

The Essential Role of Empirical Validation in Legislative Redistricting Simulation*

Benjamin Fifield[†] Kosuke Imai[‡] Jun Kawahara[§] Christopher T. Kenny[¶]

Forthcoming in *Statistics and Public Policy*

Abstract

As granular data about elections and voters become available, redistricting simulation methods are playing an increasingly important role when legislatures adopt redistricting plans and courts determine their legality. These simulation methods are designed to yield a representative sample of all redistricting plans that satisfy statutory guidelines and requirements such as contiguity, population parity, and compactness. A proposed redistricting plan can be considered gerrymandered if it constitutes an outlier relative to this sample according to partisan fairness metrics. Despite their growing use, an insufficient effort has been made to empirically validate the accuracy of the simulation methods. We apply a recently developed computational method that can efficiently enumerate all possible redistricting plans and yield an independent sample from this population. We show that this algorithm scales to a state with a couple of hundred geographical units. Finally, we empirically examine how existing simulation methods perform on realistic validation data sets.

Keywords: enumeration, gerrymandering, graph partition, Markov chain Monte Carlo, redistricting, zero-suppressed binary decision diagram

*We thank Steve Schecter, participants of the Quantitative Gerrymandering and Redistricting Conference at Duke University, and two anonymous reviewers for helpful comments. Kawahara acknowledges financial support by JSPS KAKENHI Grant Number JP18K04610. The computations in this paper were run on the FASRC Cannon cluster supported by the FAS Division of Science Research Computing Group at Harvard University. Replication materials and enumeration results are available as Fifield *et al.* (2020b).

[†]Affiliated Researcher, Institute for Quantitative Social Science, Harvard University, Cambridge, MA 02138. Email: benffield@gmail.com, URL: <https://www.benffield.com>

[‡]Professor, Department of Government and Department of Statistics, Institute for Quantitative Social Science, Harvard University, Cambridge MA 02138. Phone: 617-384-6778, Email: Imai@Harvard.Edu, URL: <https://imai.fas.harvard.edu>

[§]Associate Professor, Graduate School of Informatics, Kyoto University. Email: jkawahara@i.kyoto-u.ac.jp, URL: <http://www.lab2.kuis.kyoto-u.ac.jp/jkawahara/index-en.html>

[¶]Ph.D. Student, Department of Government, Harvard University. Email: christopherkenny@fas.harvard.edu URL: <https://www.christophertkenny.com>

1 Introduction

Congressional redistricting, which refers to the practice of redrawing congressional district lines following the constitutionally mandated decennial Census, is of major political consequence in the United States. Redistricting reshapes geographic boundaries and those changes can have substantial impacts on representation and governance in the American political system. As a fundamentally political process, redistricting has also been manipulated to fulfill partisan ends, and recent debates have raised possible reforms to lessen the role of politicians and the influence of political motives in determining the boundaries of these political communities.

Starting in the 1960s, scholars began proposing simulation-based approaches to make the redistricting process more transparent, objective, and unbiased (early proposals include Vickrey, 1961; Weaver and Hess, 1963; Nagel, 1965; Hess *et al.*, 1965). While this research agenda lay dormant for some time, recent advances in computing capability and methodologies, along with the increasing availability of granular data about voters and elections, has led to a resurgence in proposals, implementations, and applications of simulation methods to applied redistricting problems (e.g. Cirincione *et al.*, 2000; McCarty *et al.*, 2009; Altman and McDonald, 2011; Chen and Rodden, 2013; Fifield *et al.*, 2014, 2020a; Mattingly and Vaughn, 2014; Liu *et al.*, 2016; Herschlag *et al.*, 2017; Chikina *et al.*, 2017; Magleby and Mosesson, 2018; Carter *et al.*, 2019; DeFord *et al.*, 2019).

Furthermore, simulation methods for redistricting play an increasingly important role in court cases challenging redistricting plans. In 2019, simulation evidence was introduced and accepted in redistricting cases in North Carolina, Ohio, and Michigan.¹ In the few years prior, simulation methods were presented to courts in North Carolina, and Missouri.² Given these recent court cases challenging redistricting in state and federal courts, simulation methods are expected to become an

¹Declaration of Dr. Jonathan C. Mattingly, *Common Cause v. Lewis* (2019); Testimony of Dr. Jowei Chen, *Common Cause v. Lewis* (2019); Testimony of Dr. Pegden, *Common Cause v. Lewis* (2019); Expert Report of Jonathan Mattingly on the North Carolina State Legislature, *Rucho v. Common Cause* (2019); Expert Report of Jowei Chen, *Rucho v. Common Cause* (2019); Amicus Brief of Mathematicians, Law Professors, and Students in Support of Appellees and Affirmance, *Rucho v. Common Cause* (2019); Brief of Amici Curiae Professors Wesley Pegden, Jonathan Rodden, and Samuel S.-H. Wang in Support of Appellees, *Rucho v. Common Cause* (2019); Intervenor’s Memo, *Ohio A. Philip Randolph Inst. et al. v. Larry Householder* (2019); Expert Report of Jowei Chen, *League of Women Voters of Michigan v. Benson* (2019).

²Expert Report of Jowei Chen, *Raleigh Wake Citizens Assoc v. Wake County Board of Elections* (2016); Expert Report of Jowei Chen, *City of Greensboro v. Guilford County Board of Elections* (2015); Supplemental Report of Jonathan Rodden and Jowei Chen: Assessment of Plaintiffs Redistricting Proposals, *Missouri State Conference of the NAACP v. Ferguson-Florissant School District* (2017).

even more influential source of evidence for legal challenges to redistricting plans across many states after the upcoming decennial Census in 2020.

These simulation methods are designed to yield a representative sample of redistricting plans that satisfy statutory guidelines and requirements such as contiguity, population parity, and compactness.³ Then, a proposed redistricting plan can be considered gerrymandered if it constitutes an outlier relative to this sample according to a partisan fairness measure (see Katz *et al.*, 2019, for a discussion of various measures). Simulation methods are particularly useful because enumeration of all possible redistricting plans in a state is often computationally infeasible. For example, even partitioning cells of an 8×8 checkerboard into two connected components generates over 1.2×10^{11} unique partitions (see <https://oeis.org/A068416>). Unfortunately, most redistricting problems are of much greater scale.⁴ Therefore, to compare an implemented redistricting plan against a set of other candidate plans, researchers and policy makers must resort to simulation methods.

Despite the widespread use of redistricting simulation methods in court cases, insufficient efforts have been made to examine whether or not they actually yield a representative sample of all possible redistricting plans in practice.⁵ Instead, some assume that the existing simulation methods work as intended. For example, in his amicus brief to the Supreme Court for *Rucho et al. v. Common Cause*, Eric Lander declares,⁶

With modern computer technology, it is now straightforward to generate a large collection of redistricting plans that are representative of all possible plans that meet the State’s declared goals (e.g., compactness and contiguity)

And yet, if there exists no scientific evidence that these simulation methods can actually yield a representative sample of valid redistricting plans, we cannot rule out the possibility that the comparison of a particular plan against sampled plans yields misleading conclusions about gerrymandering.

We argue that the empirical validation of simulation methods is essential for the credibility of

³The outlier detection method proposed by Chikina *et al.* (2017) is a statistical test and its goal is not uniform sampling. However, the proposed enumeration method can still be useful for assessing its empirical performance.

⁴While statutory guidelines and requirements such as district contiguity, population parity, and compactness reduce the number of partitions dramatically, the resulting problem currently remains out-of-reach of full enumeration methods.

⁵For an exception, see e.g., Carter *et al.* (2019), Jonathan C. Mattingly. “Rebuttal of Defendant’s Expert Reports for *Common Cause v. Lewis*.” Andrew Chin, Gregory Herschlag, and Jonathan C. Mattingly. “The Signature of Gerrymandering in *Rucho v. Common Cause*, pp. 1261–1262.

⁶Brief for Amicus Curiae Eric S. Lander In Support of Appellees, p. 4, *Rucho et al. v. Common Cause*, No. 18-422. March 7, 2019, page 4.

academic scholarship and expert testimony in court. We apply the recently developed computational method of Kawahara *et al.* (2017), `enumpart`, that efficiently enumerates all possible redistricting plans and obtains an independent sample from this population (Section 2). The algorithm uses a compact data structure, called the zero-suppressed binary decision diagram (ZDD) (Minato, 1993). In the aforementioned 8×8 checkerboard problem, explicitly storing every partition would require more than 1 terabyte of storage. In contrast, the ZDD needs only 1.5 megabytes. Our enumeration results are available as Fifield *et al.* (2020b) so that other researchers can use them to validate their own simulation methods. In addition, we will also make the code that implements the algorithm publicly available and incorporate it as part of an open-source R software package for redistricting, `redist` (Fifield *et al.*, 2015).

We begin by showing that the `enumpart` algorithm scales to a state with a couple of hundred geographical units, yielding realistic validation data sets (Section 3). We then test the empirical performance of existing simulation methods in two ways (Section 4). First, we randomly sample many submaps of various sizes from actual state shapefiles so that we average over the idiosyncratic geographic features and voter distributions of each map. For each sampled small map, we conduct a statistical test of the distributional equality between sampled and enumerated maps under various population parity constraints. If the simulation methods yield a representative sample of valid redistricting plans, then the distribution of the resulting p -values should be uniform. Second, we exploit the fact that even for a medium-sized redistricting problem, the `enumpart` algorithm can independently sample from the population of all valid redistricting plans. We then compare the resulting representative sample with the sample obtained using existing simulation methods. This second approach is applied to the actual redistricting problem in Iowa with 99 counties and a 250-precinct subset map from Florida, both of which are too computationally intensive for enumeration.

The overall conclusion of our empirical validation studies is that Markov chain Monte Carlo (MCMC) methods (e.g., Fifield *et al.*, 2014, 2020a; Mattingly and Vaughn, 2014; Carter *et al.*, 2019) substantially outperform so-called random-seed-and-grow (RSG) algorithms (e.g., Cirincione *et al.*, 2000; Chen and Rodden, 2013). These are two types of simulation methods that are most widely used in practice. Although the currently available MCMC methods are far from perfect and have much room for improvement, it is clear that the RSG algorithms are unreliable. Of course, showing that MCMC methods work reasonably well on these particular validation data sets does not necessarily imply that they will also perform well on other data sets especially larger scale redistricting problems. Rather, failing these validation tests on small and medium-scale redistricting

problems provides evidence that RSG methods are most likely to perform poorly when applied to other larger states.

To the best of our knowledge, the only publicly available validation data set for redistricting is the 25 precinct map obtained from Florida, for which Fifield *et al.* (2014, 2020a) enumerated all possible redistricting maps for two or three contiguous districts. Other researchers have used this validation data or enumeration method to evaluate their own algorithms (e.g., Magleby and Mosesson, 2018; Carter *et al.*, 2019). However, this data set is small and represents only a particular set of precincts representing a specific political geography, and may not be representative of other redistricting problems. For example, as noted by Magleby and Mosesson (2018), this data set is not particularly balanced — only eight partitions fall within standard levels of population parity ($\pm 1.5\%$), and most fall above 10%. Our new validation data sets are much larger and hence provide unique opportunities to conduct a more realistic empirical evaluation of simulation methods.

2 The Methodology

In this section, we describe the enumeration and sampling methods used in our empirical validation studies. Our methods are based on the `enumpart` algorithm originally developed by Kawahara *et al.* (2017) who showed how to enumerate all possible redistricting plans and store them using a compact data structure, called a zero-suppressed binary decision diagram (ZDD) (Minato, 1993). We also show how the `enumpart` algorithm can be used to independently sample from the population of contiguous redistricting plans.

2.1 The Setup

Following the literature (see e.g., Altman, 1997; Mehrotra *et al.*, 1998; Fifield *et al.*, 2014), we formulate redistricting as a graph-partitioning problem. Given a map of a state, each precinct (or any other geographical units used for redistricting) is represented by a vertex, whereas the existence of an edge between two vertices implies that they are geographically contiguous to one another. Formally, let $G = (V, E)$ represent a graph with the vertex set $V = \{v_1, \dots, v_n\}$ and the edge set $E = \{e_1, \dots, e_m\}$. We consider redistricting of a state into a total of p districts where all precincts of each district are connected. This is equivalent to partitioning a graph G into p connected components $\{V_1, V_2, \dots, V_p\}$ such that every vertex in V belongs to exactly one connected component, i.e., $V_1 \cup \dots \cup V_p = V$, $V_k \cap V_{k'} = \emptyset$ for any $k \neq k'$ and all the vertices in V_k are connected.

We use the fact that a p -graph partition can alternatively be represented as an edge set S . That is, by removing certain edges from E , we can partition G into p connected components. Formally,




Figure 1: A Running Redistricting Example. We consider dividing a state with six geographical units into two districts. The original map is shown in the left panel where the shaded area is uninhabited. The middle panel shows its graph representation, whereas the right panel shows an example of redistricting map represented by an induced subgraph, which consists of a subset of edges.

for each connected component V_k , we define an induced subgraph $(V_k, S(V_k))$ as a graph whose edge set consists of all edges whose two endpoints (i.e., the two vertices directly connected by the edge) belong to V_k . Then, the p -graph partition can be defined as the union of these induced subgraphs, i.e., $\mathcal{P} = \bigcup_{k=1}^p S(V_k)$ where $S(V_k) \cap S(V_{k'}) = \emptyset$ for any $k \neq k'$. Our initial task is to enumerate *all* possible p -graph partitions of G .

Figure 1a presents the running example used throughout this section to illustrate our methodology. In this hypothetical state, we have a total of 6 precincts, represented as vertices $\{v_1, v_2, \dots, v_6\}$, which we hope to divide into 2 districts, $\{V_1, V_2\}$. A grey area is uninhabited (e.g., lake). This map can be represented as a graph of Figure 1b where two contiguous vertices share an edge. Consider a redistricting map with $V_1 = \{v_1, v_3\}$ and $V_2 = \{v_2, v_4, v_5, v_6\}$. As shown in Figure 1c, this redistricting map can be represented by an induced subgraph after removing three edges, i.e., $\{e_1, e_3, e_5\}$. Thus, we can represent each district as an induced subgraph, which is a set of edges, i.e., $S(V_1) = \{e_2\}$ or $S(V_2) = \{e_4, e_6, e_7\}$.

2.2 Graph Partitions and Zero-suppressed Binary Decision Diagram (ZDD)

A major challenge for enumerating redistricting maps is memory management because the total number of possible maps increases exponentially. We use the Zero-suppressed Binary Decision Diagram (ZDD), which uses a compact data structure to efficiently represent a family of sets (Minato, 1993). We first discuss how the ZDD can represent a family of graph partitions before explaining how we construct the ZDD from a given graph.

The ZDD that corresponds to the running example of Figure 1 is given in Figure 2. A ZDD is a directed acyclic graph. As is clear from the figure, each edge of the original graph corresponds to possibly multiple nodes of a ZDD. To avoid confusing terminology, we use a “node” rather than a




Figure 2: Zero-suppressed Binary Decision Diagram (ZDD) for the Running Example of Figure 1b. The blue path corresponds to the redistricting map represented by the induced subgraph in Figure 1c.

“vertex” to refer to a unit of ZDD, which represents an edge of the original graph. Similarly, we call an edge of the ZDD an “arc” to distinguish it from an edge of the original graph. There are three special nodes in a ZDD. The *root node*, labeled as e_1 in our example, has no incoming arc but, like other nodes, represents one of the edges in the original graph. We will discuss later how we label nodes. ZDD also has two types of *terminal nodes* without an outgoing arc, called 0-terminal and 1-terminal nodes and represented by $\boxed{0}$ and $\boxed{1}$, respectively. Unlike other nodes, these terminal nodes do not correspond to any edge in the original graph. Finally, each non-terminal node, including the root node, has exactly two outgoing arcs, *0-arc* (dashed arrow) and *1-arc* (solid arrow).

Given a ZDD, we can represent a graph partition as the set of edges that belong to a directed

path from the root node to 1-terminal node and have an outgoing 1-arc. For example, the path highlighted by blue, $e_1 \dashrightarrow e_2 \rightarrow e_3 \dashrightarrow e_4 \rightarrow e_5 \dashrightarrow e_6 \rightarrow e_7 \rightarrow \boxed{1}$, represents the edge set $\{e_2, e_4, e_6, e_7\}$, which corresponds to the 2-graph partition shown in Figure 1c. Indeed, there is one to one correspondence between a graph partition and a path of a ZDD.

2.3 Construction of the ZDD

How should we construct a ZDD for a p -graph partition from a given graph? We use the frontier-based search algorithm proposed by Kawahara *et al.* (2017). The algorithm grows a tree starting with the root node in a specific manner. We first discuss how to construct a ZDD given m labeled edges, $\{e_1, \dots, e_m\}$, where e_1 represents the root node. We then explain how we merge nodes to reduce the size of the resulting ZDD and how we label edges given a graph to be partitioned so that the computation is efficient.

2.3.1 The Preliminaries

Starting with the root node $i = 1$, we first create one outgoing 0-arc and one outgoing 1-arc from the corresponding node e_i to the next node e_{i+1} . To ensure that each enumerated partitioning has exactly p connected components, we store the number of *determined connected components* as the `dcc` variable for each ZDD node. Consider a directed path $e_1 \rightarrow e_2 \dashrightarrow e_3 \dashrightarrow e_4 \dashrightarrow e_5$. In this example, e_1 is retained whereas edges $\{e_2, e_3, e_4\}$ are not. We know that the two vertices, $\{v_1, v_2\}$, together form one district, regardless of whether or not e_5 is retained. Then, we say that a connected component is determined and set `dcc` to 1 for e_5 . If `dcc` exceeds p , then we create an arc into the 0-terminal node rather than create an arc into the next node since there is no longer a prospect of constructing a valid partition. Similarly, when creating an arc out of the final node, e_m , we point the arc into the 0-terminal node if `dcc` is less than p , which represents the total number of partitions. Finally, if the number of remaining edges exceeds $p - \text{dcc}$, we stop growing the path by creating an outgoing arc into the 0-terminal node.

How do we find out when another connected component is determined so that `dcc` needs to be increased? To do this, we need two new variables. First, for each vertex v_i , we store the *connected component number*, denoted by `comp` $[v_i]$, indicating the connected component to which v_i belongs. Thus, two vertices, v_i and $v_{i'}$, share an identical connected component number if and only if they belong to the same connected component, i.e., `comp` $[v_i] = \text{comp}[v_{i'}]$.

We initialize the connected component number as `comp` $[v_i] \leftarrow i$ for $i = 1, 2, \dots, n$ where n is the number of vertices in the original graph. Suppose that we process and retain an edge incident

to two vertices v_i and $v_{i'}$ for $i \neq i'$ by creating an outgoing 1-arc. Then, we set $\text{comp}[v_j] \leftarrow \max\{\text{comp}[v_i], \text{comp}[v_{i'}]\}$ for any vertex v_j whose current connected component number is given by $\text{comp}[v_j] = \min\{\text{comp}[v_i], \text{comp}[v_{i'}]\}$. This operation ensures that all vertices that are connected to v_i or $v_{i'}$ have the same connected component number (larger of the two original numbers).

2.3.2 The Frontier-based Search

Next, we define the *frontier* of a graph, which changes as we process each edge and grow a tree. Suppose that we have created a directed path by processing the nodes from e_1 up to e_ℓ where $\ell = 2, 3, \dots, m-1$. For each $\ell = 1, 2, \dots, m-1$, the ℓ th frontier F_ℓ represents the set of vertices of the original graph that are incident to both a processed edge (i.e., at least one of e_1, e_2, \dots, e_ℓ) and an unprocessed edge (i.e., at least one of $e_{\ell+1}, e_{\ell+2}, \dots, e_m$). Note that we define $F_0 = F_m = \emptyset$ and that the set of processed edges includes the one currently being processed. Thus, for a given graph, the frontier only depends on which edge is being processed but does not hinge on how edges have been or will be processed. That is, the same frontier results for each node regardless of paths.

The frontier can be used to check whether a connected component is determined. Specifically, suppose there exists a vertex v that belongs to the previous frontier but is not part of the current one, i.e., $v \in F_{\ell-1}$ and $v \notin F_\ell$. Then, if there is no other vertex in F_ℓ that has the same connected component number as v (that is, no vertex in F_ℓ is connected to v), there will not be another vertex in subsequent frontiers, i.e., $F_{\ell+1}, \dots, F_m$, that are connected to v . Thus, under this condition, the connected component $\text{comp}[v]$ is determined, and we increment dcc by one.

Figure 3 gives an example of computing the connected component number, constructing the frontier, and updating the determined connected components, based on the redistricting problem shown in Figure 1. In each graph, a positive integer placed next to a vertex represents its connected component number, whereas the vertices grouped by the solid line represent a frontier. For example, when processing edge e_5 (see Figure 3e), we have $F_4 = \{v_3, v_4\}$ and $F_5 = \{v_4, v_5\}$. Since there is no vertex in F_5 that shares the same connected component number as v_3 (which is 1), we can determine the first connected component and increment dcc by one.

Finally, when processing the last edge e_m represented by node n^* , if two vertices incident to the edge belong to the same connected component number, then the 0-arc from node n^* points to the 0-terminal node whereas the destination of the 1-arc is the 1-terminal node unless $\text{dcc} \neq p$. If they have different connected component numbers, the 0-arc of node n^* goes to the 1-terminal node whereas the destination of its 1-arc is the 1-terminal node so long as $\text{dcc} = p$ and the induced subgraph condition described in the next paragraph is satisfied (otherwise, it is the 0-terminal node).




Figure 3: Calculation of the Frontier, the Connected Component Number, and the Determined Connected Components. This illustrative example is based on the redistricting problem shown in Figure 1. A positive integer placed next to each vertex represents the connected component number, whereas the vertices grouped by the solid blue line represent a frontier. A connected component is determined when processing edge e_5 .

Throughout the process of building a ZDD, we must make sure that every path actually corresponds to an induced subgraph, which is defined as a subset of nodes and all arcs connecting pairs of such nodes. We call this the induced subgraph condition. Consider a path, $e_1 \rightarrow e_2 \rightarrow e_3$. Since three vertices, $\{v_1, v_2, v_3\}$, are connected, we must retain edge e_3 because its two incident vertices, v_2 and v_3 , are connected. Thus, we have $e_1 \rightarrow e_2 \rightarrow e_3 \dashrightarrow \boxed{0}$. Similarly, consider a path, $e_1 \rightarrow e_2 \dashrightarrow e_3$. We cannot retain e_3 because e_2 , which is incident to v_1 and v_3 , is not retained. This yields $e_1 \rightarrow e_2 \dashrightarrow e_3 \rightarrow \boxed{0}$.

To impose the induced subgraph condition, we introduce the *forbidden pair set* for each node. Once we decide not to use an edge that connects two distinct components, the two components must not be connected any more. Otherwise, the new component generated by connecting the two components has an unused edge, violating the induced subgraph condition. Therefore, if we determine that an edge $\{v, v'\}$ is not used, the addition of $\{\text{comp}[v], \text{comp}[v']\}$ to the forbidden pair set reminds us that the components $\text{comp}[v]$ and $\text{comp}[v']$ must not be connected. That is, if we use an edge $\{u, u'\}$ and the forbidden pair set contains $\{\text{comp}[u], \text{comp}[u']\}$, the path will be sent to the 0-terminal. In the above example, if we pass through $e_1 \rightarrow e_2 \dashrightarrow e_3$, $\{2, 3\}$ is added to the forbidden pair set, where 2 is the component number of $\{v_1, v_2\}$ and 3 is that of $\{v_3\}$. Then, since



Figure 4: An Example of Node Merging. As shown in Figure 2, these two paths merge at e_3 because the connected component numbers in F_2 are identical and the number of determined connected components is zero.

retaining e_3 violates the induced subgraph condition, we have $e_1 \longrightarrow e_2 \dashrightarrow e_3 \longrightarrow \boxed{0}$.

2.3.3 Node Merge

The above operation implies that when processing e_ℓ , the only required information is the connectivity of vertices in $F_{\ell-1}$. We can reduce the size of the ZDD by exploiting this fact. First, we can avoid repeating the same computation by merging multiple nodes if the connected component numbers of all vertices in $F_{\ell-1}$ and the number of determined connected components dcc are identical. This is a key property of the ZDD, which allows us to efficiently enumerate all possible redistricting plans by merging many different paths. Second, we only need to examine the connectivity of vertices within a frontier in order to decide whether or not any connected component is determined. Thus, we adjust the connected component number so that it equals the maximum vertex number in the frontier. That is, if some vertices in the frontier share the same connected component number, we change it to the maximum vertex index among those vertices. For example, in Figure 3b, we set $\text{comp}[v_2] = 2$ and $\text{comp}[v_3] = 3$. We need not worry about how the renumbering of $\text{comp}[v_3]$ affects the value of $\text{comp}[v_1]$ because $v_1 \notin F_2$. This operation results in merging of additional nodes, reducing the overall size of the ZDD.

Figure 4 gives an example of such a merge. Figure 4a corresponds to the path, $e_1 \longrightarrow e_2 \dashrightarrow e_3$ whereas Figure 4b represents the path, $e_1 \dashrightarrow e_2 \longrightarrow e_3$. As shown in Figure 2, these two paths are merged at e_3 because the connected component numbers in their frontier F_2 are identical and both have the same number of determined connected component, i.e., $\text{dcc} = 0$. Note that in Figure 4b the connected component number is normalized within the frontier F_2 such that the connected component number of v_2 is the maximum vertex index, i.e., 2, and that of v_3 is 3.

Node merging plays a key role in scaling up the enumeration algorithm. Although we can construct the ZDD that only enumerates graph partitions by storing the sum of population values




Figure 5: An Example of Edge Ordering by Vertex Cuts. To order edges, we choose two vertices with the maximum shortest distance and call them s and t . We then use the minimum vertex cut, indicated by the dashed oval, to create two or more connected components, which are arbitrarily ordered. The same procedure is then applied to each connected component until the resulting connected components are sufficiently small.

into each node (see Section 4 of Kawahara *et al.*, 2017), this prevents nodes from being merged, dramatically reducing the scalability of the enumeration algorithm. Therefore, we do not take this approach here.

2.3.4 Edge Ordering

How should we label the edges of the original graph? The amount of computation depends on the number of nodes in the ZDD. Recall that two nodes are merged if the stored values such as `comp` and `dcc` are identical. Since a ZDD node stores the `comp` value for each vertex in the frontier, the number of unique stored values grows exponentially as the frontier size increases (see Section 3.1 in Kawahara *et al.* (2017) for the detailed analysis). Therefore, we wish to label the edges of a graph such that the maximum size of the frontier is minimized. We take a heuristic approach here. Specifically, we first choose two vertices s, t such that the shortest distance between s and t is as large as possible across all vertex pairs. We use the Floyd-Warshall algorithm, which can find the shortest paths between all vertex pairs in $\mathcal{O}(|V|^3)$ where $|V|$ is the number of vertices of a graph. Next, we compute the minimum s - t vertex graph cut, which is the minimum set of vertices whose removal generates two or more connected components. To do this, we use a max-flow based algorithm, and arbitrarily order the resulting connected components. Finally, we recursively apply this procedure to each connected component until the resulting connected components are sufficiently small (e.g., 5 edges), at which point they are ordered in an arbitrary fashion.

Figure 5 illustrates this process. In this example, a pair, $s = v_1$ and $t = v_6$, gives the maximum shortest distance. Given this choice, there are four minimum s - t graph cuts whose size is 2, i.e., $\{v_2, v_3\}, \{v_2, v_5\}, \{v_3, v_4\}, \{v_4, v_5\}$. We arbitrarily select one of them and call it S . Suppose we set $S = \{v_4, v_5\}$. Then, this yields two connected components, i.e., $C_1 = \{v_1, v_2, v_3\}$ and $C_2 = \{v_6\}$. For each connected component C_i , let E_i represent the set of edges in C_i and between C_i and S . In the current example, $E_1 = \{\{v_1, v_2\}, \{v_1, v_3\}, \{v_2, v_3\}, \{v_2, v_4\}, \{v_3, v_5\}\}$ and $E_2 = \{\{v_4, v_6\}, \{v_5, v_7\}\}$. We order these edge sets so that all the edges E_1 will be placed before those of E_2 . To continue this process recursively, we combine all the vertices in S into a single vertex and let this new vertex be t in E_1 and s in E_2 . Now, we can apply the same procedure separately to E_1 and E_2 : computing the minimum s - t vertex cut and splitting the graph into two (or more) components.

The reason why we expect the above edge ordering procedure to produce a small frontier is that each vertex cut in the process equals one of the frontiers of the corresponding ZDD. In our example, the first vertex cut S is equal to $F_5 = \{v_4, v_5\}$. Since we choose minimum vertex cuts in each step, we expect the input graph with the edge order obtained through this procedure to have small frontiers.

2.4 Enumeration and Independent Sampling

It can be shown that every path from the root node to the 1-terminal node in the resulting ZDD has a one-to-one correspondence to a p -graph partition. This is because each p -graph partition can be uniquely represented by the union of induced subgraphs, which in turn corresponds to a unique path from the root node to the 1-terminal node. The complexity of the `enumpart` algorithm is generally difficult to characterize, but Kawahara *et al.* (2017) analyzes it in the case of planar graphs. Thus, once we obtain the ZDD as described above, we can quickly enumerate all the paths from the root node to the 1-terminal node. Specifically, we start with the 1-terminal node and then proceed upwards to the root node, yielding a unique graph partition.

In addition to enumeration, we can also independently sample p -graph partitions (Knuth, 2011). First, for each node ν of the ZDD, we compute the number of paths to the 1-terminal node. Let $c(\nu)$ be the number of such paths, and ν_0 and ν_1 be the nodes pointed by the 0-arc and 1-arc of ν , respectively. Clearly, we have $c(\nu) = c(\nu_0) + c(\nu_1)$. The values of c for the 0-terminal and 1-terminal nodes are 0 and 1, respectively. As done for enumeration, we compute and store the value of c for each node by moving upwards from the terminal node to the root node. Finally, we conduct random sampling by starting with the root node and choosing node ν_1 with probability $c(\nu_1)/\{c(\nu_0) + c(\nu_1)\}$ until we reach the 1-terminal node. Since the probability of reaching the 0-terminal node is zero, we will always arrive at the 1-terminal node, implying that we obtain a path corresponding to a p -graph

partition. Repeating this procedure will yield the desired number of independently and uniformly sampled p -graph partitions.

The reason why this procedure samples redistricting maps uniformly is that a path from the root node to the 1-terminal node corresponds to a unique p -graph partition. Since each node ν stores the number of paths from the node to the 1-terminal node as $c(\nu)$, the sampling procedure uniformly and randomly selects one path among $c(v_1)$ paths where v_1 is the root node. If non-uniform sampling is desired, one could apply the sampling-importance resampling algorithm to obtain a representative sample from a target distribution (Rubin, 1987).

3 Empirical Scalability Analysis

This section analyzes the scalability of the `enumpart` algorithm described above, and shows that the algorithm scales to enumerate partitions of maps many times larger than existing enumeration procedures. We analyze the algorithm’s scalability in terms of runtime and memory usage, and show how the memory usage of `enumpart` is closely tied to the frontier of the corresponding ZDD as explained in the previous section.

To make this empirical analysis realistic, we use independently constructed and contiguous subsets of the 2008 New Hampshire precinct map for maps ranging between 40 precincts and 200 precincts, increasing by 40, i.e., $\{40, 80, 120, 160, 200\}$. The original New Hampshire map consists of 327 precincts, which are divided into two congressional districts. To generate an independent contiguous subset of the map, we first randomly sample a precinct, and add its adjacent precincts to a queue. We then repeatedly sample additional precincts from the queue to be added to the subset map, and add the neighbors of the sampled precincts to the queue, until the map reaches the specified size. We repeat this process until the subsetted map reaches a pre-specified size.

We consider partitioning each of these maps into two, five, or ten districts and apply `enumpart` to each case. We then compute the time and memory usage of generating a ZDD for each application. For each precinct size and number of districts, we repeat the above sampling procedure 25 times, producing 25 independent and contiguous subsets of the New Hampshire map. All trials were run on a Linux computing cluster with 530 nodes and 48 Intel Cascade Lake cores per node, where each node has 180GB of RAM. Note that we do not save the results of enumeration to disk as doing so for every trial is computationally too expensive. This means that we cannot conduct an in-depth analysis of the characteristics of all enumerated maps.

Figure 6 shows the results of our scalability analysis. The top row shows scalability results for

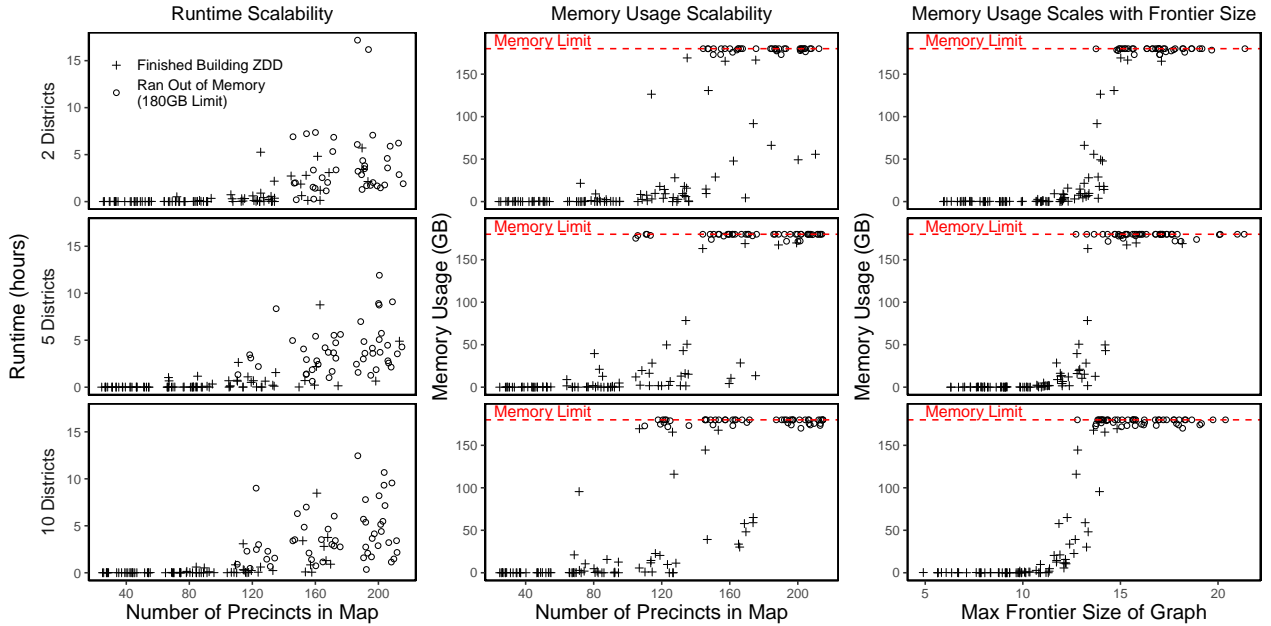


Figure 6: The Scalability of the `enumpart` Algorithm on Subsets of the New Hampshire Precinct Map. This figure shows the runtime scalability of the `enumpart` algorithm for building the ZDD on random contiguous subsets of the New Hampshire precinct map. Crosses indicate maps where the ZDD was successfully built within the RAM limit of 180GB. In contrast, open circles represent maps where the algorithm ran out of memory. For the left and middle columns, the results are jittered horizontally with a width of 20 for the clarity of presentation (The actual evaluation points on the horizontal axis are 40, 80, 120, 160, and 200). The left column shows how total runtime increases with the number of units in the underlying map, while the center column shows how the total RAM usage increases with the number of units in the underlying map. Lastly, the right-hand column shows that memory usage is primarily a function of the maximum frontier size of the ZDD. We show results for 2-district partitions (top row), five-district partitions (middle row), and 10-district partitions (bottom row).

generating a ZDD for partitions of the map into two districts, while the middle row shows the results with five districts and the bottom row shows results with ten districts. Each dot represents a run of the `enumpart` algorithm on a subset of the New Hampshire map. Crosses represent trials where the ZDD successfully built using under the 180GB RAM limit. In contrast, open circles show trials that were unable to build the ZDD with the same RAM limit. Note that for the left and middle columns, the results are jittered horizontally with a width of 20 for the clarity of presentation.

The left-hand and center columns show how the `enumpart` algorithm scales in terms of runtime and memory usage, respectively, as the number of precincts in the underlying map increases. For small maps ranging from 25 precincts to 80 precincts, runtime and memory usage are for the most part negligible. The ZDD for two-district, five-district, and ten-district partitions for these small maps can be constructed in nearly all cases in under two minutes, and using less than one gigabyte of RAM. As the number of precincts in the map starts to increase, so do the runtime and memory

usage requirements. For maps of 200 precincts, over 90% of the tested maps hit the 180GB memory limit before building the complete ZDD. For all map sizes, we also note that the runtime and memory usage requirements for building the ZDD do not appear to depend much on the number of districts that the map is being partitioned into.

What drives these patterns in scalability? While the number of units in the underlying map predicts both runtime and memory usage, there is still a great deal of variability even conditional on the number of precincts in the map. In the right-hand column, we show that the memory usage requirements for building the ZDD are closely tied to the maximum frontier size of the underlying map, as defined in Section 2.3.2. While the memory usage is minimal so long as the maximum frontier size of the graph is under 11, memory usage increases quickly once the maximum frontier size grows beyond that. This suggests that improved routines for reducing the size of a map’s frontier can allow for the enumeration of increasingly large maps.

4 Empirical Validation Studies

In this section, we introduce a set of new validation tests and datasets that can be used to evaluate the performance of redistricting simulation methods. We focus on the two most popular types of simulation methods that are implemented as part of the open-source software package `redist` (Fifield *et al.*, 2015): one based on the Markov chain Monte Carlo (MCMC) algorithm (Fifield *et al.*, 2014, 2020a; Mattingly and Vaughn, 2014) and the other based on the random-seed-and-grow (RSG) algorithm (Cirincione *et al.*, 2000; Chen and Rodden, 2013). Below, we conduct empirical validation studies both through full enumeration and independent sampling. For the sake of simplicity, we use the uniform distribution over all valid redistricting maps under various constraints. However, one could also use a non-uniform target distribution by appropriately weighting each map.

4.1 Validation through Enumeration

We conduct two types of validation tests using enumeration. We first use the `enumpart` algorithm to enumerate all possible redistricting plans using a map with 70 precincts, which is much larger than the existing validation map with 25 precincts analyzed in Fifield *et al.* (2014, 2020a). We then compare the sampled redistricting plans obtained from simulation methods against the ground truth based on the enumerated plans. The second approach is based on many smaller maps with 25 precincts. We then assess the overall performance of simulation methods across these many maps rather than focusing on a specific map.

4.1.1 A New 70-precinct Validation Map

The top left plot of Figure 7 introduces a new validation map with 70 precincts and their population, which is a subset of the 2008 Florida precinct map consisting of 6,688 precincts with 25 districts. We use the `enumpart` algorithm to enumerate every partition of this map into two districts, which took approximately 8 hours on a MacBook Pro laptop with 16GB RAM and 2.8 GHz Intel i7 processors. Nearly all of this time was spent writing the partitions to disk — building the ZDD for this map took under half a second.

The histograms of the figure shows the number of redistricting plans that satisfy the deviation from population parity up to 20 percentage points (by one percentage point increments, i.e., $[0, 0.01)$, $[0.01, 0.02)$, ..., $[0.19, 0.2)$). The deviation from population parity is defined as,

$$\max_{1 \leq k \leq p} \frac{|P_k - \bar{P}|}{\bar{P}} \tag{1}$$

where P_k represents the population of the k th district, $\bar{P} = \sum_{k=1}^p P_k/p$, and p is the total number of districts. When partitioning this map into two districts, there exist a total of 44,082,156 possible redistricting plans if we only impose the contiguity requirement.

As shown in the upper right plot, out of these, over 700,000 plans are within a 1% population parity constraint. As we relax the population parity constraint, the cumulative number of valid redistricting plans gradually increases, reaching over 3 and 7 million plans for the 5% and 10% population parity constraints, respectively. Thus, this validation map represents a more realistic redistricting problem than the validation map analyzed in Fifield *et al.* (2014, 2020a). That dataset, which enumerates all 117,688 partitions of a 25-precinct subset of the Florida map into three districts, includes only 8 plans within 1% of population parity, and 927 plans within 10% of population parity.

In addition to the population parity, we also consider compactness constraints. Although there exist a large number of different compactness measures, for the sake of illustration, we use the Relative Proximity Index (RPI) proposed by Fryer and Holden (2011). The RPI for a given plan π_s in the valid set of redistricting plans π is defined as,

$$\text{RPI}(\pi_s) = \frac{\sum_{k=1}^p \sum_{i \in V_k} \sum_{j \in V_k} P_i P_j D_{ij}^2}{\arg \min_{\pi_s \in \pi} \sum_{k=1}^p \sum_{i \in V_k} \sum_{j \in V_k} P_i P_j D_{ij}^2} \tag{2}$$

where P_i corresponds to the population for precinct i assigned to district k , and D_{ij} corresponds to the distance between precincts i and j assigned to district k . Thus, a plan with a lower RPI is more compact.

We consider two compactness thresholds based on the RPI values: 25th and 75th percentiles, which equal 1.76 and 1.44, respectively. As shown in the bottom left histogram of Figure 7, the




Figure 7: A New 70-Precinct Validation Map and the Histogram of Redistricting Plans under Various Population Parity and Compactness Constraints. The underlying data is a 70-precinct contiguous subset of the Florida precinct map, for which the `enumpart` algorithm enumerated every 44,082,156 partitions of the map into two contiguous districts. In the histograms, each bar represents the number of redistricting plans that fall within a 1 percentage point range of a certain population parity, i.e., $[0, 0.01)$, $[0.01, 0.02)$, ..., $[0.19, 0.20)$. The 25th (75th) percentile compactness constraint is defined as the set of plans that are more compact than the 25th (75th) percentile of maps within the full enumeration of all plans for the 70-precinct map, using the Relative Proximity Index to measure compactness. The annotations reflect the exact number of plans which meet the constraints. For example, when no compactness constraint is applied, there are 3,678,453 valid plans when applying a 5% population parity constraint, and 717,060 valid plans when applying a 1% population parity constraint. Under the strictest constraints, the 1% population parity constraint and 75th percentile compactness constraint, there are 271,240 valid plans.

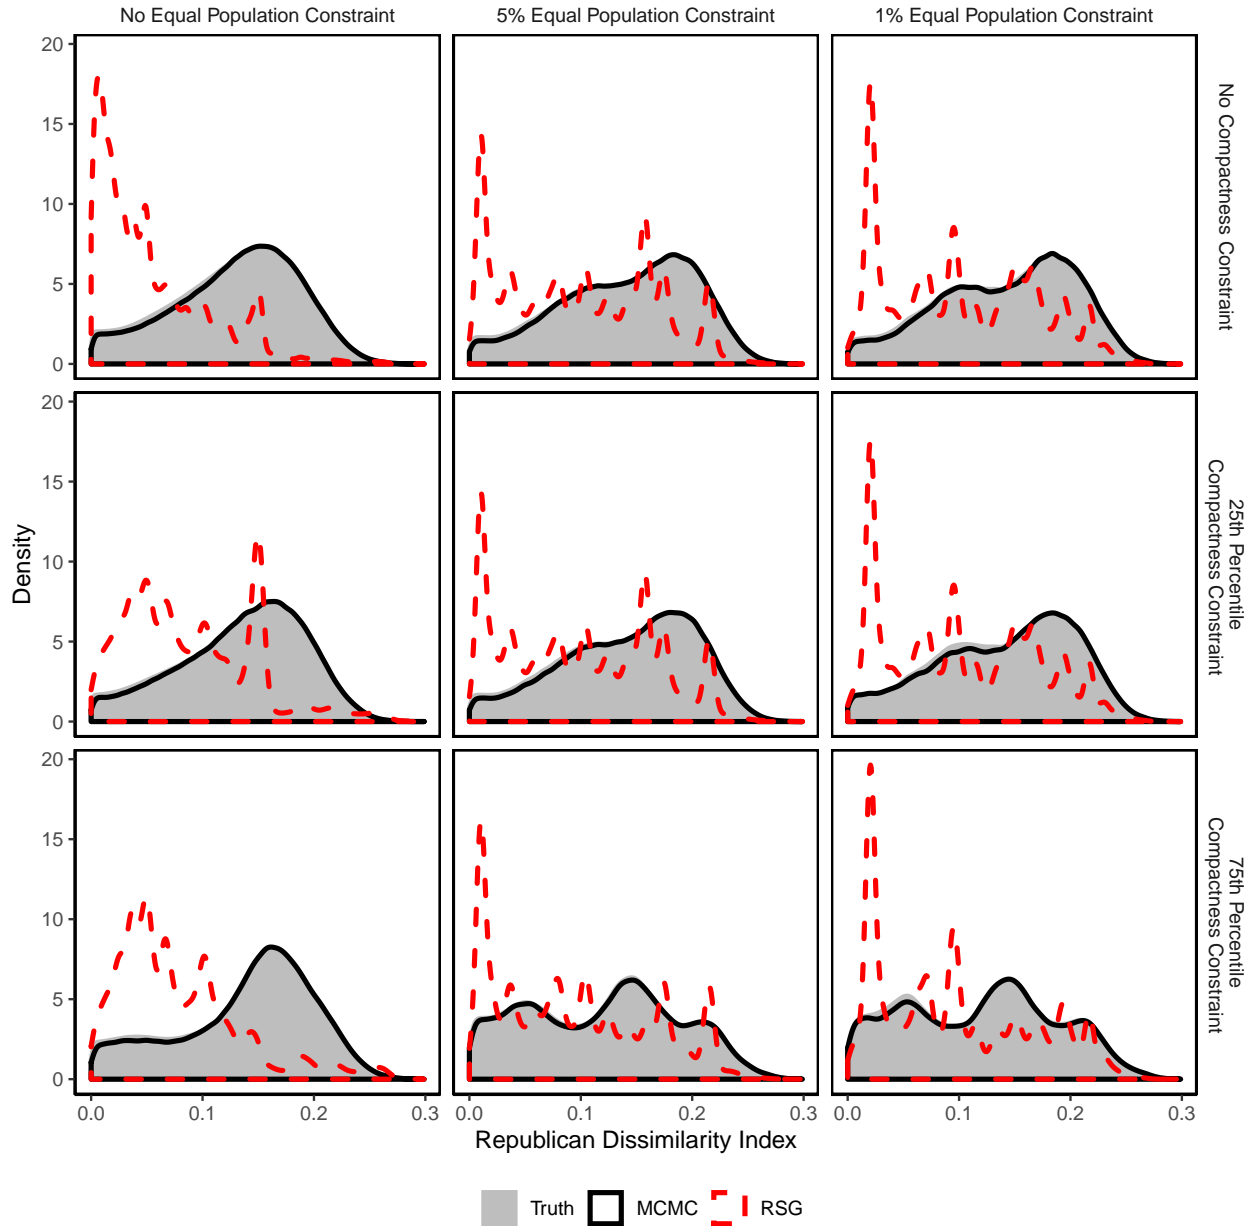


Figure 8: A Validation Study Enumerating all Partitions of a 70-Precinct Map into Two Districts. The underlying data is the 70-precinct contiguous subset introduced in the left plot of Figure 7. Unlike the random-seed-and-grow (RSG) method (red dashed lines), the Markov chain Monte Carlo (MCMC) method (solid black line) is able to approximate the target distribution. The 25th percentile (75th percentile) Compactness Constraint is defined as the set of plans that are more compact than the 25th (75th) percentile of maps within the full enumeration of all plans for the 70-precinct map, using the Relative Proximity Index to measure compactness.

25th percentile constraint does little beyond the population constraint. The number of valid plans that satisfy a 5% population parity threshold remains identical even after imposing this compactness constraint. However, the 75th percentile compactness constraint dramatically reduces the number of valid plans as seen in the bottom right histogram. For example, it reduces the total number of plans that meet the 1% population parity threshold by more than 70 percent.

Figure 8 shows the performance of the MCMC and RSG simulation methods using the new 70-precinct validation map (see Algorithms 1 and 3 of Fifield *et al.*, 2014, 2020a, respectively, for the details of implementation). The solid grey density shows the true distribution of the Republican dissimilarity index (Massey and Denton, 1988) on the validation map, which is defined as

$$D = \frac{1}{2} \sum_{k=1}^p \frac{P_k}{P} \cdot \frac{|R_k - R|}{R(1 - R)} \quad (3)$$

where k indexes districts in a state, P_k is the population of district k , R_k is the share of district k that voted for the Republican presidential candidate in 2008, P is the total population in the state, and R is the voteshare for the Republican presidential candidate across all districts.

The red dashed lines show the distribution of the Republican dissimilarity index of the RSG algorithm. Solid black lines show the distribution of the Republican dissimilarity index on plans drawn by the MCMC algorithm. In cases where we impose a population parity target, we specify a target distribution of plans using the Gibbs distribution where plans closer to population parity are more likely to be sampled by the algorithm (see Fifield *et al.*, 2014, 2020a, for details).

Similarly, when a compactness constraint is imposed, we specify a target Gibbs distribution such that more compact plans are sampled. Note that in typical redistricting applications, we would not know the denominator of equation (2). Fryer and Holden (2011) derive a power-diagram approach to finding a plan that approximately minimizes the denominator. Since we have enumerated all possible plans in the current setting, we simply use the true minimum value. However, this has no impact on the performance of the algorithm, since it is absorbed into the normalizing constant of the target distribution.

Specifically, we use the following target Gibbs distribution,

$$f_{\beta}(\pi_s) = \frac{1}{z(\beta)} \exp \left\{ - \sum_{k=1}^p (\beta_p \psi_k^p + \beta_c \psi_k^c) \right\}, \quad (4)$$

where

$$\psi_k^p = \frac{|P_k - \bar{P}|}{\bar{P}} \quad \text{and} \quad \psi_k^c = \frac{\sum_{i \in V_k} \sum_{j \in V_k} P_i P_j D_{ij}^2}{\arg \min_{\pi_s \in \pi} \sum_{k=1}^p \sum_{i \in V_k} \sum_{j \in V_k} P_i P_j D_{ij}^2}.$$

In this formulation, the strength of each constraint is governed by separate temperature parameters β_p (for population parity) and β_c (for compactness), where higher temperatures increase the likelihood that plans closer to the population parity or compactness target will be sampled. Once the algorithm is run, we discard sampled plans that fail to meet the target population and compactness constraints, and then reweight and resample the remaining plans so that they approximate a uniform sample from the population of all plans satisfying the constraints. After some initial tuning, we

selected $\beta_p = 10$ for the 5% equal population constraint, and $\beta_p = 50$ for the 1% equal population constraint. We selected $\beta_c = .001$ for the 25th percentile compactness constraint and $\beta_c = .01$ for the 75th percentile compactness constraint. When the population or compactness constraints are not applied, we set their corresponding temperature parameter to 0.

The RSG algorithm was run for 1,000,000 independent draws for each population constraint, while the MCMC algorithms were run for 250,000 iterations using 8 chains for each pair of constraints. Starting plans for each MCMC chain were independently selected using the RSG algorithm. The Gelman-Rubin diagnostic Gelman and Rubin (1992), a standard diagnostic tool for MCMC methods based on multiple chains, suggests that all MCMC chains had converged after at most 30,000 iterations. Unfortunately, the RSG algorithm does not come with such a diagnostic and hence we simply run it until it yields the same number of draws as the MCMC algorithms for the sake of comparison.

It is clear that on this test map, the RSG algorithm is unable to obtain a representative sample of the target distribution, at any level of population parity or compactness. This finding is consistent with the fact that the RSG algorithm is a heuristic algorithm with no theoretical guarantees and no specified target distribution. In contrast, the MCMC algorithm is able to approximate the target distribution, across all levels of population parity and compactness tested.

4.1.2 Many Small Validation Maps

A potential criticism of the previous validation study is that it is based on a single map. This means that even though it is of much larger size than the previously available validation map, the results may depend on the idiosyncratic geographical and other features of this particular validation map. To address this, we conduct another study based on many small validation maps. Specifically, we use our algorithm to enumerate all possible redistricting plans for each of 200 independent 25-precinct subsets of the 2008 Florida map. We then evaluate the performance of simulation methods for each validation map. Since we do not tune the temperature parameter of the MCMC algorithm for each simulated map unlike what one would do in practice, this yields a simulation setting that poses a significant challenge for the MCMC algorithm.

To assess the overall performance across these validation maps, we use the Kolmogorov-Smirnov (KS) statistic to test the distributional equality of the Republican dissimilarity index between the enumerated plans and the simulated plans. To increase the independence across simulated plans, we run the MCMC and RSG algorithms for 5 million iterations each on every map and then thinning by 500 (i.e., taking every 500th posterior draw). Without thinning, there is a significant amount

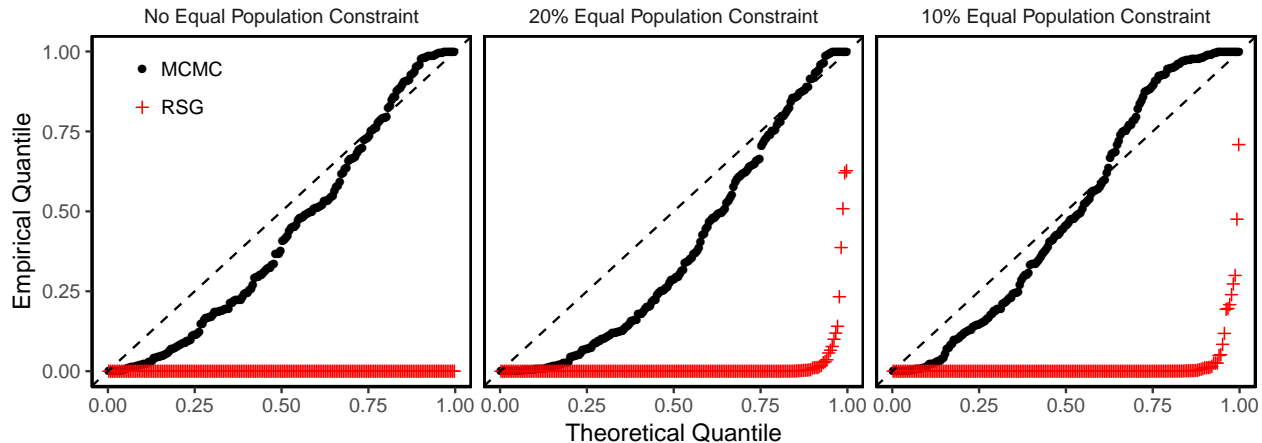


Figure 9: Quantile-Quantile Plot of p -values based on the Kolmogorov-Smirnov (KS) Tests of Distributional Equality between the Enumerated and Simulated Plans across 200 Validation Maps and under Different Population Parity Constraints. Each dot represents the p -value from a KS test comparing the empirical distribution of the Republican dissimilarity index from the simulated and enumerated redistricting plans. Under independent and uniform sampling, we expect the dots to fall on the 45-degree line. The MCMC algorithm (black dots), although imperfect, significantly outperforms the RSG algorithm (red crosses). See Figure 14 in the appendix for discussion of thinning values.

of autocorrelation across draws, with the autocorrelation typically ranging between 0.75 and 0.85 between adjacent draws and between 0.30 and 0.60 for draws separated by 5 iterations. In contrast, when thinning the Markov chain by 500, the autocorrelation between adjacent draws falls to under 0.05. When thinning the Markov chain by 1000, the results are approximately the same, as seen in Figure 14.

Although this does not make simulated draws completely independent of one another, we compute the p -value under the assumption of two independently and identically distributed samples. If the simulation methods are successful and the independence assumption holds, then we should find that the distribution of p -values across 200 small validation maps should be approximately uniform. After some initial tuning, we set the temperature parameter of the MCMC algorithm such that $\beta_p = 1$ for the 20% equal population constraint, and $\beta_p = 5$ for the 10% equal population constraint. These values are used throughout the simulations. In other words, unlike what one would do in practice, no data-specific tuning is performed, leading to a setting that is not favorable to the MCMC algorithm. After running the simulations, we again discard plans falling outside of the specified parity threshold and then reweight and resample the remaining plans to approximate a uniform draw from the target distribution of plans satisfying the specified parity constraint. We then calculate the KS test p -value by comparing the reweighted and resampled set of plans against the true distribution.

Figure 9 shows the results of this validation study. The left plot shows how the MCMC (black dots) and RSG (red crosses) algorithms perform when not applying any population parity constraint.

Each dot corresponds to the p -value of the KS test for a separate 25-precinct map. Under the assumption of independent sampled plans, if a simulation algorithm is successfully approximating the target distribution, these dots should fall roughly on the 45 degree line.

It is clear from this validation test that the RSG algorithm consistently fails to obtain a representative sample of the target distribution. That the red dots are concentrated near the bottom of the graph indicates that the KS p -value for the RSG algorithm is near zero for nearly every map tested. When population parity constraints of 20% and 10% are applied, the RSG algorithm continues to perform poorly compared to the MCMC algorithm. By using a soft constraint based on the Gibbs distribution, we allow the Markov chain to traverse from one valid plan to another through intermediate plans that may not satisfy the desired parity constraint. We find that although imperfect, the MCMC algorithm works much better than RSG algorithm.

4.2 Validation through Independent Sampling

Next, we conduct larger-scale validation studies by leveraging the fact that the `enumpart` algorithm can independently sample from the population of all possible redistricting plans. This feature allows us to scale up our validation studies further by avoiding for larger maps the computationally intensive task of writing to the hard disk all possible redistricting plans, which exponentially increases as the map size gets larger or as we try to partition a map into more districts. We independently sample a large number of redistricting plans and compare them against the samples obtained from simulation methods. Below, we present two validation studies. The first study uses the actual Congressional district maps from Iowa, where by law redistricting is done using 99 counties. The second study is based on a new 250-precinct validation map obtained from the Florida map.

4.2.1 The Iowa Congressional District Map

We first analyze a new validation dataset constructed on the redistricting map from the state of Iowa. In Iowa, redistricting is conducted using a total of 99 counties instead of census blocks to piece together districts, in order to avoid splitting county boundaries in line with the Iowa State Constitution.⁷ As a result, the Iowa redistricting problem is more manageable than other states.

The left plot of Figure 10 shows the Iowa map, where the shading indicates the population of each county. In 2016, Republicans won three districts while Democrats won one district, while in 2018, Democrats won three districts and the Republicans held only one. We use the `enumpart`

⁷Article 3, Section 37 of the Iowa State Constitution states “When a congressional, senatorial or representative district shall be composed of two or more counties, it shall not be entirely separated by any county belonging to another district; and no county shall be divided in forming a congressional, senatorial, or representative district.”

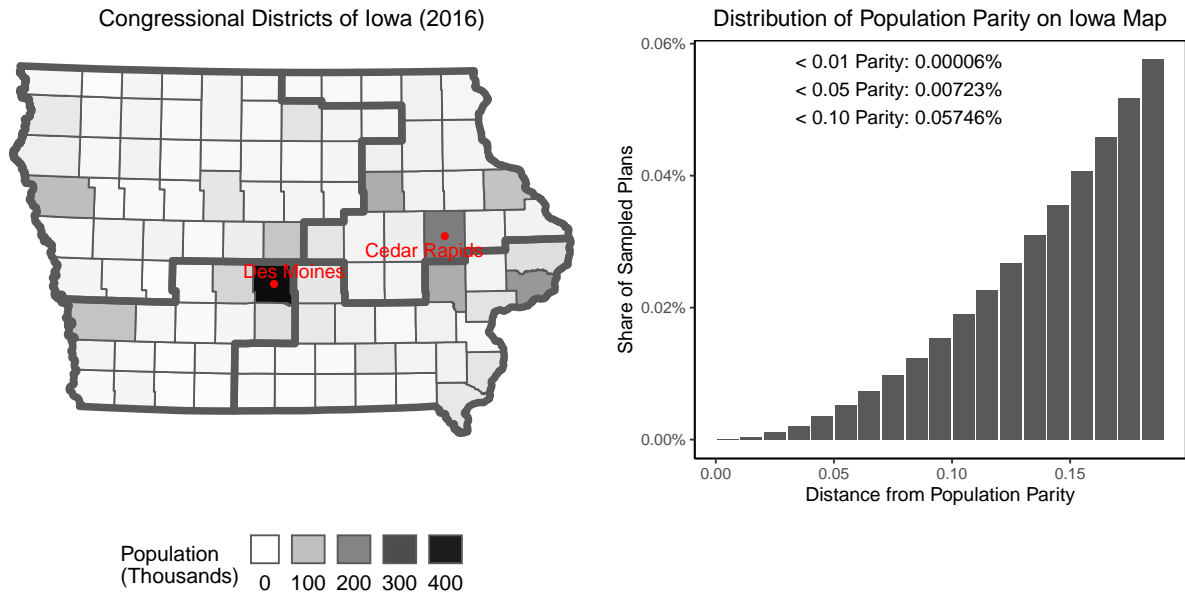


Figure 10: Iowa’s 2016 Congressional Districts and the Histogram of a Random Sample of Redistricting Plans under Various Population Parity Constraints. The underlying data is the Iowa county map, for which the `enumpart` algorithm generated an independent and uniform random sample of 500 million partitions of the map into four contiguous districts. In the histogram, each bar represents the number of redistricting plans that fall within the 1 percentage point range of a certain population parity, i.e., $[0, 0.01)$, $[0.01, 0.02)$, ..., $[0.19, 0.20)$. There are 36,131 valid plans when applying a 5% population parity constraint, and only 300 valid plans when applying a 1% population parity constraint.

algorithm to independently and uniformly sample 500 million contiguous partitions of this map into four districts. This number is minuscule relative to the total number of valid partitions of the map into four districts, of which there are approximately 10^{24} , but still is more than enough to use it as the target distribution. We note that while it took around 36 hours to sample 500 million partitions on the aforementioned cluster computer using significant parallelization, building the ZDD for this map took less than half a second on our MacBook Pro laptop mentioned earlier. Nearly all of the runtime of the enumeration was spent writing the solutions to harddisk.

The histogram in Figure 10 shows the share of the sampled redistricting plans that satisfy the deviation from population parity up to 20 percentage points. Of the 500 million plans we have randomly sampled, only 300, or less than 0.00006%, satisfy a 1% population parity constraint, illustrating the sheer scale of the redistricting problem and how much the population equality constraint alone shrinks the total solution space of valid redistricting plans. There are 36,131 plans, or less than 0.001%, satisfying a 5% population parity constraint, which is still a minuscule share compared to the total number of enumerated plans.

Figure 11 shows the performance of the MCMC and RSG simulation methods on the state-




Figure 11: A Validation Study, Uniformly Sampling from the Population of all Partitions of the Iowa Map into Four Districts. The underlying data is Iowa’s county map in the left plot of Figure 10, which is partitioned into four congressional districts. As in the previous validation exercises, the Markov chain Monte Carlo (MCMC) method (solid black line) is able to approximate the independently and uniformly sampled target distribution, while the random-seed-and-grow (RSG) method (red dashed line) performs poorly.

sized redistricting problem for Iowa. The solid grey density shows the distribution of the Republican dissimilarity index based on the independently and uniformly sampled set of 500 million redistricting plans. The red dashed lines show the distribution of the Republican dissimilarity index on plans sampled by the RSG algorithm, while the solid black lines shows the distribution for plans sampled by the MCMC algorithm. As in the previous validation exercise, where we impose 5% and 1% population parity constraints, we specify a target distribution of plans using the Gibbs distribution. Here, we set the temperature parameter $\beta_p = 25$ for the 5% parity constraint, and $\beta_p = 50$ for the 1% parity constraint, which we selected after initial tuning. After discarding plans not satisfying the constraint and then reweighting, we ended up with 629,729 plans for the 5% parity constraint, and 93,046 plans for the 1% parity constraint. The RSG algorithm was run for 2 million independent draws, while the MCMC algorithms were run for 250,000 iterations and initializing 8 chains for each algorithm. The chains were run without a burn-in period, and the Gelman-Rubin diagnostic suggested that the Markov chains had converged after at most 30,000 iterations.

As with the previous validation test, the MCMC algorithm outperforms the RSG algorithm across all levels of the population parity constraint. When no equal population constraint is applied or a 5% population parity constraint is applied, the MCMC algorithm samples from the target distribution nearly perfectly. Even with the 1% parity map, where there are only 300 valid plans in the target distribution, the MCMC algorithm approximates the target distribution reasonably well, missing by

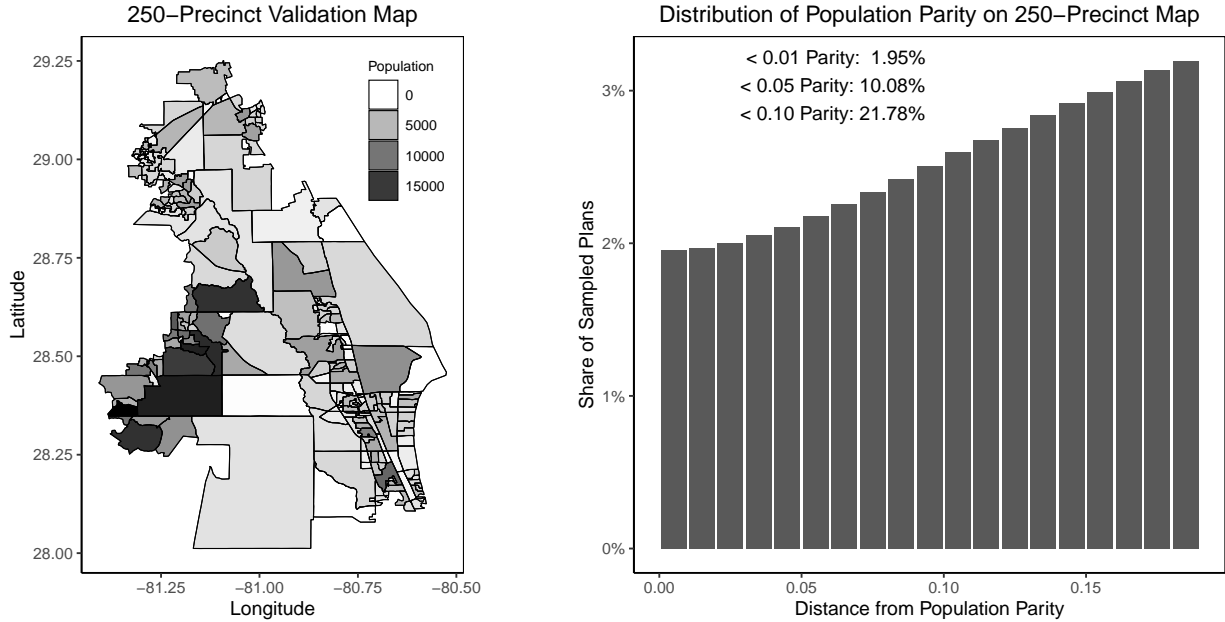


Figure 12: A New 250-Precinct Validation Map and the Histogram of Redistricting Plans under Various Population Parity Constraints. The underlying data is a 250-precinct contiguous subset of the Florida precinct map, for which the `enumpart` algorithm generated an independent and uniform random sample of 100 million partitions of the map into two contiguous districts. In the histogram, each bar represents the number of redistricting plans that fall within the 1 percentage point range of a certain population parity, i.e., $[0, 0.01)$, $[0.01, 0.02)$, ..., $[0.19, 0.20)$. There are 10,082,542 valid plans when applying a 5% population parity constraint, and 1,953,736 valid plans when applying a 1% population parity constraint.

only slightly in portions of the distribution. In contrast, at all levels of population parity, the RSG algorithm is unable to draw a representative sample of plans from the target distribution.

4.2.2 A New 250-Precinct Validation Map

Next, we present the results of validation tests that use a new, 250-precinct validation map, which is constructed from a contiguous subset of the 2008 Florida precinct map. As with the previous validation exercise, we use the `enumpart` algorithm to independently sample 100 million partitions of this map into two districts. This is still a minuscule number of plans relative to about 5×10^{39} possible partitions of this map into two districts. However, given the ability of the `enumpart` algorithm to independently sample plans using the ZDD, we are able to approximate an accurate target distribution arbitrarily well.

The left plot of Figure 12 shows the validation map, where the shading indicates the population of each of the precincts. Unlike the Iowa map, this map has geographical units of various sizes. This validation map also has a slightly larger frontier size (maximum frontier of 14) than that of the Iowa map (maximum frontier of 11), making it more likely to run out of memory due to the size of ZDD




Figure 13: A Validation Study Enumerating all Partitions of a 250-Precinct Map into Two Districts. The underlying data is the 250-precinct contiguous subset introduced in the left plot of Figure 12. As in the previous validation exercises, the Markov chain Monte Carlo (MCMC) method (solid black line) is able to approximate the target distribution based on the independent and uniform sampling, while the random-seed-and-grow (RSG) method (red dashed line) performs poorly.

and thereby also increasing computational time. The histogram on the right gives the distribution of population parity distance among the sampled plans, through 20% parity. Of the sampled plans, 1.95% (1.95 million plans) satisfy the 1% population parity constraint, while 21.8% of the sampled plans (21.8 million plans) satisfy the 10% population parity constraint.

We sample 4 million plans for each population parity level using the MCMC and RSG algorithms. For the MCMC algorithm, we initialized 8 chains running for 500,000 iterations each, and where a population parity constraint is imposed, we specify the target distribution of plans using the Gibbs distribution. We set the temperature parameter $\beta_p = 25$ for sampling plans within 5% of parity, and $\beta_p = 50$ when sampling plans within 1% of parity. After discarding invalid plans and reweighting, these parameter settings yielded 3,088,086 plans satisfying the 5% parity constraint, and 1,881,043 plans satisfying the 1% parity constraint. All eight chains were run without burn-in, and the Gelman-Rubin convergence diagnostic suggested the chains had converged after approximately 75,000 iterations.

Results for the validation test using the 250-precinct validation map are shown in Figure 13. As with the previous validation exercise, the solid grey density shows the target distribution of the Republican dissimilarity index on the 100 million plans sampled by the `enumpart` algorithm, while the red dashed lines show the distribution of the Republican dissimilarity index for the RSG algorithm. Finally, the solid black lines show the distribution of the Republican dissimilarity index for the MCMC algorithm. Across all levels of population parity, including the 1% constraint, the MCMC

algorithm is able to successfully sample from the target distribution and return a representative sample of redistricting plans. In contrast, where no population parity constraint is applied or where a 5% parity constraint is applied, the RSG algorithm is not able to sample from the target distribution with any accuracy. While it performs somewhat better on the 1% constraint, it is still biased towards plans with higher values on the dissimilarity index, and fails to capture the bimodality of the target distribution.

5 Concluding Remarks

More than a half century after scholars began to consider automated redistricting, legislatures and courts are increasingly relying on computational methods to generate redistricting plans and determine their legality. Unfortunately, despite the growing popularity of simulation methods in the recent redistricting cases, there exists little empirical evidence that these methods can in practice generate a representative sample of all possible redistricting maps under the statutory guidelines and requirements.

We believe that the scientific community has an obligation to empirically validate the accuracy of these methods. In this paper, we show how to conduct empirical validation studies by utilizing a recently developed computational method that enables the enumeration and independent sampling of all possible redistricting plans for maps with a couple of hundred geographical units. We make these validation maps publicly available, and implement our methodology as part of an open source software package, `redist`. These resources should facilitate researchers' efforts to evaluate the performance of existing and new methods in realistic settings.

Indeed, much work remains to be done in order to understand the conditions under which a specific simulation method do and does not perform well. A real-world redistricting process is complex. Distinct geographical features and diverse legal requirements play important roles in each state. It is far from clear how these factors interact with different simulation methods. Future work should address these issues using the data from various states.

It is also important to further improve the capabilities of the `enumpart` algorithm and of the MCMC algorithm. The maximum frontier size of our largest validation maps, which predicts the computational difficulty for the `enumpart` algorithm, is 14, which is far less than that of other states. For example, the maximum frontier size for New Hampshire (2 districts, 327 precincts) and Wisconsin (8 districts, 6895 precincts) are 21 and 84, respectively. These are much more challenging redistricting problems than the validation studies presented in this paper.

As the 2020 Census passes, lawsuits challenging proposed redistricting plans will inevitably be brought to court, and simulation evidence will be used to challenge and defend those plans. Thus, it is necessary that the empirical performance of these methods be rigorously evaluated. This paper introduces what we hope will be the first of many future complementary validation tests used to ensure that this evidence is of the highest possible quality according to scientific standards.

A Appendix




Figure 14: Quantile-Quantile Plot of p -values based on the Kolmogorov-Smirnov (KS) Tests of Distributional Equality between the Enumerated and Simulated Plans across 200 Validation Maps and under Different Population Parity Constraints. Each dot represents the p -value from a KS test comparing the empirical distribution of the Republican dissimilarity index from the simulated and enumerated redistricting plans. Under independent and uniform sampling, we expect the dots to fall on the 45-degree line. The MCMC algorithm (black dots), although imperfect, significantly outperforms the RSG algorithm (red crosses).

Figure 14 provides comparison to Figure 9. In Figure 9, thinning for the MCMC runs was set to 500. For this run, thinning for the MCMC runs was set to 1000. We find no significant difference between the two values. Thinning at 500 should then be sufficient and more efficient for most cases.

References

- Altman, M. (1997). The computational complexity of automated redistricting: Is automation the answer. *Rutgers Computer & Technology Law Journal* **23**, 81–142.
- Altman, M. and McDonald, M. P. (2011). BARD: Better automated redistricting. *Journal of Statistical Software* **42**, 4, 1–28.
- Carter, D., Herschlag, G., Hunter, Z., and Mattingly, J. (2019). A merge-split proposal for reversible monte carlo markov chain sampling of redistricting plans. Tech. rep., arXiv:1911.01503.
- Chen, J. and Rodden, J. (2013). Unintentional gerrymandering: Political geography and electoral bias in legislatures. *Quarterly Journal of Political Science* **8**, 239–269.
- Chikina, M., Frieze, A., and Pegden, W. (2017). Assessing significance in a Markov chain without mixing. *Proceedings of the National Academy of Sciences of the United States of America* **114**, 11, 2860–2864.
- Cirincione, C., Darling, T. A., and O’Rourke, T. G. (2000). Assessing South Carolina’s 1990s congressional districting. *Political Geography* **19**, 189–211.
- DeFord, D., Duchin, M., and Solomon, J. (2019). Recombination: A family of Markov chains for redistricting. Tech. rep., arXiv:1911.05725.
- Fifield, B., Higgins, M., Imai, K., and Tarr, A. (2014). A new automated redistricting simulator using Markov chain Monte Carlo. Tech. rep., Presented at the 31th Annual Meeting of the Society for Political Methodology. University of Georgia.
- Fifield, B., Higgins, M., Imai, K., and Tarr, A. (2020a). Automated redistricting simulator using Markov chain Monte Carlo. *Journal of Computational and Graphical Statistics* Forthcoming.
- Fifield, B., Imai, K., Kawahara, J., and Kenny, C. T. (2020b). Replication data for: The essential role of empirical validation in legislative redistricting simulation. <https://doi.org/10.7910/DVN/NH4CRS>.
- Fifield, B., Tarr, A., and Imai, K. (2015). *redist*: Markov chain monte carlo methods for redistricting simulation. available at the Comprehensive R Archive Network (CRAN). <https://CRAN.R-project.org/package=redist>.

- Fryer, R. G. and Holden, R. (2011). Measuring the compactness of political districting plans. *The Journal of Law and Economics* **54**, 3, 493–535.
- Gelman, A. and Rubin, D. B. (1992). Inference from iterative simulations using multiple sequences (with discussion). *Statistical Science* **7**, 4, 457–472.
- Herschlag, G., Ravier, R., and Mattingly, J. C. (2017). Evaluating partisan gerrymandering in Wisconsin. Tech. rep., Department of Mathematics, Duke University.
- Hess, S. W., Weaver, J. B., Siegfeldt, H. J., Whelan, J. N., and Zitlau, P. A. (1965). Nonpartisan political redistricting by computer. *Operations Research* **13**, 6, 998–1006.
- Katz, J., King, G., and Rosenblatt, E. (2019). Theoretical foundations and empirical evaluations of partisan fairness in district-based democracies. *American Political Science Review* Forthcoming.
- Kawahara, J., Horiyama, T., Hotta, K., and Minato, S. (2017). Generating all patterns of graph partitions within a disparity bound. In *Proceedings of the 11th International Conference and Workshops on Algorithms and Computation*, vol. 10167, 119–131.
- Knuth, D. E. (2011). *The Art of Computer Programming: Combinatorial Algorithms, Part I*, vol. 4A. Addison-Wesley, New York.
- Liu, Y. Y., Tam Cho, W. K., and Wang, S. (2016). PEAR: a massively parallel evolutionary computation approach for political redistricting optimization and analysis. *Swarm and Evolutionary Computation* **30**, 78–92.
- Magleby, D. and Mosesson, D. (2018). A new approach for developing neutral redistricting plans. *Political Analysis* **26**, 2, 147–167.
- Massey, D. S. and Denton, N. A. (1988). The dimensions of residential segregation. *Social Forces* **67**, 2, 281–315.
- Mattingly, J. C. and Vaughn, C. (2014). Redistricting and the will of the people. Tech. rep., Department of Mathematics, Duke University.
- McCarty, N., Poole, K. T., and Rosenthal, H. (2009). Does gerrymandering cause polarization? *American Journal of Political Science* **53**, 3, 666–680.
- Mehrotra, A., Johnson, E., and Nemhauser, G. L. (1998). An optimization based heuristic for political districting. *Management Science* **44**, 8, 1100–1114.

- Minato, S. (1993). Zero-suppressed BDDs for set manipulation in combinatorial problems. In *Proceedings of the 30th ACM/IEEE Design Automation Conference*, 272–277.
- Nagel, S. S. (1965). Simplified bipartisan computer redistricting. *Stanford Law Journal* **17**, 5, 863–899.
- Rubin, D. B. (1987). Comment: A noniterative sampling/importance resampling alternative to the data augmentation algorithm for creating a few imputation when fractions of missing information are modest:the SIR algorithm. *Journal of the American Statistical Association* **82**, 398, 543–546.
- Vickrey, W. (1961). On the prevention of gerrymandering. *Political Science Quarterly* **76**, 1, 105–110.
- Weaver, J. B. and Hess, S. W. (1963). A procedure for nonpartisan districting: Development of computer techniques. *Yale Law Journal* **73**, 2, 288–308.