GerryProjects

Daryl DeFord

January 19, 2021

This document contains some suggestions for projects and open questions about computational redistricting. This is not intended to be a comprehensive resource to all of the interesting questions in this field but rather some starting points that I think are particularly appealing or useful. Many of these descriptions grew out of an email that I wrote to a group of senior math majors who were interested in tackling a gerrymandering related problem for their senior theses, but as you will see the problems stretch from things that can be tackled with little mathematical/computational background to research problems suitable for graduate students.

Most of the ideas below are presented as sketches, rather than completely formulated and explicit conjectures or theorem statements. If one of them seems interesting to you, please feel free to reach out if you would like more information and detail. I can also hopefully connect you to other people who have supervised projects on similar topics, as (older versions of) this list has been used as a starting point for several project-based, semester-long courses around the country.

Background Material: As a fairly new field, there is not yet a ton of material summarizing the current state of computational redistricting as a research enterprise. When the book "Political Geometry" (edited by Moon Duchin) appears later this year, it will hopefully provide a comprehensive source but until then most of the relevant readings are in the form of research papers, some of which are linked below. Additionally, as this is an interdisciplinary research area, really digging into this material requires engaging with literature from political science and other social science perspectives as well as mathematics, statistics, and computer science.

For background about discrete MCMC in this setting, I have prepared a set of notes[1] here along with several interactive software widgets exploring the ideas. In addition to the package documentation, for help getting started with the GerryChain package for computational work I have prepared these templates and this boot-camp style worksheets. Finally, I taught an IAP course on this topic at MIT in 2019 and the materials for that class can be found here.

1 Local Baselines

While much of the initial work and interest in computational redistricting was motivated by the adversarial court setting (i.e. this map is a gerrymander because ... or this new court-ordered map remedies the initial harms because ... Excellent examples of these sorts of analysis are presented in [5, 16, 17]), from my perspective much of the important work now relates to questions about understanding the space of potentially valid plans and the potential impacts of reform language on eventually enacted plans. One way to think about these concepts is in terms of baselines and exploitability.

Understanding or creating a reasonable baseline of metrics for a given state requires first understanding and operationalizing the laws that govern redistricting. The operationalization step is crucially important, since legislation is rarely written with mathematical specificity, and frequently requires doing some historical digging to determine how the stated guidelines have been applied in practice over the decades. Once a set of modeling choices has been made, an ensemble of plans can be constructed that reflect potentially permissible districting plans under those rules. A variety of partisan, demographic, geographic, and other metrics can then be evaluated on those plans to determine a starting point for understanding the expected partisan (demographic, geographic, etc.) behavior of a plan drawn only according to the stated rules.

[1] beware typos!

I focus on the key role of operationalization above both because it is a critical step to performing analyses that are relevant to the real-world concerns but also to highlight the inherent flexibility within the problem formulation[2] and the relationship of language choices to the quantitative analysis. This is becoming increasingly relevant as states look to reform their constitutional language governing redistricting, sometimes forced by citizen's initiative efforts as we saw in 2018.

These ensemble tools give us a way to measure the potential impacts, and more importantly, the unintended consequences of specific choices of language or line-drawing priorities. We saw this play out several times with the 2018 reform bills and this will continue to be an important perspective throughout the current cycle (some examples here [12]). The question of analyzing specific language also raises questions of optimization and exploitability, particularly as states start to adopt more quantitative language, as in Missouri (since repealed) and Utah, and potentially begin to run afoul of Goodhart's Law. In these cases, it is important to understand whether achieving the stated goals is possible as well as whether achieving those goals actually constrains the line-drawers from creating unfair maps in the intended fashion [9]. Modern computational tools have made it possible to start to address these issues but there is still much that we don't understand about optimizing over the relevant state space.

Thus, the first set of projects and questions that I suggest is focused on tackling these state or area specific questions. The specific questions of interest vary between states and the choices of modeling procedures are also likely to lead to differences, so there is a lot of value in multiple groups attempting these analyses. Beyond the specific examples presented here, many of the problems presented in the other sections below could also benefit from simply being evaluated locally. For example, the questions about municipality splitting could be just as easily asked about splitting school districts or precincts in local elections. Some previous examples of this sort of work include [3, 6, 10] as well as these reports focused on city council races: mggg.org/{chicago/lowell/SantaClara}.

2 Project Blurbs

Most of the problems described below are raised in the context of using MCMC sampling to generate ensembles of districting plans. The 50,000 foot overview of the procedure is to create a dual graph from geographic units with relevant data (like voting precincts or census blocks) which reduces the problem to studying graph partitions. Then, the state space must be delineated by operationalizing the relevant state rules and a distribution and Markov chain must be selected. One of the most common choices for the proposal method in the underlying Markov chain is the ReCom algorithm introduced in [11], which creates a new partition by redrawing a pair of adjacent districts at every step. I've tried to separate the problems below into categories but unsurprisingly many of the problems are related to all steps of this process.

2.1 Data and Dual Graphs

The underlying data for redistricting problems starts as a dual graph to a planar partition, derived from geographic units. These graphs tend to look fairly structured, with many lattice-like components, but contain several features not common to combinatorial graphs or social networks. I have created .json files and visualizations of all of the census unit dual graphs with demographic information here. This first set of problems concerns the properties of the underlying data.

- **Graph and Network Theory** From the perspective of graph theory one of the basic questions is characterizing the set of graphs that arise as dual graphs to plane partitions. These graphs need not be planar[3] but they are mostly "close" to planar. There are plenty of formal definitions of these types of properties but proving exactly which properties are common to the graph we observe would be valuable. Additionally, even without a formal description it would be useful to consider generative models for building "similar-looking" graphs as null models. Finally, given that very little is known about these

[2] My favorite example of this occurred in a court case in the previous cycle where there were three separate groups that presented ensembles and each of them chose different functions to measure each of the measured variables, including population balance, county splitting, and compactness

[3] They can come from disconnected units, have weird behavior near water, or fail to be simply connected, etc.

graphs as a class there are a lot of interesting starting points, computing the various standard graph properties - colorings, matchings, homomorphisms, path lengths, centralities, etc.

- **Graph Homomorhpisms** This is mentioned in a couple of different places in this document but deserves its own special bullet point. In order to measure how close to grid–like a particular dual graph is, you might define a measure capturing the number of homomorhpisms to/from various regular lattices/grid graphs into/out-of the graph. Additionally, it may be possible to bound the number of matchings in these graphs using pullbacks of these same homomorphisms.

- **Counting Partitions** A seemingly straightforward combinatorial question that turns out to be quite difficult is enumerating the possible contiguous districtings of a graph that satisfy certain additional properties (for arbitrary graphs, we don't even know how to just compute the number of contiguous partitions efficiently). Looking at specific families of graphs (grids, wheels, etc.) would be a good starting point here for the exact case while trying to prove some upper/lower bounds on the number of admissible partitions is the more likely course for arbitrary graphs. As a concrete example, the number of partitions of a 10×10 grid into 10 contiguous pieces is currently unknown. Some very interesting recent work on efficient algorithms for this problem is used in [15].

- **Generating partitions** Given a graph with population weights and political/geographic node labels can you construct reasonable samples from the space of all permissible plans? Given that we can't enumerate these (see above) it isn't reasonable to expect to be able to choose a particular distribution, thus you have a lot of flexibility in making use of any sort of method for constructing the partitions. A list of example methods I tested out several years ago is here: https://v2.overleaf.com/read/zpmyzqmpvmnx but a) these haven't been analyzed carefully b) there are undoubtedly more methods available and c) these only balance populations and don't consider political boundaries, etc. These methods can also serve as sub-algorithms for bipartitioning as a part of the Markov proposal for ReCom MCMC.

- **Vote Distributions** We have samples of voting data in annual or biannual snapshots but very little understanding of the spatial structure of these voter distributions. There are many open questions here in terms of how the distributions evolve. Additionally, understanding how the vote distributions can change the outcome under various districting plans is under explored. This is really a question of robustness - how stable are the voting outcomes under perturbations of the vote data, given that the plan is fixed? Additionally, as with the local work projects in Section 1 there are descriptive questions that haven't been answered about the properties of these distributions across the states and years.

- **Vote Models** In addition to the perturbation question above there is also the question of defining simple models for generating reasonable voting data sets. Although this seems like a challenging task, the increasing urban/rural polarization in the country makes this tractable in some states. Expectation maximization on 2d gaussians seems to work quite well for states with well separated population centers but could more complex discrete optimization also work?

- **Proration and Roundoff** Frequently, the data used in the redistricting process is heterogeneous in the sense that it is most naturally defined on mis-matched geographic units. For example, the census reports population at the block level but the states usually report votes at the precinct or county level. In order to combine the information at a single level of resolution we frequently aggregate or interpolate the results between units. However, no one has studied what types of distortions might get encoded by this process and whether or not they should be expected to be systematic or random.

- **Urban vs. Rural** One of the fundamental topological questions is how much of the full information is contained in the dual graph. From a geometric perspective it is relatively easy to determine which units are urban vs. rural (although implementing this in practice would be a nice machine learning project) but using only the topology this becomes much more complex. Given just the graph can you determine the same classification? If not, what additional information is necessary to distinguish these regions?

2.2 Markov Chains

This next set of problems is more directly concerned with the properties of the algorithms for generating samples from distributions over partitions. The open questions lists at the end of these papers are also good places to start [11, 18, 19].

- **Tree splitting** One method currently in use for creating new partitions is drawing a spanning tree for the dual graph and deleting $k-1$ edges to get a k partition [4, 11]. Since we want the partitions to be population balanced the sets of edges that can be cut is much smaller than the $\binom{n-1}{k-1}$ naive options. This procedures leads to both computational and theoretical questions, including:
  - What is the best algorithm for finding a cuttable edge / all cuttable edges.
  - What are the expected properties of uniform sampling from all permissible cuts on a fixed tree (how different are they?)
  - Can we bound the # possible trees given a fixed ε (both graphs and and geographic)
  - How different are the distributions between deleting k edges at once vs. deleting a single edge and redrawing a tree.
  - What proportion of trees are splittable?
  - How long do you have to wait to find a splittable tree when cycle–basis walking?
  - How bad is the "best" cut, even if not permissible?
  - What is the distribution of best cuts over trees?
- **Municipality Splitting** Many states require the mapmaker to minimize splits of municipalities. However, this is a difficult concept to operationalize, since "as few as possible" is not mathematically well defined. Previous attempts have tried to simply bound the number of splits but this does not take in to account population differences (some must be split for population reasons, some states prefer to require that small counties by kept together, etc.) or the question of how many splits should be allowed in a county that touches more than one district. A couple of recent papers [2, 6] propose algorithmic methods for sampling these more easily but there is plenty of room for new and creative approaches.
- **Sample Hardness** Some recent papers have looked at this problem from a complexity theory perspective to prove hardness results about generating samples from distributions over connected graph partitions [1, 7, 18, 19]. One key question is relating these results to the types of data (and preferred distributions) actually encountered in empirical redistricting work. Additionally, some of these results raise concerns about the stability of sampling methods to adversarial actors with control over various aspects of the initial data or sampling method.
- **Spectral ReCom** Instead of using spanning trees to do the bipartioning in ReCom, we might instead use a variant of spectral clustering, partitioning using the Fiedler verctor. An example of using this in practice is included in the GerryChain Templates linked above and the apparent convergence of that example leads to some natural theoretical questions. Can we prove that this always converges? How does the final solution compare to the continuous 2d soap bubble problem in the limit? How does this algorithm behave under refinement?
- **Districting tradeoffs** There are many types of constraints placed on permissible maps but how they interact within the sampling methodology is mostly unknown. A starting point here might be to see what the extreme plans look like and use that to investigate the relations, establishing a Pareto frontier for the chosen metrics. For example, constructing partitions that are maximally population balanced or compact or partisanly unfair must constrain the possibilities for the other values to vary. A recent example of this occurred in PA where one party managed to construct a particularly compact partisan map but the partisan advantage was very unstable - is this always the case? This is interesting both in terms of binary constraints placed on the state space, as well as metrics used in the energy function controlling the distribution.

- **New Proposals**: While there are a couple of commonly used proposal distributions for these problems, there is still plenty of room for new ideas. This is another wild west type problem where almost any proposal type is fair game for analysis. In particular generating interesting reversible proposals or constructing methods for converting current proposals to reversible ones would allow for the application of some recently proved theorems for redistricting analysis. A related problem is determining when the space that you are walking in is disconnected by a particular choice of constraints. Finally, there has been some recent work on finding the nearest reversible chain for irreversible transition matrices. Maybe this could be used to salvage the previously discarded ideas.

- **Functional Convergence** One real surprise in our analysis has been that we frequently see what appears to be convergence behavior in distributions of functions long before it is possible for the Markov chain to have actually mixed. For example, there are (probably!) quintillions of possible partitions of the 10x10 grid but the distributions of number of seats won by a given part seem to be reasonably consistent after tens of thousands of steps, using a very slow mixing walk. Characterizing the properties of functions or walks that have this behavior would be a great problem.

- **Aggregate measures** When performing Gibbs energy versions of MCMC it is common to constrain the behavior of the entire plan by some measures like population balance or compactness. However, it is also true that most of the laws enforce these conditions on individual districts, not on the plan as a whole. The impact of this aggregation is not well understood and should definitely be looked at particularly in the overall behavior of chains as the aggregation changes from l0 to l1 to l2.

- **Benchmarking.** There are several different methods in the literature for generating districting plans by computer but currently we don't have a great way to compare them in terms of their abilities to explore the entire space of permissible plans. There are three main avenues here. First, it would be useful to have a set of methods for computing the distance between two districting plans - people have considered several but there aren't any systematic studies yet. Relatedly, for methods that start with an initial plan and iteratively perturb it, being able to compute a distance measure that quantifies how far the partition has moved after k steps would be useful. Finally, viewing the collection of plans generated by a given method as an ensemble or distribution over the full space of plans, generating a method for comparing the distributions would be useful. There are some standard tests and metrics for this but again, no systematic approach and plenty of room for experimentation.

2.3 Metrics

The final category of problems concerns the measurements that we apply to districting plans to understand their properties. One of the fascinating parts of this application domain is that there are many types of metrics that are relevant in different situations, from partisan imbalance to geography. Many of the commonly used approaches were defined in the political science literature long before the recent resurgence of interest in ensemble methods and there is plenty of room for creative new ways to evaluate (ensembles of) districting plans.

- **Generalized matchings** Some states require that senate districts be formed out of 2 or 3 contiguous house districts. An example analysis of this sort is presented in [3]. This offers possibilities for gerrymandering at both the level of defining the house districts as well as in pairing together the smaller districts to make large ones. Enumerating all such partitions is a hard computational task in general but for specific families of examples there are undoubtedly clever combinatorial methods that could be used. Once the partitions are determined there is also the related problem of understanding how the vote distribution can change the outcome. This concept is also related to one of the game-theoretic approaches discussed below.

- **CVAP Malapportionment** At the end of the previous census cycle there was a push by some states to use citizenship data to draw their district boundaries, with respect to population balancing. Very little is known about the likely impacts of such proposals both in terms of consequences for malapportionment, as well as VRA representation issues.

- **Competitiveness** One question that is likely to come up several times in the current crop of initiatives is the question of whether enforcing "competitive" districts is a good requirement to include. We would like to study whether competitive districts are actually better for overall representation. As an example, Arizona has a mandate for competitive districts in their constitution as well as a non-partisan drawing committee - does this lead them to have better outcomes overall? are there unintended consequences? In states without such a requirement, is having more competitive districts correlated with better representation? This fits in to a broader set of questions about how different requirements on districts interact and the extent to which you can optimize for a particular constraint while still satisfying the others. This question has been addressed in a couple of recent papers [6, 12] but there is still plenty of work to be done, particularly in understanding the tradeoffs of competitiveness with other measures of partisan balance.

- **Compactness measures.** Many state constitutions have rules requiring that districts be made as compact as possible but the definition of compactness is rarely specified. This means it is kind of the wild west in terms of coming up with interesting measures of how to tell when a district is "circular enough". That is, you could generate some really imaginative measures using pretty much any mathematical tool and compare how they perform on various plans. A comprehensive set of examples showing where the current measure perform well and poorly (with respect to their intuitive definitions/purposes) would also be useful. Beyond this question of measuring values on specific districts there is also the related question of how you might compute a measure for an entire districting plan at once. People have used things like the worst score or some kind of weighted average of the scores in the plan but this is definitely an issue crying out for some thoughtful analysis, both for new measures as well as the old ones. This question is also interesting since it applies for both continuous and discrete measures.

  Recently, we have been advocating for multiscale approaches that incorporate a geometric perspective [13, 20] and there are plenty of open questions surrounding these ideas in both pure and applied mathematics. A simple version of this would be to use a pointwise 'distance from boundary' as a multiscale measure where you measure the proportion of the district that is within x units distance from the boundary. This has a nice discrete and continuous formulation and many variants, like measuring out to the boundary from centroid or measuring the average proportion of boundary reachable from points inside.

- **Skinny Compactness** Many of the geometric measures used to evaluate the compactness of legislative districts focus on the large scale details but aren't as sensitive to the internal structure of the district. One interesting project would be constructing measures that identify "skinny" regions of districting plans both for continuous regions in a plane as well as in the discrete case of graphs representing the individual geographic units. A useful measure should do more than just identify a single place where the district is pinched (which may be necessitated by external factors) but instead try to capture whether or not the skinny region persists for a long distance within the district.

- **Game theory** We usually think of map drawing as being controlled by a single party but there has been some recent research on methods for fair division type analyses of districting and related impossibility theorems. Variations of "I cut you choose" and similar approaches are being explored but there is a lot of room for creativity here. Additionally, the matching problem may fit well into this framework: Let one party draw k*n subdistricts and then allow the other party to merge these into n final districts.

- **Voting theory** Most of the current research has focused on the winner take all districts of the US congressional system but there are many other types of voting systems that still interact with district boundaries. Analyzing how different voting systems perform under different districting plans might be interesting. For example, multi representative districts or transferable vote schemes might influence the type of partisan balances that map makers would like to draw. There have been some preliminary analyses of this sort at the level of city councils (see e.g.: mggg.org/{chicago/lowell/SantaClara}) but there is still plenty of modeling and simulation to be done just to get started.

- **Geometry.** This is similar to the questions about graphs above but using the actual plane partition geometry. A characterization of the types of regions that occur in the basic census and their properties would be very useful. Adding the shape of the boundaries adds a bunch of interesting questions to the graph setting, since in a graph you assume all the edges and nodes are abstract but the full geometry has a great deal of extra structure - coastlines, river boundaries, state boundaries, roads, etc. A great deal is known about operators (and their spectrum) on surfaces but relatively little of this machinery has been leveraged in the context of redistricting. Viewing the population as a measure on the geometry also opens several extra possibilities for analysis.

- **Partisan measures** What are the 'right' or 'useful' notions of fairness or the amount of partisan imbalance. Can we get away from linear transformations of the seats–votes curve? Here is an example of a measure that incorporates geography [8, 14]. Additionally, even for the commonly used and studied measures (mean median, partisan bias, efficiency gap, etc.) there is still plenty we don't know about how they play out in practice [9].

- **Vote distribution** While much of the work to this point has focused on varying the districting lines, one of the key future directions is figuring out how to including variants or uncertainty in the underlying vote data with the ensemble approach. Some interesting related questions include:

  – How sensitive is a given districting plan to perturbations in vote distributions? For example, if you are allowed to transfer at most x% of the population across a single district boundary (or collection of boundaries) how much can you change the electoral outcome? Relatedly, how well correlated geographically are the proportions of voters in each party? Given a districting plan can you come up with a method for determining the smallest amount of voter preference change that would need to occur in order to achieve a particular partisan outcome in results? Are there currently districts in the US that are particularly sensitive to this or that are particularly unfair?

  – What do good null models for voting data look like? Particularly if we want to incorporate the underlying graph structure? Can low dimensional representations of voting patterns be used to generate useful counterfactuals? What balance between sampling and interpolation is most useful for combining multiple sets of vote data in the same setting?

  – Some people have proposed multi representative districts with some sort of transferable vote mechanism as a solution to gerrymandering. This is a little tough to analyze given the current voting data that exists, is there a way to formalize this given the current data? One model for this is to take a given districting plan ans form "superdistricts" by merging adjacent districts together (several states like Alaska and Iowa form their state senate districts from state house districts in this fashion) and then electing multiple representatives for the superdistricts equal to the number of districts that were merged together. Can you prove that this is always at least as far as voting in the original districts? Alternatively, can you find or characterize examples where this method is particularly unfair?

2.4 Example Project

Here is an example of a project that a student used as the starting point for their senior thesis in 2019, sketching out some ways to generalize an initial problem. I definitely wouldn't set it up exactly the same way now but it does give a sense of how easy it is to branch out from a simple looking starting point.

- To begin with let's assume that our districting plan is an n x n grid with n districts, where each district has sizes in [n-k,n+k] for some fixed k. A common walk on this space is to randomly select a pair of adjacent squares that belong to different partitions and then change the assignment of one of them (again chosen at random) so that they match. The move is accepted if the districts remain contiguous and the sizes remain in [n-k,n+k]. We can measure the compactness of the districting by summing the perimeters of each district or by counting the total number of edges that lie between squares with different assignments.

- MCMC properties. In this setting we have two measurables: population imbalance and cut-edge compactness. You could compare a few types of MCMC walks for reasonably small n
    1. Sample uniformly
    2. Use Metropolis-Hastings to prioritize compact or noncompact districts
    3. Use Metropolis-Hastings to prioritize population (im)balanced districts
- For each of these you can compute the distribution of the two statistics as the chain proceeds. An interesting result might start by proving positive correlation between the two and then move on to characterizing the distribution of one measure in terms of the steady state of a MH walk that (de)prioritizes the other.

  Possible extensions:

    - Heterogeneous population. First, how do the results above vary as you vary k? How does this problem change if instead of requiring that the district sizes are within [n-k,n+k] you instead place an integer value on each node and require that the populations are balanced to within some fraction epsilon of the mean? Can you construct examples of weights that make it impossible to have districts? Can you prove or characterize the properties of such weightings?
    - Including voting data. Use fixed assigned votes on the grid or randomly select the votes (flip a weighted coin for each square) and record the number of districts one by each party as a third measure. How does it relate to the other two?
    - Mixing times. For small n you can explicitly construct the transition matrix and compute its eigenvalues. How does the mixing vary with the proposal choice? How does this appear to scale as the grid gets larger?
- Generating partitions. Let's think about either the hierarchical or flood fill methods from here: https://v2.overleaf.com/read/zpmyzqmpvmnx in our grid setting. Can you prove bounds on the expected perimeter or population of a district grown with one of these methods on a grid? As a preliminary step you could look at the small n case and actually construct the distribution over partitions. That would be super interesting and useful in its own right.
- This recent paper: https://www.cambridge.org/core/journals/political-analysis/article/new-approach-for-developing-neutral-redistricting-plans/31F8EB3FFB7A8F5B3F7C2171BE016D47 has a description of the "fairness" of a method for generating plans in terms of the behavior of the median of an expected plan on a grid. You could repeat their analysis with these methods and compare the results to the one that they present. For some of these, you should also be able to prove that the method is balanced, in expectation.
- New MCMC proposal. Something that we have had some success with recently is replacing the proposal above with the following procedure: At each step randomly choose an adjacent pair of districts and merge all of their nodes into a single one. Then, use one of these generating methods to split the superdistrict in half to get another valid plan. Can you prove anything about the distribution of measurables of this type of walk for a specific choice of method?
- Min-Cut methods. In order to use a min-cut method one approach is to use random weights on the graph edges. Empirically, it looks like exp(minimum distance -3) seems to works pretty well as a weight but I'm sure a results could be proved showing that an exponential function like this almost always leads to good partitions.
- Build your own. Can you invent your own method that provably creates population balanced districts, whose complements are connected and that assign non-zero probability to each contiguous partition?

References

[1] Hugo A Akitaya, Matthew D Jones, Matias Korman, Christopher Meierfrankenfeld, Michael J Munje, Diane L Souvaine, Michael Thramann, and Csaba D Tóth. Reconfiguration of Connected Graph Partitions. *arXiv preprint arXiv:1902.10765*, 2019.

[2] Eric A. Autry, Daniel Carter, Gregory Herschlag, Zach Hunter, and Jonathan C. Mattingly. Multi-scale merge-split markov chain monte carlo for redistricting. *arxiv:2008.08054*, 2020.

[3] Sophia Caldera, Daryl DeFord, Moon Duchin, Samuel C Gutekunst, and Cara Nix. Mathematics of Nested Districts: The Case of Alaska. *Statistics and Public Policy*, 7:39–51, 2020.

[4] Daniel Carter, Gregory Herschlag, Zach Hunter, and Jonathan Mattingly. A Merge-Split Proposal for Reversible Monte Carlo Markov Chain Sampling of Redistricting Plans. *arXiv:1911.01503*, 2019.

[5] Maria Chikina, Alan Frieze, and Wesley Pegden. Assessing significance in a Markov chain without mixing. *Proceedings of the National Academy of Sciences*, 114(11):2860–2864, March 2017.

[6] J. Clelland, H. Colgate, D. DeFord, B. Malmskog, and F. Sancier-Barbosa. Colorado in Context: Congressional Redistricting and Competing Fairness Criteria in Colorado. *arXiv:2011.06049*, 2020.

[7] Vincent Cohen-Addad, Philip N Klein, and Dániel Marx. On the computational tractability of a geographic clustering problem arising in redistricting. *arXiv preprint arXiv:2009.00188*, 2020.

[8] D. DeFord, N. Eubank, and J. Rodden. Partisan dislocation: A precinct-level measure of representation and gerrymandering. *Political Analysis*, to appear.

[9] Daryl DeFord, Natasha Dhamankar, Moon Duchin, Varun Gupta, Mackenzie McPike, Gabe Schoenbach, and Ki Wan Sim. Implementing partisan symmetry: Problems and paradoxes. *arxiv:2008.06930*, 2020.

[10] Daryl DeFord and Moon Duchin. Redistricting Reform in Virginia: Districting Criteria in Context. *Virginia Policy Review*, 12(2):120–146, 2019.

[11] Daryl DeFord, Moon Duchin, and Justin Solomon. Recombination: A family of Markov chains for redistricting. *arXiv preprint arXiv:1911.05725*, 2019.

[12] Daryl DeFord, Moon Duchin, and Justin Solomon. A Computational Approach to Measuring Vote Elasticity and Competitiveness. *Statistics and Public Policy*, 7:39–51, 2020.

[13] Daryl DeFord, Hugo Lavenant, Zachary Schutzman, and Justin Solomon. Total Variation Isoperimetric Profiles. *arXiv:1809.07943 [cs, math]*, September 2018.

[14] N. Eubank and J. Rodden. Who is my neighbor? the spatial efficiency of partisanship. *Statistics and Public Policy*, 2020.

[15] Benjamin Fifield, Kosuke Imai, Jun Kawahara, and Christopher T. Kenny. The Essential Role of Empirical Validation in Legislative Redistricting Simulation. *Statistics and Public Policy*, 7:52–68, 2020.

[16] Gregory Herschlag, Han Sung Kang, Justin Luo, Christy Vaughn Graves, Sachet Bangia, Robert Ravier, and Jonathan C. Mattingly. Quantifying Gerrymandering in North Carolina. *arxiv:1801.03783*, 2018.

[17] Gregory Herschlag, Robert Ravier, and Jonathan C. Mattingly. Evaluating Partisan Gerrymandering in Wisconsin. *arXiv:1709.01596 [physics, stat]*, September 2017.

[18] Lorenzo Najt, Daryl Deford, and Justin Solomon. Complexity and geometry of sampling connected graph partitions. *arXiv preprint arXiv:1908.08881*, 2019.

[19] Lorenzo Najt, Daryl DeFord, and Justin Solomon. Empirical sampling of connected graph partitions for redistricting. *arxiv:2012.04564*, 2020.

[20] Paul Zhang, Daryl Deford, and Justin Solomon. Medial axis isoperimetric profiles. *arxiv:2007.01428*, 2020.