention to municipal and county lines. We found that modifying the algorithm to prefer to locate district boundaries on or near such administrative lines significantly increases the electoral advantage of rural voters. In many states, it is expected that district plans take into account these administrative lines.

To understand how election results depend on the parameters of rural-verse-urban geography, we apply the same analysis to

[1]We analyzed all states for which there are high-quality 2016 presidential precinct election results.

Conference'17, July 2017, Washington, DC, USA

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

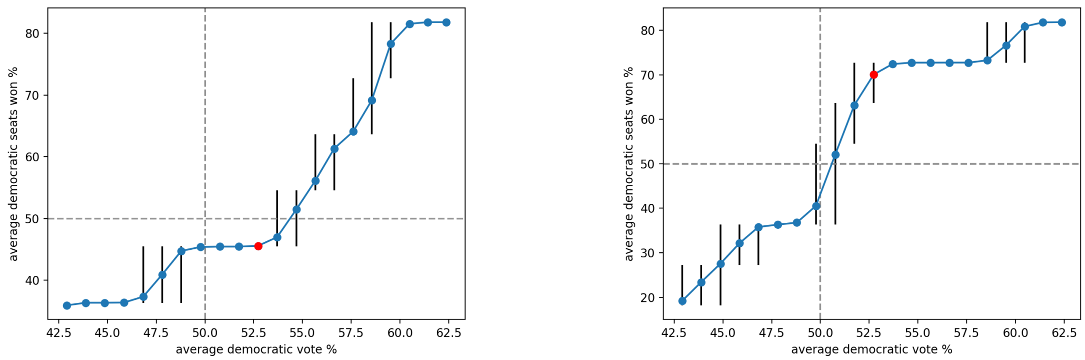

Figure 2: Left: Election results for the current Virginia districts. Right: Election results for the algorithmic district plan. Both simulations were modeled using the 2016 presidential election. Error bars show the 95% confidence interval of all outcomes. The point marked in red is the average result for the 2016 election with no change in the popular vote.

synthetic data. This enables us to understand the effects of three factors: number of districts, urban density, and party preference distribution. We find that the most important factors are the number of districts and the population density distribution (as opposed to the party preference distribution). When the number of districts is below a threshold (around five), the rural party has an advantage regardless of other parameters. Above that threshold, contrary to what one might expect, increasing urban population density *advantages* the urban party

2 BACKGROUND

2.1 Related Work

Significant academic work has been done addressing the related problem of measuring gerrymandering. A primary goal of such measurements has been to quantifiably measure gerrymandering in order to provide tools for legal precedent [1, 10, 15]. One popular measure is the "Efficiency Gap," which measures the number of waisted votes unneeded to achieve a majority [2]. Another approach has been to use Monte Carlo Markov Chains (MCMC) to build simulated random districts in order to show that partisan gerrymandering outcomes are nearly impossible probabilistically [10]. For instance, Massachusetts has a nonzero efficiency gap, but due to population density it is nearly impossible to build a district plan where any Republican representatives could win [9]. These approaches, however, have largely sought to measure existing districts and have not proposed algorithmic solutions to the problem. There has been some recent work to show that techniques from MCMC could be extended to automate district building [12].

2.2 Prior Work

Chen and Rodden [4] explore the electoral outcomes of algorithmic redistricting using historical data. They find, for instance, that even when using the notoriously close 2000 Florida presidential election results, Republicans would be heavily favored by their method of algorithmic redistricting. However, Chen and Rodden's method of redistricting neither maximizes an empirical measure of compactness, nor builds plans on the basis of being population balanced. [2] Their techniques are similar to that of recent MCMC work as they sample results from a family of *random* district plans. While they use randomness to show consistency, their distribution of compact districts is heavily determined by their model which does not use a rigorous definition of compactness. We expand on this discussion in Section 7. We hope to further this line of research by exploring electoral trends inherent to compactness.

Borodin et al. [3] introduce the concept of gerrymandering power as it relates to the geographic rural-verse-urban advantage. They show that maps drawn to *intensionally* advantage rural voters favor them more than the opposite advantage urban voters receive when maps are drawn in their favor. However, Borodin et al. use a simplistic model where voter locations are confined to a grid and capture no notion of density [3].

Previous work by Cohen-Addad, Klein and Young used balanced centroidal power diagrams to algorithmically build compact districts [6]. The system uses a capacitated k-means algorithm to build districts by minimizing the sum of distances between residents within each district. By minimizing this measure, the algorithm also maximizes the RPI compactness score of the district [7, 13]. Their hope is that compact districts are inherently fair and difficult to gerrymander [6]. In this paper we use some of the same techniques to examine resulting election outcomes.

3 GEOGRAPHIC ADVANTAGE

It is a hotly debated topic over what constitutes a fair, or even unbiased, district plan [13, 14]. However, most experts believe that plans should at least be compact and preserve certain communities. [13, 14]. We believe it is foolish to mathematically define what constitutes fair redistricting, and any plan would likely be subject to unique local political issues. However, we posit that it may be possible to algorithmically design districts in a manner which is *unbiased* to certain political variables. Of note, it is possible that a district plan drawn in an unbiased manner may still result in

[2] Chen and Rodden do ensure districts are relatively close in population by swapping border precincts after the principle construction of the districts.

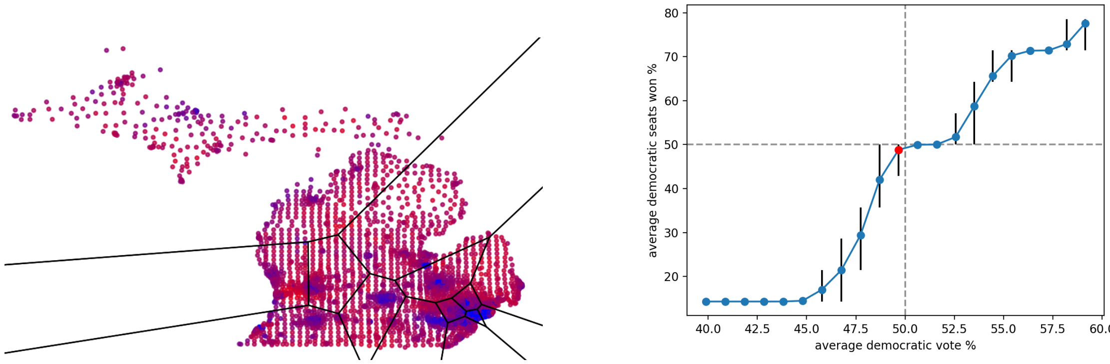

Figure 3: Algorithmic redistricting and election outcomes for Michigan.

communities or groups being unfairly advantaged. Our goal, in fact, is to explore the degree of political advantage that results from unbiased compact redistricting.

In this paper, we build districts to maximize compactness. This approach only takes into account the locations of residents and has no notion of political affiliation or ethnicity. It is in this manner that we describe compactness to be unbiased. We define the inherent geographic advantage of compact redistricting to be the expected party based representation of the optimally compact district plan. By considering the outcome of such a district plan, one can determine wether compactness inherently advantages a particular party. See Figure 2 for an example of outcomes of compact redistricting in Virginia. We do not claim that our algorithmically generated district plans should be adopted: merely that they illuminate political trends inherent to compact redistricting.

4 METHODS

We use a measure of compactness, RPI, proposed by Fryer and Holden [13] based on the locations of residences within a state. Under this measure, an optimally compact district plan is one that minimizes the sum of mean squared distances between voters in each district.

There have been many proposed quantitative measures of compactness [8, 14]. We find that techniques which measure border length are highly sensitive to geographic features such as rivers or state boundaries. Since our goal is to measure rural-verse-urban advantage, we think our measure of compactness should be defined by resident locations rather than the shape of resulting districts. Fryer and Holden argue that RPI is equivalent to any measure based on resident locations which maintain three desired properties [13].

Grouping residents and minimizing mean squared distance, however, reduces to an NP-hard problem, k-means. In practice, Lloyd's algorithm is a well known approach to solving k-means which is efficient and suspected to find near-optimal solutions [7]. For redistricting, however, one needs to solve a more complicated capacitated version of k-means [7]. To solve this problem we build on earlier work by Cohen-Addad, Klein and Young which presents a modified version of Lloyd's algorithm using balanced power diagrams [7]. The resulting district plans are similar to weighted voronoi diagrams.

4.1 Construction

We use 2010 census block data to build United States House of Representative districts using the power diagram capacitated k-means algorithm. We assume that all voters within a census block live at the centroid of the block's polygon. While this is not true in practice, census blocks are the finest grained population information released by the US census. In cities, for instance, they are often the size of a city block. In fact, many previous works use significantly more course datasets for district building such as precincts [4, 10, 12]. In the real world a district could hypothetically be drawn to split a census block, but we have no way of knowing where within a block residents live. We do allow for some census blocks lying on the border of districts to split their population. However, Cohen-Addad, Klein & Young show that this can be fixed in practice with small perturbations in blocks along the border while still achieving perfect population balance [7]. Additionally, the population counts provided by the census are statistically perturbed to preserve privacy [21]. Therefore we think it is reasonable to split a small number of blocks as long as it is not done in a biased way.

While our balanced power diagram technique does use a randomized start, it consistently finds the same result. For instance, over a hundred runs on Virginia our algorithm found 99 identical results and one result with slightly different weights, but which was nearly indistinguishable. The district plans differed only in handling a tiny number of census blocks along the border. This property is consistent across different states. Because of this, we suspect that the algorithm is finding a nearly optimal solution.

4.2 Simulating Elections

Voting results, however, are not reported by census block. For historical voting outcomes we use precinct level results since they are the most geographically fine grained data available. We used

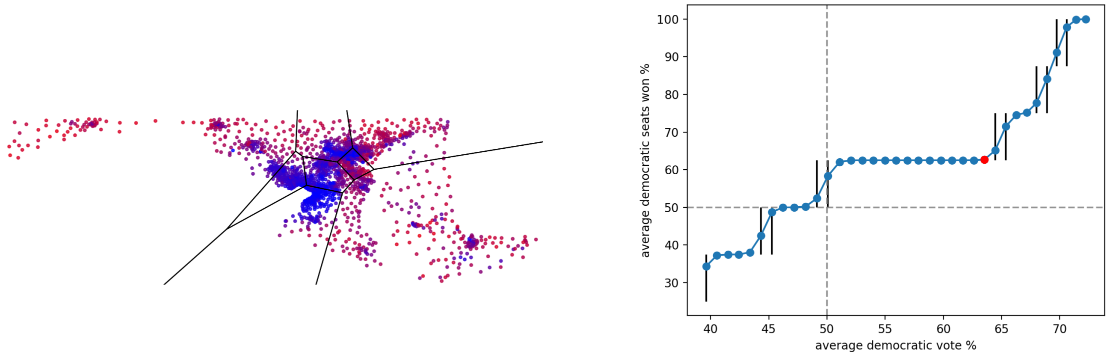

Figure 4: Algorithmic redistricting result and election outcomes for Maryland

data from openprecincts.org which is rigorously compiled and validated [18]. It is important to note that high quality results are not available for many states and elections.

We take these historical precinct results and place all voters within the district where the centroid of the precinct lies. In order to create a robust result, we simulate random elections under this new map. For each precinct we transform the historical result into a normal distribution and independently select a random value. We then average the number of House seats won by each party over a thousand different simulated elections. We compute the expected value of the number of seats each party would have won during that election and the 95% confidence interval of outcomes.

While we measure U.S. House of Representative results, we use party-based voting results from statewide elections, such as the 2016 presidential race. This is necessary since in many states there are districts where candidates ran unopposed. Since those districts are now drawn differently, our result must capture how other parties might have performed in such precincts. Our goal is to build a baseline party breakdown, and we do not account for individual candidate appeal.

This simulation gives us a view into what the results of a specific election might have been under our algorithmic map. While interesting, historical results are only a single data point. Therefore we also compute results for hypothetical elections with different popular votes. To do this, we add a delta increase in votes for one party across all districts in the state. Election outcomes across precincts are correlated: if the democratic party wins big they will likely outperform across many districts [20]. This technique is similar to those used by planscore.org and election forecaster Nate Silver [16, 20]. This modification maintains a statewide voter distribution similar to historical results, while simulating what could happen at different statewide popular votes. Importantly, it provides a picture of what an election result would be under a close election even if that has not happened historically. Importantly, we do not attempt to include complicated voter models as our goal is *not* to predict elections rather to portray a hypothetical results given the 2016 election voter distribution.

We graph these results both for our algorithmic maps and the current district plan for Virginia in Figure 2. Results for additional states can be found in Appendix A. The error bars show the 95% confidence interval of electoral outcomes. Since the number of seats won is always an integer value, adjacent confidence ranges are often the same. These bounds show the number of highly competitive districts for that election, which is sometimes zero.

5 STATE RESULTS

The election outcome diagrams shown in Figures 2, 3 and 4 show how each party would likely perform over a range of potential outcomes. For each simulated election we measure the total vote assigned to each party across all districts. We call this the popular vote. For each popular vote result we average the number of seats won by each party over all simulated elections.

It may be tempting to assume that the ideal fair result should be a linear relation between voter split and expected outcome. However, this result is only indicative of a proportional election system. Under the United States' "first pass the post" system, one should expect a nonlinear relationship with a sharp difference in outcome after receiving a majority of the votes. This is because the probability of winning a district is not independent from other districts' outcome [20]. As soon as a party is likely to win one competitive district, it is likely to win a number of districts and outperform its popular vote.

However, there is little justification for winning a majority of districts given a minority of the popular vote. The importance of this measure is argued and formalized by the legal scholars Grofman and King [15]. There are two important points to consider in these charts: the results given an equal 50/50 percent split between voters, and the percentage of overall votes needed before one party is expected to win a majority of districts. We mark these as dotted lines in our election outcome graphs. See Figure 2. This breaks election outcomes into four quadrants. Results in the upper left or lower right quadrants indicate hypothetical elections where a party wins a majority of the seats with a minority of the vote.

We see that the existing Virginia district plan exhibits significant bias toward Republican voters as Democrats need roughly 55% of the vote before they are expected to win a majority of the seats. Additionally, in any reasonably close election Republicans should be expected to win a majority of the districts.

It is important to note that this current map was redrawn by a court in 2016 [17]. The purposed algorithmic map also appears to show a small bias towards Republicans, but it is significantly less pronounced.

5.1 Competitiveness

Somewhat surprisingly, we find that many real-world district maps resulting from balanced power diagrams are highly competitive. In Virginia, using voting results from the 2016 election four out of the eleven districts would have had outcomes where each party would be within 5% of each other. Many non-partisan policy experts and redistricting committees find competitiveness to be an important factor [14]. For some states, such as Maryland (Figure 4), there exists a plateau where no districts are competitive within a range of popular votes. However, algorithms which intentionally build districts to optimize competitiveness necessarily require a prior expectation of how voters will vote. This raises ethical issues that are sidestepped by unbiased methods which only consider voter location.

5.2 State Outcomes

We present a case study on three states of differing density demographics: Virginia (Figures 1 & 2), Michigan (Figure 3) and Maryland (Figure 4). While no state appears to heavily favor a single party, they do seem to show evidence of some geographic advantage. Michigan shows an example of a state which appears to have a small geographic advantage for the Republican party. However, Maryland shows a clear trend towards the more urban Democratic party. Both results continued to hold for their respective party when redistricted for the more numerous state legislative districts.

While the outcomes differ, this is perhaps not surprising given each state has different population densities. Maryland is highly dominated both in population and area by Baltimore [21]. In Michigan, however, the largest city, Detroit, makes up a much smaller fraction of the population and geographic area [21]. A question then arises: what are the factors of population density which advantage urban or rural voters? We attempt to answer this by generating and evaluating synthetic voter distributions.

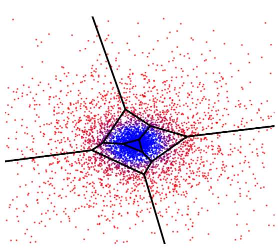

Figure 5: An example district plan for a simulated city. Each dot represents a precinct of 100 voters and is colored according to combined vote.

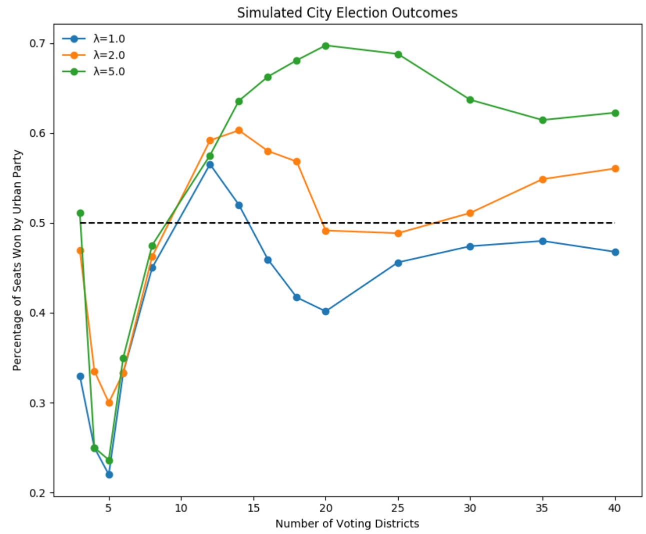

Figure 6: Average election outcomes with a 50/50 popular vote split given the number of districts drawn. The parameter λ controls the population distribution with a higher λ representing a denser population.

6 WHEN CITES LOSE

6.1 Simulations

Since in the United States political affiliation is highly correlated to urban density [19], we focus our attention to outcomes around that of a hypothetical city. To do this, we fix a model and examine population density's effect on the outcome of election results. We use a power law exponential distribution to model population density's falloff from an urban center [5]. The density of this distribution is determined by a parameter λ with

$$p(x) = \lambda e^{-\lambda x}$$

This distribution is used to place voters' location *independently of how they will vote.*

We also need a second distribution determining the probability that a voter will vote for each party as a function of their distance from the city center. Once voter locations are fixed, we select urban party voters without replacement proportionally to $\frac{1}{1+||x||^{\alpha}}$ with parameter α and distance from the city center $||x||$. We set voting behavior in this manner to ensure an exactly even split between the two parties. The last parameter k is the number of districts drawn. We show an example of such a synthetic district map in Figure 5.

6.2 Results

Given our model for cities, we graph the results in Figure 6. Values for $\alpha > 1$ had little impact on the results, therefore we present results in Figure 6 for a fixed α with $\alpha = 2$. It is important to stress that the outcomes in Figure 6 differ solely on the population density parameter λ which is independent of voter political preference. Additionally, this general trend is repeatable for different types of models. We found similar results with added noise, asymmetric

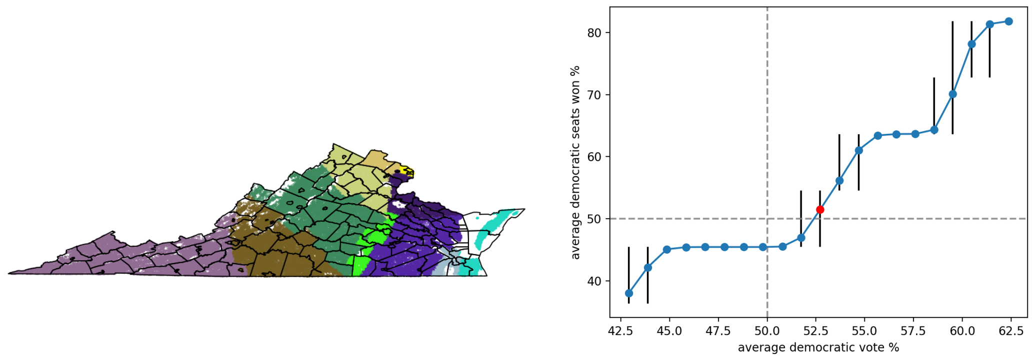

Figure 7: Results for Virginia when incorporating county lines.

population density and voter preference modeled using a power law distribution.

Interestingly, asymmetry appears to help the urban party, where as the amount of noise introduced had varied results, but maintained the same general trend. The urban party consistently performed better by increasing the density parameter λ. Additionally, when k was around five the urban party consistently underperformed.

It is possible that this dip is an artifact of synthetic data. At low values of k the algorithm tends to place a single district entirely within the city. All other districts look like long wedges cutting into the city. As k increases, however, the packed urban centers are balanced by entirely rural districts outside of the city. Districts drawn on real data, either on a city or a state level, more closely resemble synthetic redistricting with larger k. This effect could be reduced by using more realistic asymmetric population distributions which more closely resembles the real world.

6.3 County Lines

Many states require or expect that district plans largely preserve county and municipal lines [14]. Our approach uses census blocks

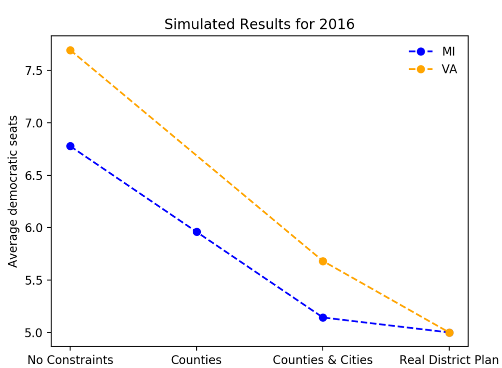

Figure 8: Comparison of simulated results for the 2016 election under various district plans. In Virginia cities are their own county.

as building blocks which are significantly more fine grained than counties [21]. For instance, in Virginia there are 145,045 census blocks but only 133 counties and cities [22]. Because counties are so much larger, using them as indivisible building blocks does not yield high quality district plans. However, incorporating county lines in some manner could be an important factor for the rural-verse-urban balance. County and city boundaries both shape and are shaped by the populations which they contain.

One technique to measure such an effect is to penalize districts which split counties. Our approach maintains a similar power diagram algorithm, but changes the metric space by introducing a penalty for districts which cross an administrative line. To do this we modify the metric space by artificially increasing the distance between census blocks separated by a county line. This means there is more room for power diagram borders to fall between counties and are less likely to split them. Intuitively, this measure is equivalent to adding a "wall" to county boundaries. A line which crosses any wall is longer as it has to travel up and down each side. This technique produces district plans that are affected by administrative lines, but are likely non-ideal for real world use. Due to the non-euclidian metric space, edge cases such as cites wholly contained inside counties can result in some noncontiguous districts. These district plans do, however, provide a window into how such lines could change election outcomes. The number of counties which are split is reduced by about 10-15% when the penalty is introduced.

We find that in Virginia, incorporating county boundaries results in simulated elections that significantly advantage rural voters. See Figure 7 and 8. However, Michigan does not exhibited as strong an effect when only accounting for counties. Importantly, in Virginia city boundaries are themselves counties. In Michigan, however, counties are independent of cities and drawn to be relatively even in size —not population— and are all rectangular and evenly shaped. Because of this, county boundaries for Michigan do not include city municipalities. After including census designated Urban Areas, we see a similar political advantage for rural areas. See Figure 10 and 8.

These results suggest that incorporating existing administrative lines into district plans could introduce unintended political bias.

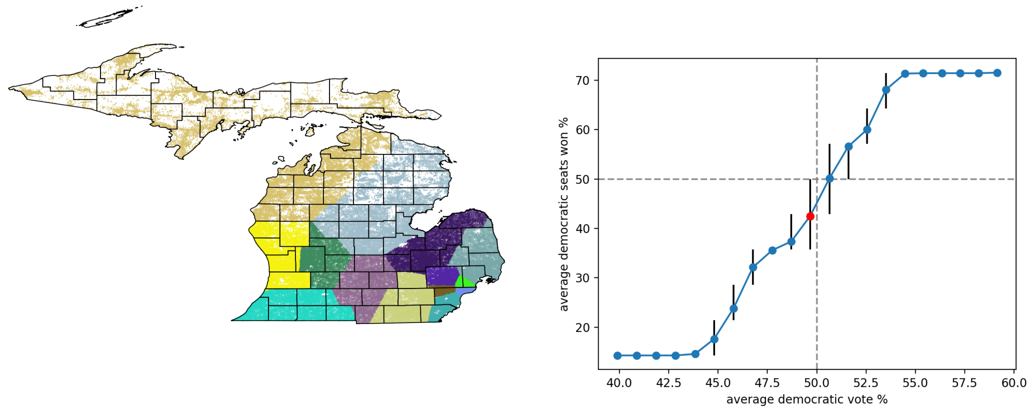

Figure 9: Results for Michigan when incorporating county lines.

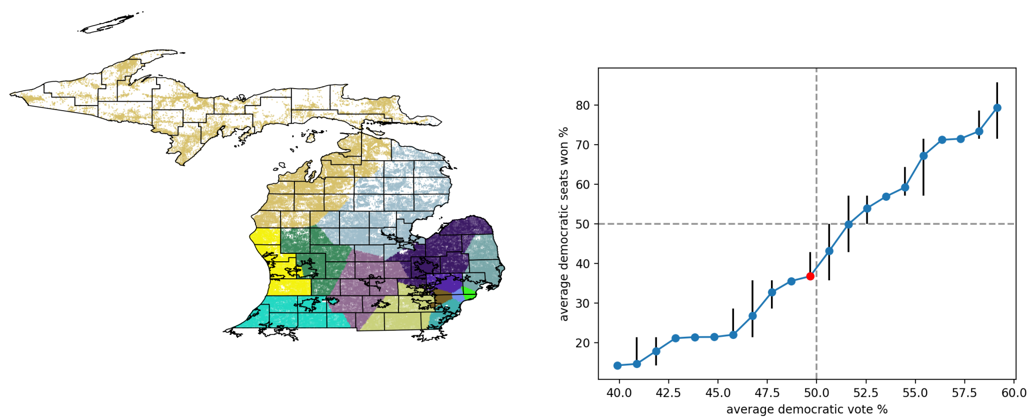

Figure 10: Results for Michigan when incorporating the boundaries of urban areas as well as counties.

Even if city and county lines are not politically motivated, using them could result in politically biased districts.

7 DISCUSSION

Likely one of the reasons our maps do not heavily favor rural voters is that they tend to split dense urban areas across multiple districts. While some districts do lie at the heart of large cities, our approach tends to generate competitive districts that cut across urban, suburban and rural areas. This is contrary to the expectation that compact district plans will place cities inside of a single district or group of districts, thereby packing urban voters and waisting their votes.

In previous work by Chen and Rodden [4] exploring the electoral outcomes of algorithmically generated districts, they find split cities to be rare in their model. Their work, and related results using MCMC, builds districts by randomly combining precincts [4, 10, 12]. Chen and Rodden [4], for instance, argue their maps are compact since they only combine precincts that are nearest to each other. However, we suspect this model *assumption* is highly sensitive to local features. Since urban areas are dense, urban precincts will always be combined with other urban precincts. This process likely results in first placing cities into their own district and then building rural districts from what is left over. While their notion of compactness makes sense from the perspective of a single district, there is no optimization which balances compactness across all districts. This de facto creates a rural-versus-urban divide.

Our approach, however, iteratively reduces voter dispersion across all districts balancing compactness from both local and global features. This means that both urban *and* rural districts are equally optimized for compactness. It is this objective to maximize global compactness that makes our maps significantly different from human drawn maps that only appear compact.

Additionally, our maps may look unnatural or overly simplistic to those familiar with existing district plans. We stress that we are not advocating that our maps should be implemented as is, rather that they show electoral trends inherent to maximizing compactness. There are many additional factors district plans need to achieve such as maintaining majority-minority districts and communities of interest.

8 CONCLUSION

We quantitatively measure the geographic rural-verse-urban balance using an unbiased algorithm to generate compact district plans. Additionally, we provide a nuanced account for the effects of population density. Our results suggest that contrary to assumption, population density may actually advantage the urban party in certain cases.

We expect that these results differ from previous studies for two reason: we ignore administrative lines and our algorithm tends to split cities across multiple districts. When we modify our approach to respect administrative lines, we see an advantage for the rural party. Moreover, the results of our algorithm suggest that splitting cities across multiple districts is actually a result of fully maximizing compactness across all districts. This empirical definition of compactness surprisingly runs counter to many assumptions rooted in a more vague notion of compactness.

We hope to counter the assumption that compactness on its own inherently favors rural voters. We suggest instead that it is attempts to respect municipal and county lines which introduce bias.

ACKNOWLEDGMENTS

This work was supported by National Science Foundation grant CCF-1841954.

REFERENCES

[1] Robert Barnes. Supreme Court to hear cases on partisan gerrymandering. *The Washington Post*, 2019.

[2] Mira Bernstein and Moon Duchin. A formula goes to court: partisan gerrymandering and the efficiency gap. *Notices of the AMS*, 64(9):1020–1024, 2017.

[3] Allan Borodin, Omer Lev, Nisarg Shah, and Tyrone Strangway. Big city vs. the great outdoors: Voter distribution and how it affects gerrymandering. In *IJCAI*, pages 98–104, 2018.

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

[5] Colin Clark. Urban population densities. *Journal of the Royal Statistical Society. Series A (General)*, 114(4):490–496, 1951.

[6] Vincent Cohen-Addad, Philip N. Klein, and Neal E. Young. Balanced centroidal power diagrams for redistricting. In *SIGSPATIAL/GIS*, 2018.

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

[8] Moon Duchin. Geometry versus gerrymandering. *Scientific American*, 2018.

[9] Moon Duchin. Gerrymandering metrics: How to measure? what's the baseline? *arXiv preprint arXiv:1801.02064*, 2018.

[10] Moon Duchin, Taissa Gladkova, Eugene Henninger-Voss, Ben Klingensmith, Heather Newman, and Hannah Wheelen. Locating the representational baseline: Republicans in massachusetts. *arXiv preprint arXiv:1810.09051*, 2018.

[11] Jordan Ellenberg. How computers turned gerrymandering into a science. *The New York Times*, 2017.

[12] Benjamin Fifield, Michael Higgins, Kosuke Imai, and Alexander Tarr. A new automated redistricting simulator using markov chain monte carlo. *Work. Pap., Princeton Univ., Princeton, NJ*, 2015.

[13] Roland G Fryer Jr and Richard Holden. Measuring the compactness of political districting plans. *The Journal of Law and Economics*, 54(3):493–535, 2011.

[14] Justin Levi. All about redistricting: Professor Justin Levi's guide to drawing the electoral lines, 2012.

[15] Michael D McDonald and Robin E Best. Unfair partisan gerrymanders in politics and law: A diagnostic applied to six cases. *Election Law Journal*, 14(4):312–330, 2015.

[16] Michal Migurski, Eric McGhee, Simon Jackman, Nicholas Stephanopoulos, and Ruth Greenwood.

[17] Associated Press. Virginia: Court again redraws voting districts. *https://www.nytimes.com/2016/01/08/us/virginia-court-again-redraws-voting-districts.html*, Jan 2016.

[18] Princeton Gerrymandering Project. Openprecincts. openprecincts.org.

[19] Jonathan Rodden. *Why cities lose: the deep roots of the urban-rural political divide.* Basic Books, 2019.

[20] Nate Silver. A user's guide to fivethirtyeight's 2016 general election forecast, Jun 2016.

[21] U.S. Census Bureau. 2010 census. U.S. Department of Commerce, February 2011.

[22] U.S. Census Bureau. Tiger/line shapefile, 2016, series information for the current county subdivision state-based shapefile. U.S. Department of Commerce, 2019.

[23] Michael Wines. What is gerrymandering? What if the supreme court bans it? *The New York Times*, Mar 2019.

A ADDITIONAL STATES

We present our simulated election results for all states with at least two districts and sufficiently high quality precinct level election data (as of this time) from openprecincts.org [18]. All results use the 2016 presidential election as a baseline. Error bars show the 95% confidence interval of all outcomes. There are slightly different ranges of popular votes for some states as historical outcomes differ. Points marked in red are the average result for the 2016 election with no change in the popular vote.

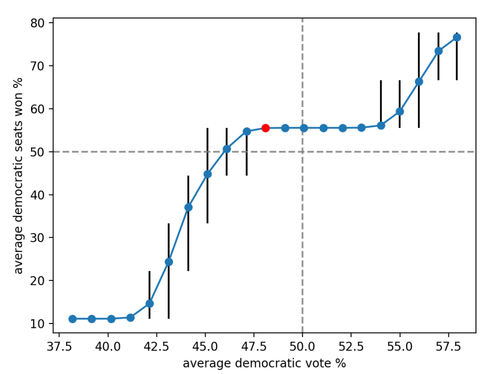

Figure 11: Arizona

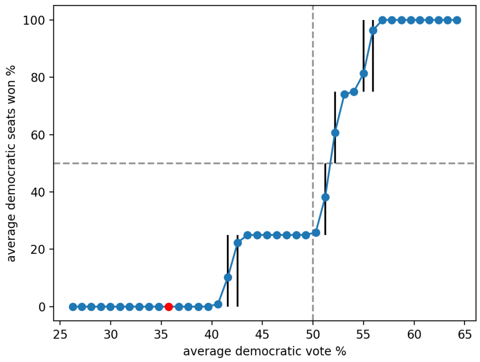

Figure 12: Arkansas

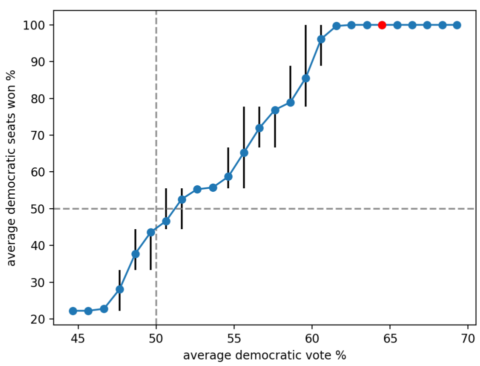

Figure 13: Massachusetts

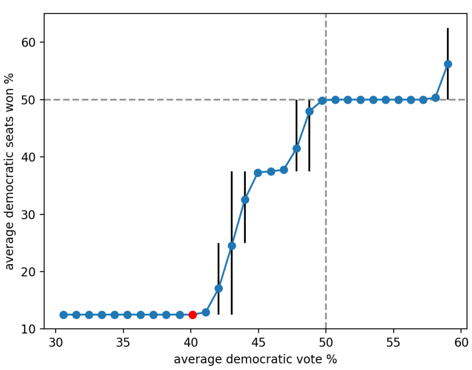

Figure 14: Missouri

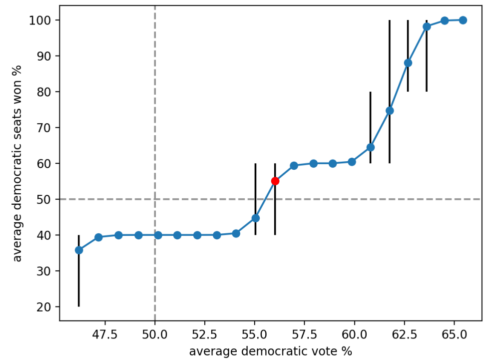

Figure 15: Oregon

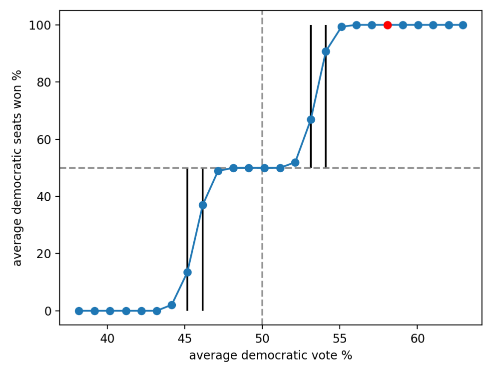

Figure 16: Rhode Island

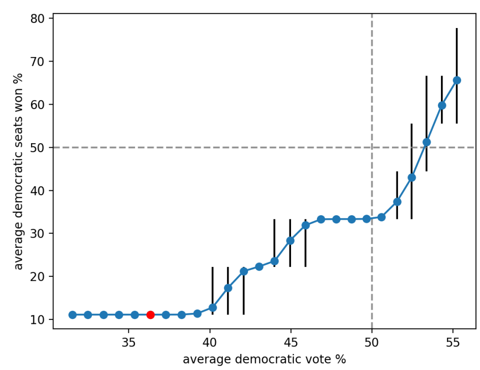

Figure 17: Tennessee

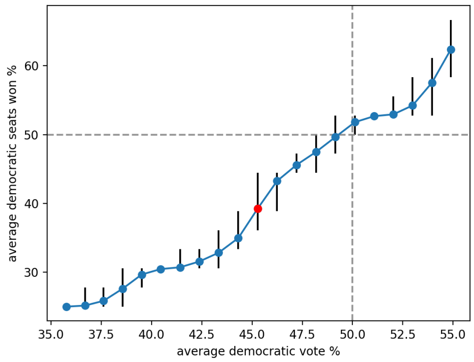

Figure 18: Texas

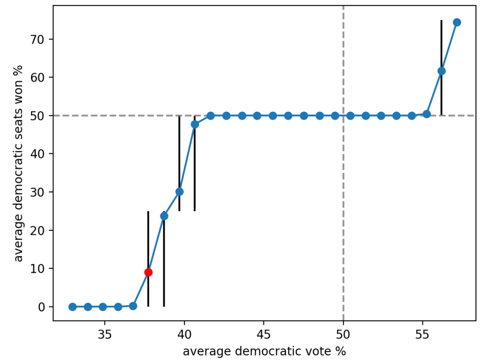

Figure 19: Utah