Èdouard Lucas:

The theory of recurrent sequences is an inexhaustible mine which contains all the properties of numbers; by calculating the successive terms of such sequences, decomposing them into their prime factors and seeking out by experimentation the laws of appearance and reproduction of the prime numbers, one can advance in a systematic manner the study of the properties of numbers and their application to all branches of mathematics.

Preview

Mathematical Challenges in Neutral Redistricting

Daryl DeFord

MIT – CSAIL
Geometric Data Processing Group

Math Department Colloquium
Yale University
August 7, 2019

Outline

1. Introduction
2. Political Redistricting Problems
3. Total Variation Isoperimetric Profiles
4. Hardness of Partition Sampling
5. Conclusion

Key Collaborators

• Prof. Moon Duchin	**Tufts Math**
• Prof. Justin Solomon	**MIT CSAIL**
• Lorenzo Najt	**University of Wisconsin Math**
• Hugo Lavenant	**Universitè Paris–Sud Math**
• Zachary Schutzman	**University of Pennsylvania CIS**

Key Collaborators

- Prof. Moon Duchin — **Tufts Math**
- Prof. Justin Solomon — **MIT CSAIL**
- Lorenzo Najt — **University of Wisconsin Math**
- Hugo Lavenant — **Universitè Paris–Sud Math**
- Zachary Schutzman — **University of Pennsylvania CIS**

- Resources:
  - mggg.org
  - github.com/{gerrymandr,mggg,mggg-states,vrdi}
  - tinyurl.com/gerryprojects

Main Theorems

Theorem (D., Lavenant, Schutzman, Solomon (2018))

The Total Variation relaxation of the isoperimetric profile satisfies an isoperimetric inequality, is the lower convex envelope of the regular profile, and admits a family of distinguished solutions that extend the Cheeger set.

Theorem (D., Najt, Solomon (2019))

If there is a polynomial time algorithm to sample uniformly from the space of connected k partitions on the class of connected planar graphs then $RP = NP$.

MORAL #1:

MORAL #1:

Computational Redistricting is NOT a solved problem!

MORAL #2:

MORAL #2:

Computational Redistricting is NOT a solved problem!

MORAL #2:

Computational Redistricting is NOT a solved problem!

More Background:
Computational Approaches for Political Redistricting
Introduction to Discrete MCMC for Redistricting (with Scrabble)

Research Projects:
tinyurl.com/gerryprojects

Research Papers:
mggg.org/work

Political Partitioning

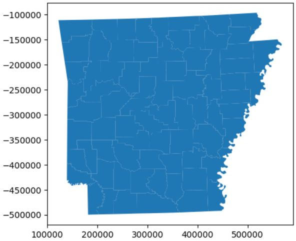

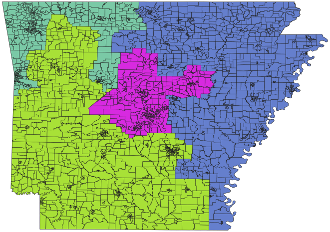

Permissible Districting Plans

We want to partition a given geography (graph), at a given scale, into k pieces, satisfying some constraints:

- Contiguity
- Population Balance
- Compactness
- Communities of Interest
- Municipal Boundaries
- Competitiveness/Responsiveness
- Partisan Symmetry
- Incumbency Protection
- ...

Permissible Districting Plans

We want to partition a given geography (graph), at a given scale, into k pieces, satisfying some constraints:

- **Contiguity**
- **Population Balance**
- Compactness
- Communities of Interest
- Municipal Boundaries
- Competitiveness/Responsiveness
- Partisan Symmetry
- Incumbency Protection
- ...

Discrete Districting

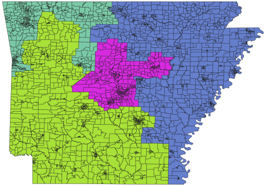

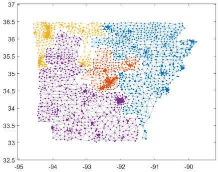

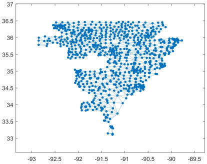

(c) District # 1

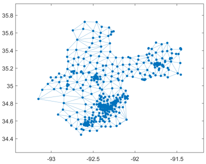

(d) District # 2

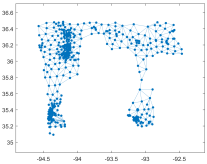

(e) District # 3

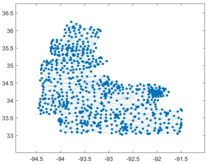

(f) District # 4

Mathematical Formulation

Given a (connected) graph $G = (V, E)$:

- A k-**partition** $P = \{V_1, V_2, \ldots, V_k\}$ of G is a collection of disjoint subsets $V_i \subseteq V$ whose union is V.
- A partition P is **connected** if the subgraph induced by V_i is connected for all i.
- The **cut edges** of P are the edges (u, w) for which $u \in V_i$, $w \in V_j$, and $i \neq j$
- A partition P is ε-**balanced** if $\mu(1 - \varepsilon) \leq |V_i| \leq \mu(1 + \varepsilon)$ for all i where μ is the mean of the $|V_i|$'s
- An **equi–partition** is a 0-balanced partition

Which graphs?

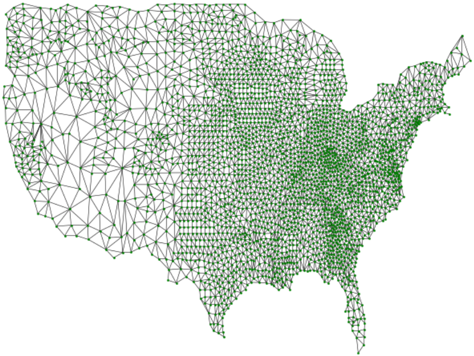

Census Dual Graphs

IN Dual Graphs

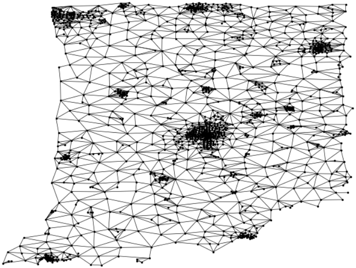

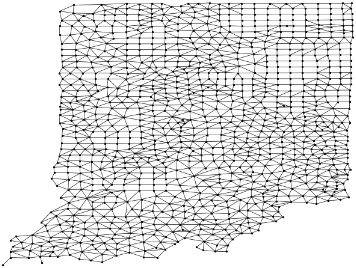

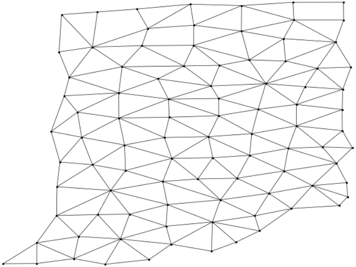

IA Dual Graphs

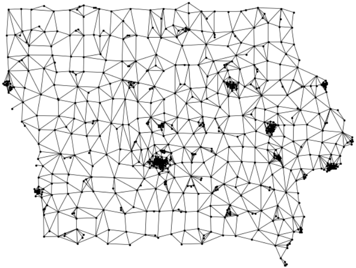

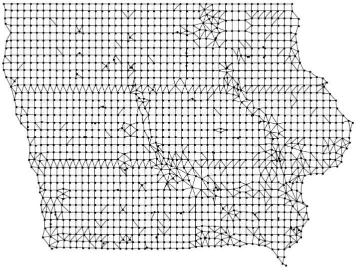

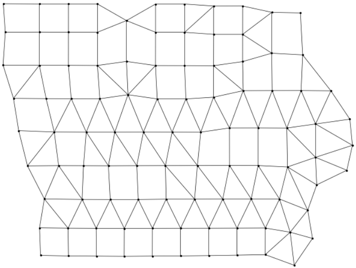

KS Dual Graphs

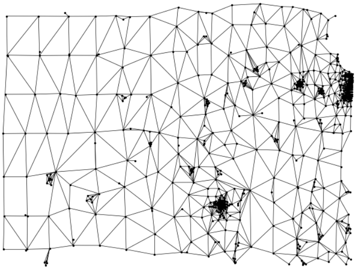

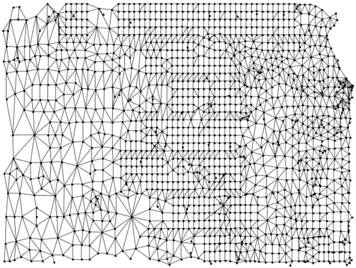

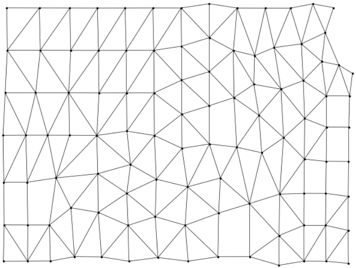

KY Dual Graphs

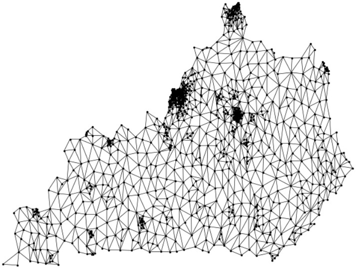

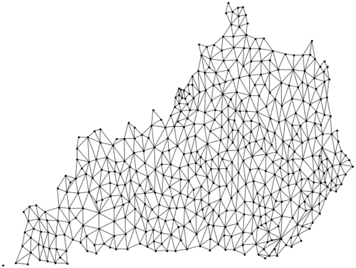

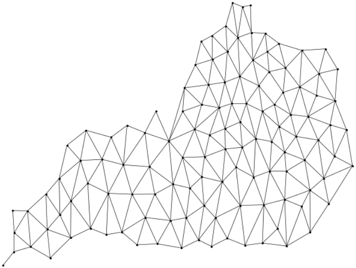

LA Dual Graphs

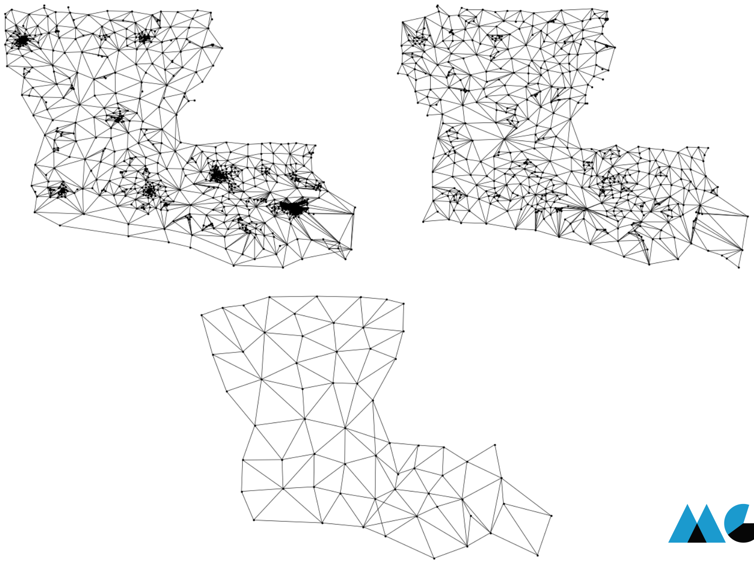

Electoral Data

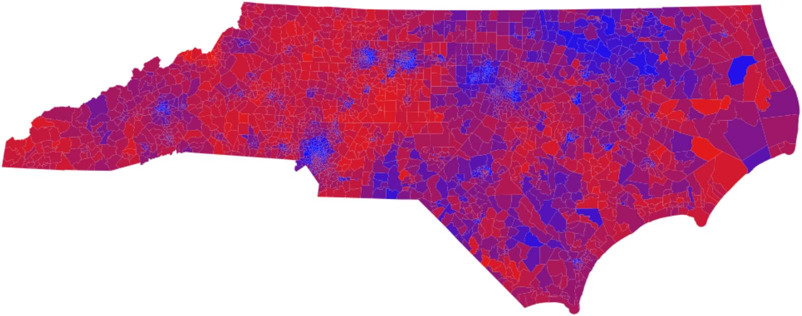

MORAL:

Computational Redistricting is NOT a solved problem!

Example: Iowa

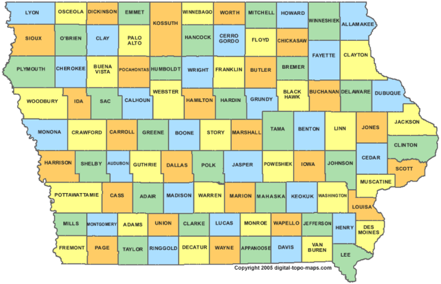

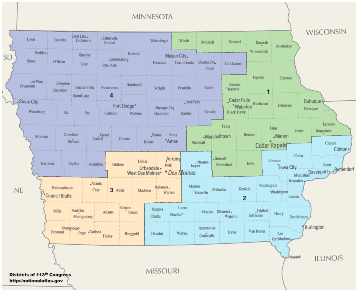

Example: Iowa

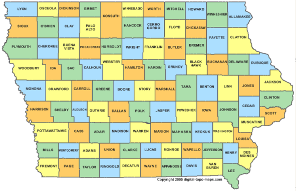

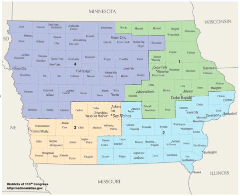

- 4 Congressional Districts, 100 House Districts, 50 Senate Districts
- House districts nest into Senate districts
- Congressional districts made out of counties
- Independent committee with legislative approval
- No partisan data allowed

Example: Pennsylvania

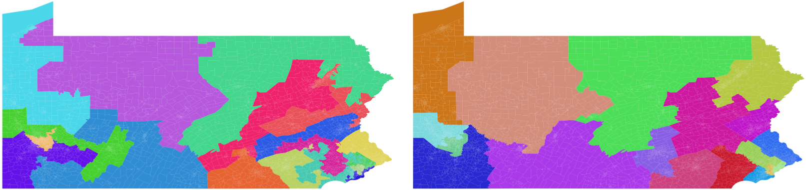

- 18 Congressional Districts, 203 House Districts, 50 Senate Districts
- Zero–balanced population
- Legislature draws congressional districts committee draws legislative districts
- Partisan behavior allowed

MORAL:

Computational Redistricting is NOT a solved problem!

What can go wrong?

What can go wrong?

- Geometry

What can go wrong?

- Geometry
- Partisan imbalance

What can go wrong?

- Geometry
- Partisan imbalance
- Outlier analysis

Ugly Shapes

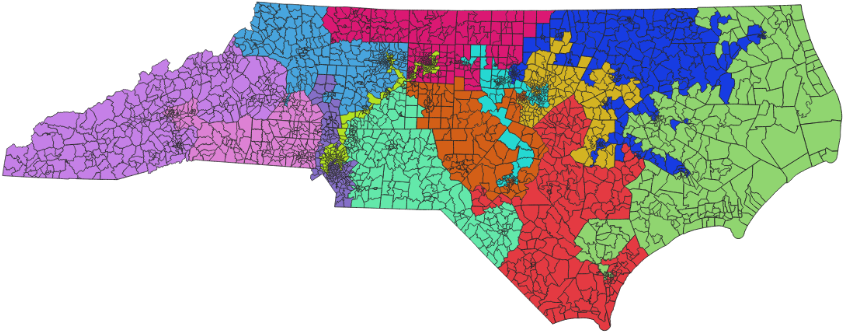

Ugly Shapes

(a) NC12 #1

(b) NC12 #2

(c) NC12 #12

Ugly Shapes

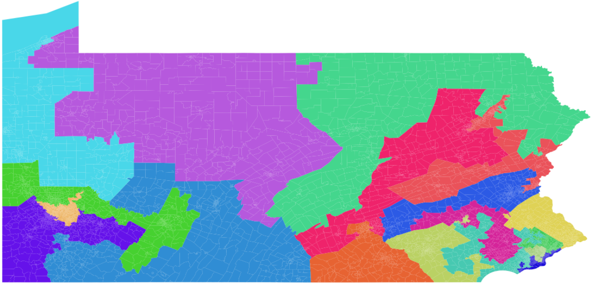

Measurement Problems

Theorem (Bar-Natan, Najt, and Schutzman 2019)

There is no local homeomorphism from the globe to the plane that preserves the rankings of your favorite compactness measure.

Problem (Barnes and Solomon 2018)

Compactness scores can be distorted by:

- *Map projection*
- *Data resolution*
- *State borders and coastline*
- *Topography*
- *...*

Partisan Imbalance

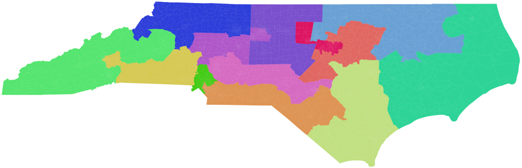

(a) NC16

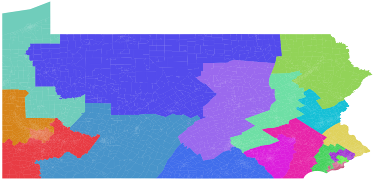

(b) PA TS-Proposed

Partisan Fairness Scores

Types of metrics:

- Mean-median
- Partisan Bias
- Partisan Gini
- Efficiency Gap

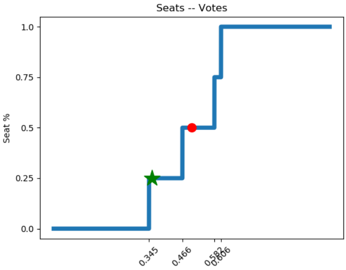

Outlier Examples: NC

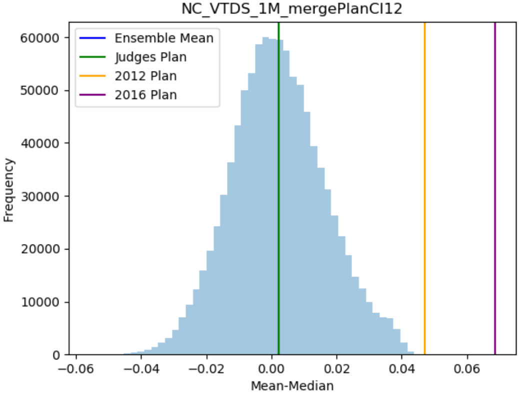

Outlier Examples: NC

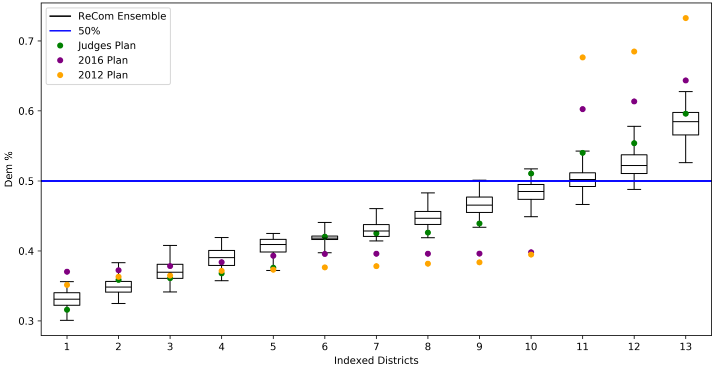

Baseline Example: VA

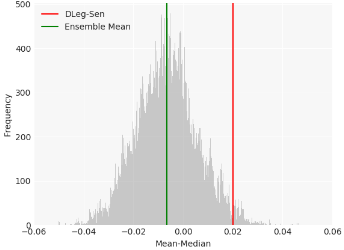

(a) Mean–Median

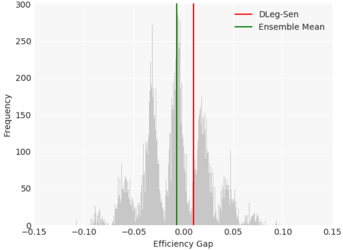

(b) Efficiency Gap

Baseline Example: PA

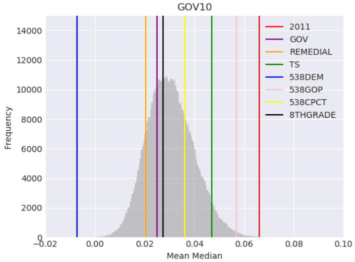

(a) Mean–Median

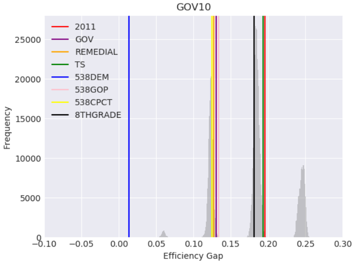

(b) Efficiency Gap

Reform Example: VA

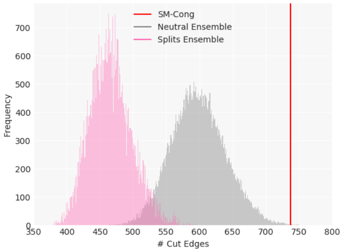

(a) Compactness

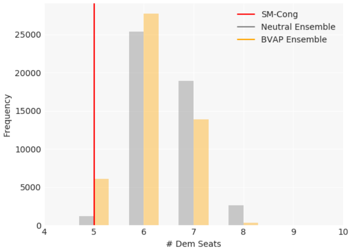

(b) Dem Seats

Which ensembles?

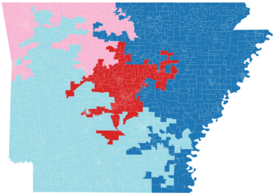

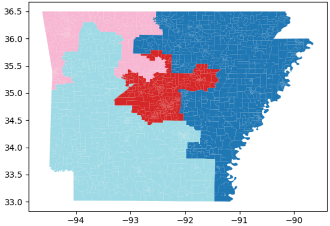

Redistricting Wrapup

- Shape based measures are hard to implement
- "Unfair" plans no longer need to appear ugly
- Partisan metrics are exploitable and vary in meaning from state to state
- Outlier approaches can be sensitive to choice of distribution
- Blindness isn't necessarily fairness

MORAL:

Computational Redistricting is NOT a solved problem!

Polsby–Popper

Theorem (Isoperimetry)

Let Ω be a bounded open subset of $\mathbb{R}^2$ with finite perimeter. Then:

$$4\pi A \leq P^2.$$

Definition (Polsby–Popper)

The Polsby–Popper score of a district is:

$$PP(\Omega) = \frac{4\pi A}{P^2}$$

Polsby–Popper

Theorem (Isoperimetry)

Let Ω be a bounded open subset of $\mathbb{R}^2$ with finite perimeter. Then:

$$4\pi A \leq P^2.$$

Definition (Polsby–Popper)

The Polsby–Popper score of a district is:

$$PP(\Omega) = \frac{4\pi A}{P^2}$$

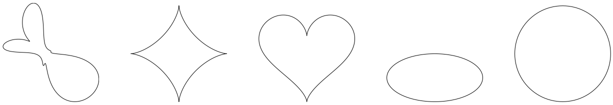

$\frac{4\pi A}{L^2} =$ 0.359 0.411 0.680 0.841 1.000

Boundary Perturbation

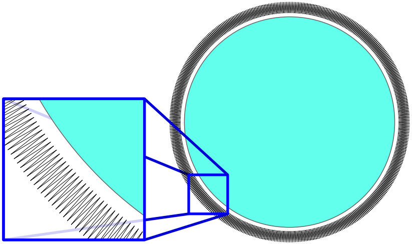

$$\frac{4\pi A}{L^2} = 0.004$$

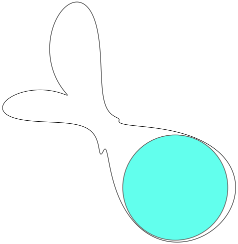

$$\frac{4\pi A}{L^2} = 0.359$$

Multiscale Desiderata

- Disambiguate different types of “badness”
- Stability under practical constraints
- Interpolate well–studied single measures
- Continuous and discrete versions
- Internal vs. external

Isoperimetric Profile

Definition (Isoperimetric Inequality)

Let $\Omega \subseteq \mathbb{R}^n$ to be a compact region whose boundary $\partial\Omega \subseteq \Omega$ is an $(n-1)$-dimensional hypersurface in $\mathbb{R}^n$

$$n \cdot \mathrm{vol}(\Omega)^{\frac{(n-1)}{n}} \cdot \mathrm{vol}(B(1, \mathbf{0}))^{\frac{1}{n}} \leq \mathrm{area}(\partial\Omega).$$

Isoperimetric Profile

Definition (Isoperimetric Inequality)

Let $\Omega \subseteq \mathbb{R}^n$ to be a compact region whose boundary $\partial\Omega \subseteq \Omega$ is an $(n-1)$-dimensional hypersurface in $\mathbb{R}^n$

$$n \cdot \mathrm{vol}(\Omega)^{\frac{(n-1)}{n}} \cdot \mathrm{vol}(B(1, \mathbf{0}))^{\frac{1}{n}} \leq \mathrm{area}(\partial\Omega).$$

Definition (Isoperimetric Profile)

With Ω as above and $t \in [0, \mathrm{vol}(\Omega)]$ we ask for the smallest surface area needed to enclose volume t completely within Ω:

$$I_\Omega(t) := \min\{\mathrm{area}(\partial\Sigma) : \Sigma \subseteq \Omega \text{ and } \mathrm{vol}(\Sigma) = t\}.$$

Geometric Properties

Theorem (Flores and Nardulli (2016) [1)

] Let M^n be a complete smooth Riemannian manifold with $\mathrm{Ric}_M \geq (n-1)k$, with $k \in \mathbb{R}$ and $V(B(p,1)) \geq v_0 > 0$. Then the isoperimetric profile is continuous on $[0, V(M)[$

[1] A. Flores and S. Nardulli: Continuity and differentiability properties of the isoperimetric profile in complete noncompact Riemannian manifolds with bounded geometry, https://arxiv.org/abs/1404.3245.

Geometric Properties

Theorem (Flores and Nardulli (2016) [1]

] Let M^n be a complete smooth Riemannian manifold with $\mathrm{Ric}_M \geq (n-1)k$, with $k \in \mathbb{R}$ and $V(B(p,1)) \geq v_0 > 0$. Then the isoperimetric profile is continuous on $[0, V(M)[$

Question

Identify a polynomial-time algorithm or NP-hardness result for computing isoperimetric profiles. The simplest open problem is computing the isoperimetric profile of a polygon in the plane $\mathbb{R}^2$.

[1] A. Flores and S. Nardulli: Continuity and differentiability properties of the isoperimetric profile in complete noncompact Riemannian manifolds with bounded geometry, https://arxiv.org/abs/1404.3245.

Total Variation

Definition (Three formulations of TV)

$$\mathrm{TV}[f] =$$

1. $$\sup\left\{\int_{\mathbb{R}^n} [f(x)\nabla\cdot\phi(x)]\,dx : \phi \in C_c^1(\mathbb{R}^n \to \mathbb{R}^n) \text{ and } \|\phi\|_\infty \leq 1\right\}$$
2. $$\int_{\mathbb{R}^n} \|\nabla f\|_2\,dx$$
3. $$\int_0^{+\infty} \mathrm{area}(\partial\{f \geq s\})ds$$

Perimeter as Total Variation

Definition

For a region $\Sigma \subseteq \mathbb{R}^n$, denote its *indicator function* $\mathbb{1}_\Sigma$ via

$$\mathbb{1}_\Sigma(x) := \begin{cases} 1 & \text{if } x \in \Sigma \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

Then, a consequence of the co-area formula is that

$$\text{area}(\partial\Sigma) = \text{TV}[\mathbb{1}_\Sigma]. \tag{2}$$

TV Relaxation

Definition (Isoperimetric Profile)

$$I_\Omega(t) = \begin{cases} \inf_{f \in L^1(\mathbb{R}^n)} & \mathrm{TV}[f] \\ \text{subject to} & \int_{\mathbb{R}^n} f(x)\,dx = t \\ & 0 \leq f \leq \mathbb{1}_\Omega \\ & f(x) \in \{0,1\}\ \forall x \in \mathbb{R}^n. \end{cases}$$

TV Relaxation

Definition (Isoperimetric Profile)

$$I_\Omega(t) = \begin{cases} \inf_{f \in L^1(\mathbb{R}^n)} & \mathrm{TV}[f] \\ \text{subject to} & \int_{\mathbb{R}^n} f(x)\,dx = t \\ & 0 \le f \le \mathbb{1}_\Omega \\ & f(x) \in \{0,1\}\ \forall x \in \mathbb{R}^n. \end{cases}$$

Definition (TV Profile)

$$I_\Omega^{\mathrm{TV}}(t) := \begin{cases} \min_{f \in L^1(\mathbb{R}^n)} & \mathrm{TV}[f] \\ \text{subject to} & \int_{\mathbb{R}^n} f(x)\,dx = t \\ & 0 \le f \le \mathbb{1}_\Omega. \end{cases}$$

Example: Circle

Proposition

For all (Ω, t), *we have* $I_\Omega^{\mathrm{TV}}(t) \leq I_\Omega(t)$.

Examples (Circle)

Suppose $\Omega \subset \mathbb{R}^2$ *is a circle of radius* R, *and take* $t = \pi r^2$ *for* $r \in (0, R)$. *In this case, by the isoperimetric inequality we know* $I_\Omega(t) = 2\pi r$. *But suppose we take* $f(x) \equiv \dfrac{r^2}{R^2}$. *By the co-area formula*

$$I_\Omega^{\mathrm{TV}}(t) \leq \mathrm{TV}[f] = 2\pi R \cdot \frac{r^2}{R^2} = 2\pi r \cdot \frac{r}{R} < I_\Omega(t).$$

Hence, our relaxation is not tight.

Isoperimetry and Convexity

Proposition (Isoperimetry)

Suppose $B \subset \mathbb{R}^n$ is a ball whose volume matches $\mathrm{vol}(\Omega)$. Then, for all $t \in [0, \mathrm{vol}(\Omega)]$, we have $I_B^{\mathrm{TV}}(t) \leq I_\Omega^{\mathrm{TV}}(t)$, and if the equality holds for some $t > 0$ then Ω is a ball.

Proposition (Convexity)

$I_\Omega^{\mathrm{TV}}(t)$ is a convex function of t.

Proposition (Convex Envelope)

The function I_Ω^{TV} is the lower convex envelope of I_Ω.

Dual Optimization

Dual Formulation:

$$I_\Omega^{TV}(t) = \begin{cases} \sup_{\phi \in C_c^1(\mathbb{R}^n \to \mathbb{R}^n), \lambda \in \mathbb{R}} & \lambda t - \int_\Omega \max(\lambda - \nabla \cdot \phi(x), 0)\, dx \\ \text{subject to} & \|\phi\|_\infty \leq 1 \end{cases}$$

Proof.

With the dual in hand, the convexity results follow from constructing an auxilliary function:

$$h(\lambda) = \inf_{\|\phi\|_\infty \leq 1} \int_\Omega \max(\lambda - \nabla \cdot \phi(x), 0)\, dx.$$

and computing some Legendre transforms. □

Minimizer Structure

Proposition (Distinguished Solutions)

There exists a family $(f_t)_{t\in[0,1]}$ *such that:*

- *For any* $t \in [0,1]$*, the function* $f_t \in L^1(\mathbb{R}^n)$ *satisfies* $0 \le f_t \le \mathbb{1}_\Omega$*,* $\int_{\mathbb{R}^n} f_t(x)\, dx = t$ *and* $\mathrm{TV}(f_t) = I_\Omega^{\mathrm{TV}}(t)$*.*
- *For any* $t \in [0,1]$*, there exist* $v_t \in (0,1)$ *such that* f_t *takes its values in* $\{0, v_t, 1\}$*.*
- *For a.e.* $x \in \Omega$*, the function* $t \to f_t(x)$ *is increasing.*

NC 12 # 9

NC 12 # 2

NC 12 # 12

Cheeger Sets

Definition (Cheeger Constant)

The Cheeger constant of Ω, denoted by $h_1(\Omega)$, is defined as

$$h_1(\Omega) := \inf_{\tilde{\Sigma} \subseteq \Omega} \frac{\operatorname{area}(\partial \tilde{\Sigma})}{\operatorname{vol}(\tilde{\Sigma})},$$

and a subset $\Sigma \subseteq \Omega$ such that $h_1(\Omega) = \frac{\operatorname{area}(\partial \Sigma)}{\operatorname{vol}(\Sigma)}$ is known as a Cheeger set of Ω.

Proposition (Small t)

Let Ω be compact, let $h_1(\Omega)$ be the Cheeger constant of Ω, and let Σ be a Cheeger set of Ω. Then for any $t \leq \operatorname{vol}(\Sigma)$, we have $I_\Omega^{\mathrm{TV}}(t) = h_1(\Omega)t$, and a solution f is given by $f := \frac{t}{\operatorname{vol}(\Sigma)} \cdot \mathbb{1}_\Sigma$.

Cheeger Proof

Proof.

We start with $\hat{f} = \frac{t}{\mathrm{vol}(\Sigma)} \cdot \mathbb{1}_\Sigma$ which satisfies the constraints of the problem defining $I_\Omega^{\mathrm{TV}}(t)$ as soon as $t \leq \mathrm{vol}(\Sigma)$, which ensures $0 \leq \hat{f} \leq \mathbb{1}_\Sigma \leq \mathbb{1}_\Omega$. Hence, $I_\Omega^{\mathrm{TV}}(t) \leq h_1(\Omega)t$. On the other hand, using the co-area formula, if f is any competitor then

$$\begin{aligned}
\mathrm{TV}(f) &= \int_0^{+\infty} \mathrm{area}(\partial\{f \geq s\}) ds \\
&= \int_0^{+\infty} \mathrm{vol}(\{f \geq s\}) \cdot \underbrace{\frac{\mathrm{area}(\partial\{f \geq s\})}{\mathrm{vol}(\{f \geq s\})}}_{\geq h_1(\Omega) \text{ by definition}} ds \\
&\geq h_1(\Omega) \int_0^{+\infty} \mathrm{vol}(\{f \geq s\}) ds = h_1(\Omega) \int_{\mathbb{R}^d} f(x) dx = h_1(\Omega)t.
\end{aligned}$$

Hence, for $t \leq \mathrm{vol}(C)$, we have $I_\Omega^{\mathrm{TV}}(t) = h_1(\Omega)t$. □

Synthetic Examples

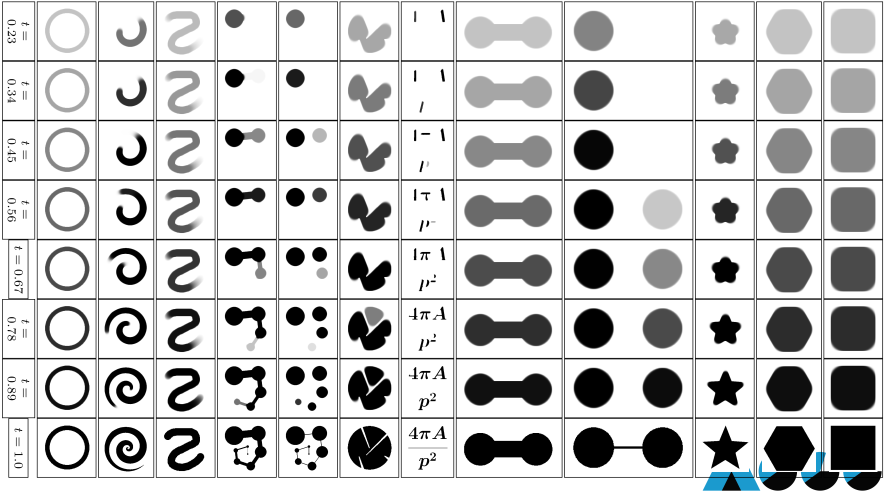

Synthetic Profiles

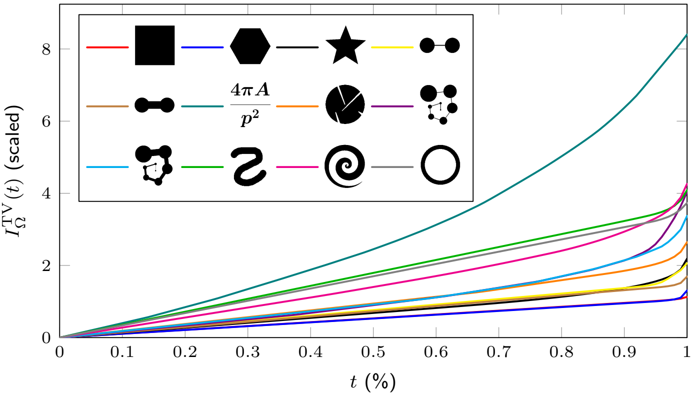

Synthetic Derivatives

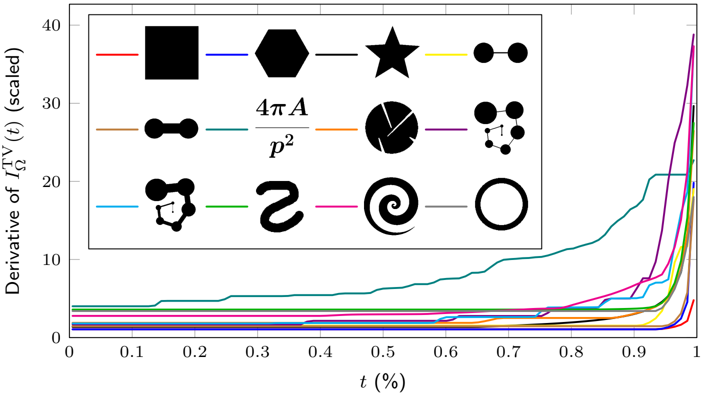

North Carolina

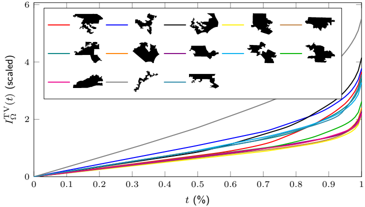

(a) 2011 districts

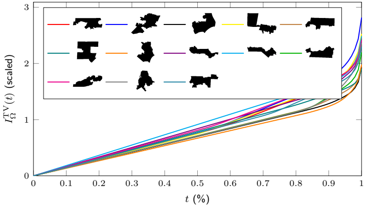

(b) 2016 districts

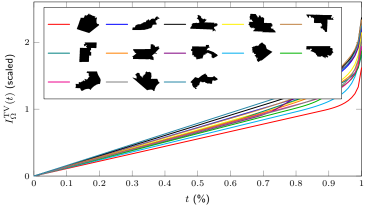

(c) Judges plan

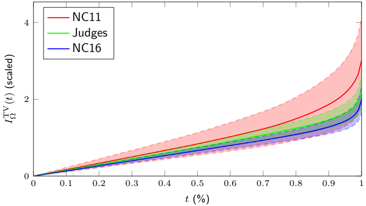

(d) Comparison

NC 2011 Districts

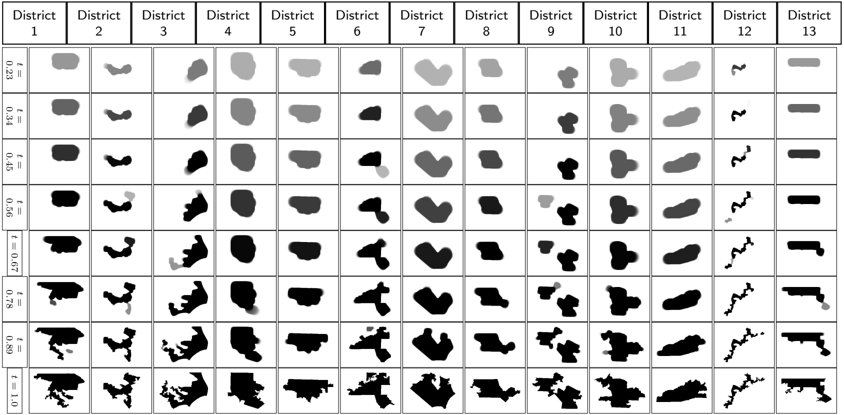

NC 2016 Districts

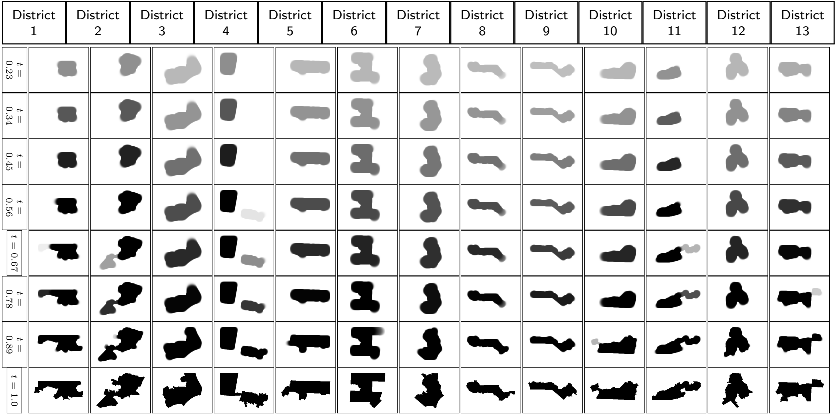

Judge's Plan

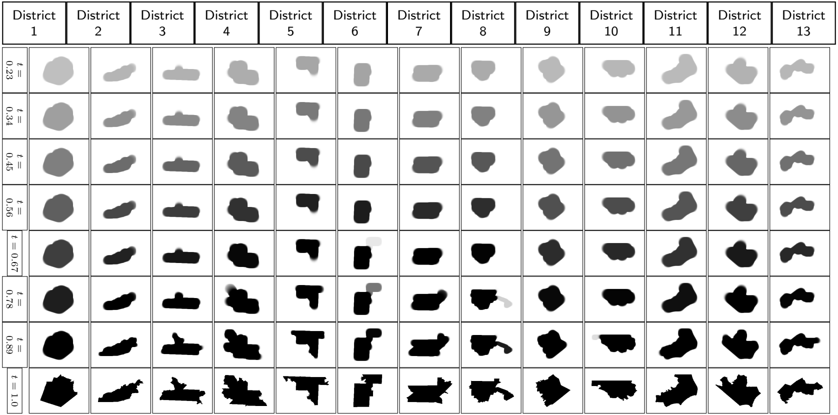
District 1
District 2
District 3
District 4
District 5
District 6
District 7
District 8
District 9
District 10
District 11
District 12
District 13
t = 0.23
t = 0.34
t = 0.45
t = 0.56
t = 0.67
t = 0.78
t = 0.89
t = 1.0

Higher Dimensions

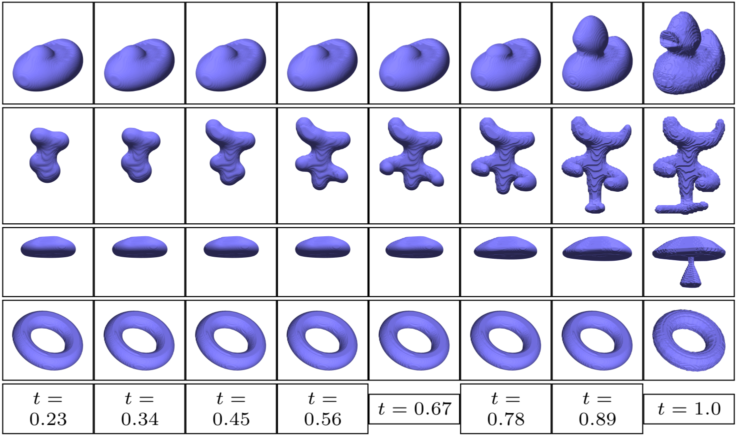

Other Formulations

Definition (Population Measure)

$$I^{\mathrm{TV}}_{\Omega,\rho}(t) := \begin{cases} \min_{f \in L^1(\mathbb{R}^n)} & \mathrm{TV}[f] \\ \text{subject to} & \int_{\mathbb{R}^n} f(x)\, d\rho(x) = t \\ & 0 \le f \le \mathbb{1}_\Omega . \end{cases}$$

Definition (Discrete)

$$I^{\mathrm{TV}}_{V_0}(t) := \begin{cases} \min_{f \in \mathbb{R}^V} & \sum_{(v,w)\in E} |f(v) - f(w)| \\ \text{subject to} & \sum_{v \in V_0} f(v) = t|V_0| \\ & f(v) = 0 \ \forall v \notin V_0 \\ & f(v) \in [0,1] \ \forall v \in V. \end{cases}$$

Synthetic Cities

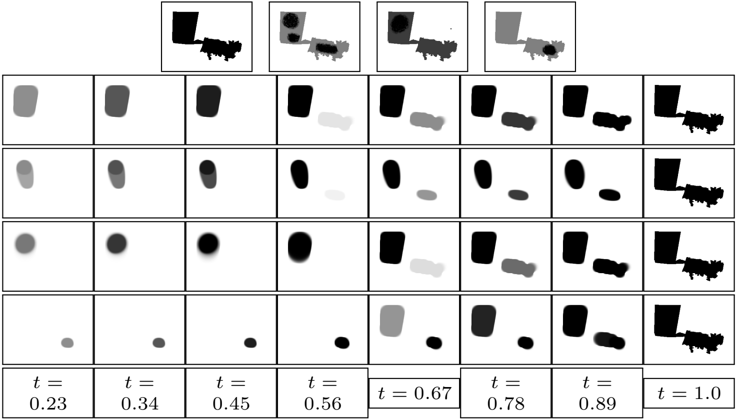

Discrete Animation

Multiscale Wrapup

Open questions:

- How much can we learn about the full profile from the relaxed version?
- Can the medial axis be computed from the TV-Profile?
- What is the right way to compare regions of the profiles?

- Mean curvature flow?
- Spectral Versions (i.e. how to make the heat kernel useful)
- Random walk versions (absorbing boundary nodes)
- Distance based measures
- ...

Which ensembles?

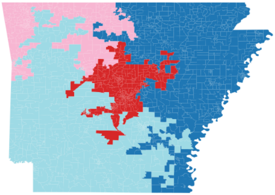

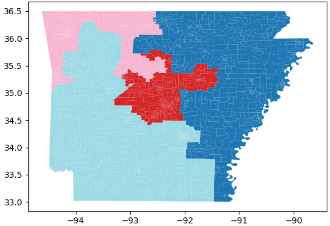

Ensembles in Practice

- The appeal of an ensemble method is that you get to control the input data very carefully
- However, just because a particular type of data was not considered doesn't mean that the outcome is necessarily "fair"
- There are lots of "random" methods for constructing districting plans
- Most don't offer any control over the distribution that you are drawing from

MCMC on partitions

1. Set constraints to define the state space
2. Start with an initial plan
3. Propose a modification
4. Verify that the modification satisfies the constraints
5. Accept using MH criterion
6. Repeat

MCMC on partitions

1. Set constraints to define the state space
2. Start with an initial plan
3. Propose a modification
4. Verify that the modification satisfies the constraints
5. Accept using MH criterion
6. Repeat

Why?

MCMC on partitions

1. Set constraints to define the state space
2. Start with an initial plan
3. Propose a modification
4. Verify that the modification satisfies the constraints
5. Accept using MH criterion
6. Repeat

Why?

- Control over sampling distribution and input data
- Possibility of local sampling
- Ergodic Theorem

MCMC on partitions

1. Set constraints to define the state space
2. Start with an initial plan
3. **Propose a modification**
4. Verify that the modification satisfies the constraints
5. Accept using MH criterion
6. Repeat

Why?

- Control over sampling distribution and input data
- Possibility of local sampling
- Ergodic Theorem

Single Edge Flip Proposals

1. Uniformly choose a cut edge
2. Change one of the incident node assignments to the other

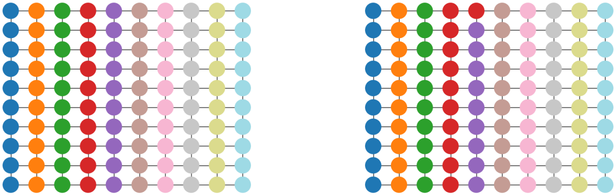

- Mattingly et al. (2017, 2018) Court cases in NC and WI.
- Pegden et al. Assessing significance in a Markov chain without mixing, PNAS, (2017). Court case in PA and NC.

Single Edge Ensembles

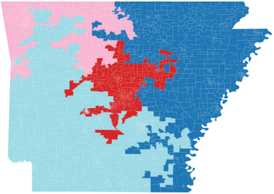

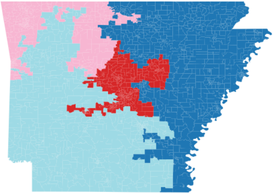

PA Single Edge Flip

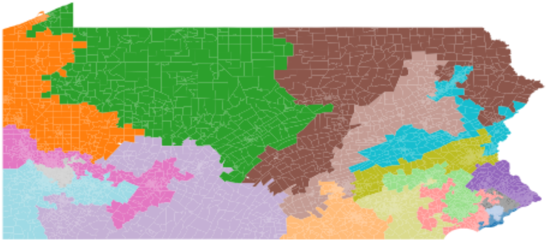

Constraints

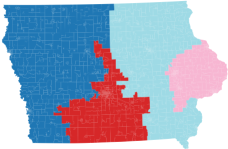

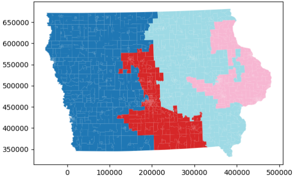

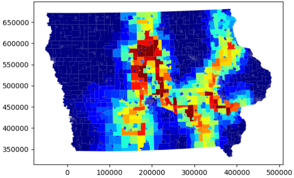

Annealing

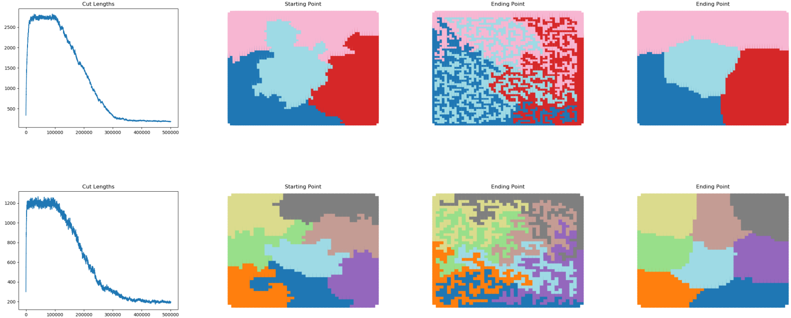
Cut Lengths
2500
2000
1500
1000
500
0
100000
200000
300000
400000
500000
Starting Point
Ending Point
Ending Point
Cut Lengths
1200
1000
800
600
400
200
0
100000
200000
300000
400000
500000
Starting Point
Ending Point
Ending Point

KY Annealing

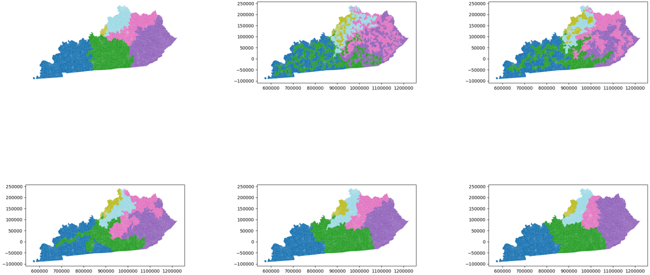

Sensitivity to Topology

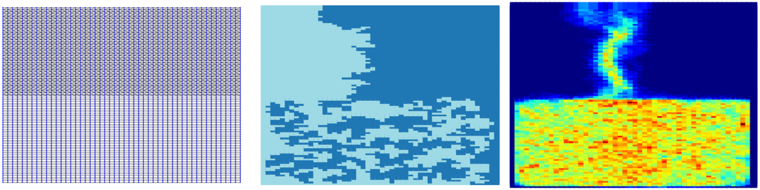

Uniform Sampling of Contiguous Partitions

Theorem (D., Najt, and Solomon 2019)

Suppose that $\mathscr{C}$ is the class of connected planar graphs and $k \geq 2$. If there is a polynomial time algorithm to sample uniformly from:

- *the connected k-partitions of graphs in $\mathscr{C}$,*
- *or the connected, 0-balanced k-partitions of graphs in $\mathscr{C}$.*

then $RP = NP$.

Uniform Sampling of Contiguous Partitions

Theorem (D., Najt, and Solomon 2019)

Suppose that $\mathscr{C}$ is the class of connected planar graphs and $k \geq 2$. If there is a polynomial time algorithm to sample uniformly from:

- *the connected k-partitions of graphs in $\mathscr{C}$,*
- *or the connected, 0-balanced k-partitions of graphs in $\mathscr{C}$.*

then $RP = NP$.

Remark

This theorem has various interesting extensions, including:

- *Connectivity constraints on $\mathscr{C}$*
- *Degree bounds*
- *Distributions proportional to cut length*
- *TV distribution approximation*

Stronger Version Example

Theorem (D., Najt, and Solomon 2019)

Let $\mathscr{C}$ be the class of cubic, planar 3-connected graphs, with face degree bounded by $C = 60$. Let $\mu_x(G)$ be the probability measure on $P_k(G)$ such that a partition P is drawn with probability proportional to $x^{cut(P)}$. Fix some $x > 1/\sqrt{2}$, $\epsilon > 0$ and $\alpha < 1$. Suppose that there was an algorithm to sample from $P_2^{\epsilon}(G)$ according to a distribution $\nu(G)$, such that $||\nu_G - \mu_x(G)||_{TV} < \alpha$, which runs polynomial time on all $G \in \mathscr{C}$. Then $RP = NP$.

Proof Outline Sketch

Following technique of Jerrum, Valiant, and Vazirani[1].

1. Show that uniformly sampling simple cycles is hard on some class $\mathscr{C}$
   1. Choose a gadget that respects $\mathscr{C}$ and allows us to concentrate probability on long cycles
   2. Count the proportion of cycles as a function of length
   3. Reduce to Hamiltonian path on the graph class
2. Show closure of class under planar dual
3. Identify partitions with cut edges $\mapsto$ simple cycles (via planar duality)
4. Conclude that sampling partitions would allow you to sample from cycles which would allow you to find Hamiltonian cycles

□

[1] M. Jerrum, L. Valiant, and V. Vazirani, Random generation of combinatorial structures from a uniform distribution, Theoretical Computer Science, 43 (1986), 169–188.

Proof Sketch – Planar 2–Partitions

Still following technique of Jerrum, Valiant, and Vazirani.

1. Let $\mathscr{C}$ be the planar connected graphs
   1. Replace the edges with chains of dipoles
   2. Hamiltonian hardness for $\mathscr{C}$ given by [1]
2. $\mathscr{C}$ closed under planar duals
3. Identify partitions with cut edges (via planar duality)

□

[1] M. Garey, D. Johnson, and R. Tarjan, The Planar Hamiltonian Circuit Problem is NP-Complete, SIAM Journal on Computing, 5, (1976), 704–714.

Slowly Mixing Graph Families

Theorem (D., Najt, and Solomon 2019)

Let G be any connected graph. Then let $G^{(d)}$ be the graph obtained by replacing each edge by a doubled d-star. Then the flip walk on partitions of family of graphs $G^{(d)}_{d \geq 1}$ is slowly mixing, in the sense the Cheeger constant is decaying exponentially fast. More specifically:

$$H(Partition\ Graph(G^{(d)})) = O(2^{-d})$$

Slowly Mixing Graph Families

Theorem (D., Najt, and Solomon 2019)

Let G be any connected graph. Then let $G^{(d)}$ be the graph obtained by replacing each edge by a doubled d-star. Then the flip walk on partitions of family of graphs $G^{(d)}_{d \geq 1}$ is slowly mixing, in the sense the Cheeger constant is decaying exponentially fast. More specifically:

$$H(Partition\ Graph(G^{(d)})) = O(2^{-d})$$

Remark

There are many similar constructions that give rise to equivalent mixing results.

Slow Mixing Example

Slow Mixing Example

MORAL:

Computational Redistricting is NOT a solved problem!

Tree based methods

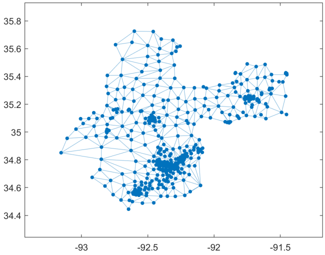

(a) District

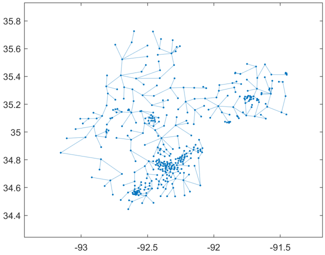

(b) Spanning Tree

Tree Seeds Ensemble

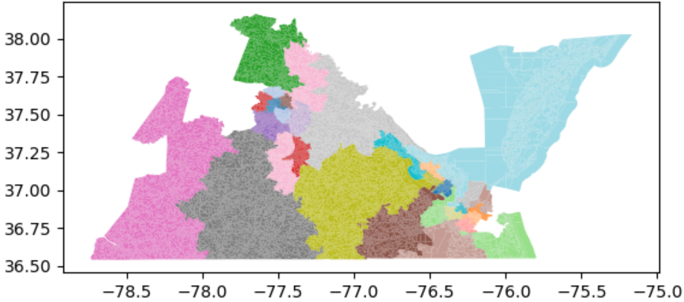

Recombination Steps

1. At each step, select two adjacent **districts**
2. Merge the subunits of those two districts
3. Draw a spanning tree for the new super–district
4. Delete an edge leaving two population balanced districts
5. Repeat
6. (Optional) Mix with single edge flips

Recombination Step Example

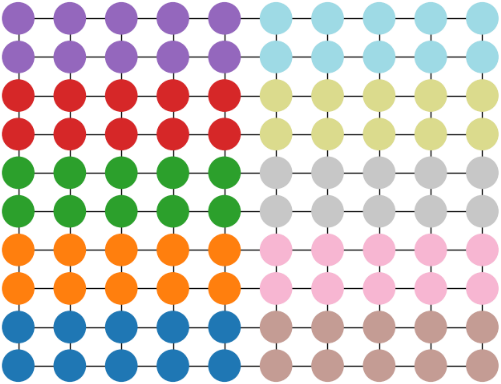

Recombination Step Example

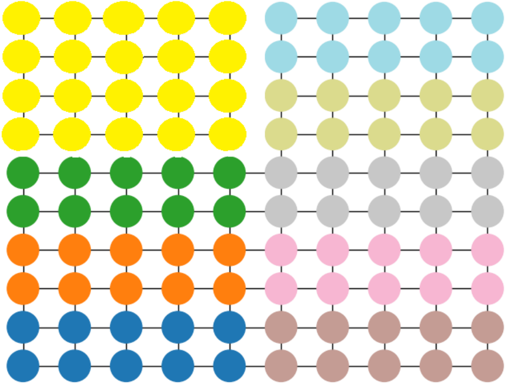

Recombination Step Example

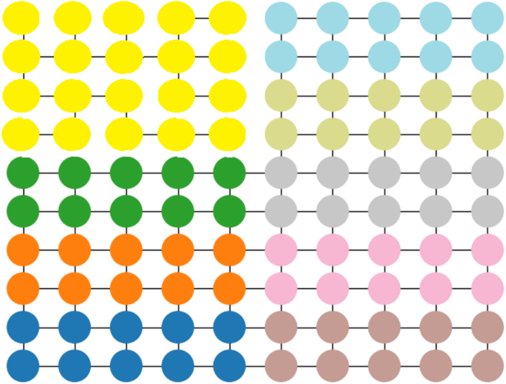

Recombination Step Example

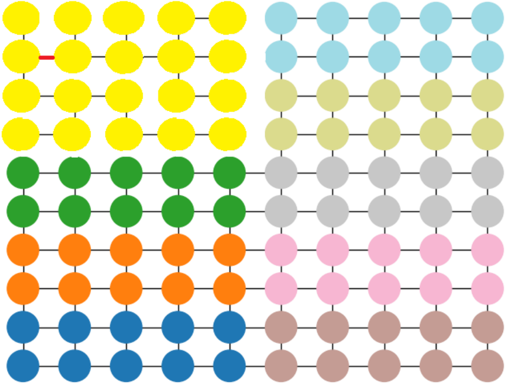

Recombination Step Example

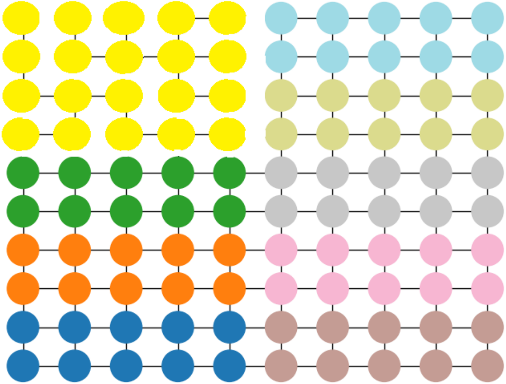

Recombination Step Example

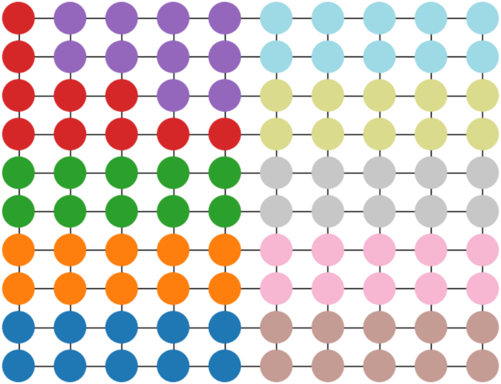

AR Ensembles

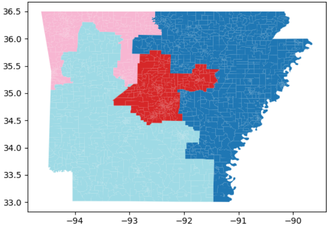

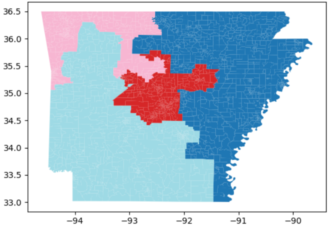

PA Recombination Steps

Recombination Distribution

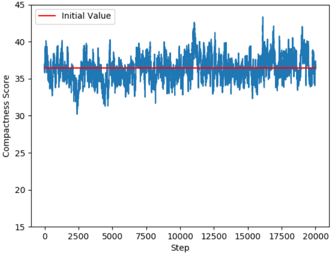

(a) Compactness

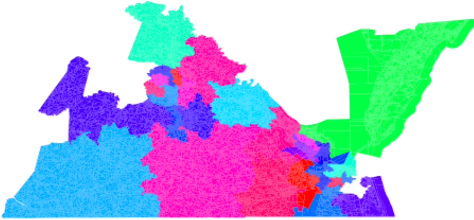

(b) 5702 cut edges

Recombination Mixing

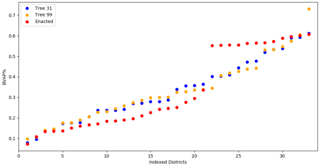

(a) Initial

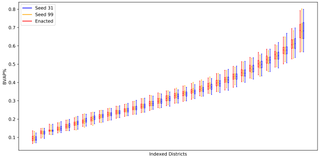

(b) 20,000 Recombination Steps

Recombination Mean–Median

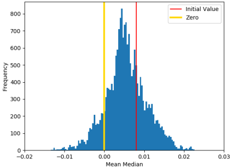

(a) ReCom Seed31

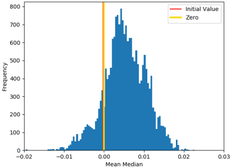

(b) ReCom Seed99

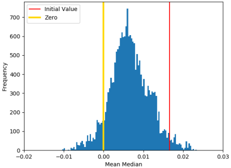

(c) ReCom Enacted

General Tree Proposals

1. Form the induced subgraph on the complement of the cut edges
2. Add some subset of the cut edges
3. Uniformly select a maximal spanning forest
4. Apply a Markov chain on trees
5. Partition the spanning forest into k population balanced pieces

Special Cases

- Uniform Trees: Add all cut edges
- k–edges: Uniformly add k cut edges
- Recombination: Add all cut edges between one pair of districts.
- Super-Recombination: Take a maximal matching on the dual graph to the districts and add all cut edges between matched districts.
- Bounce Walk: Add a single cut edge between enough pairs of districts to make a tree in the dual graph of districts.

Special Cases

- Uniform Trees: Add all cut edges
- k–edges: Uniformly add k cut edges
- Recombination: Add all cut edges between one pair of districts.
- Super-Recombination: Take a maximal matching on the dual graph to the districts and add all cut edges between matched districts.
- Bounce Walk: Add a single cut edge between enough pairs of districts to make a tree in the dual graph of districts.

Question

What are the steady state distributions (and mixing times) of these walks?

Tree Partitioning Questions

- Characterizing the distribution on partitions defined by cutting trees!
- How bad is the best cut?
- Criteria for determining when a tree is ε cuttable?
- Criteria for determining when all spanning trees of a graph are ε cuttable?
- How hard is it to find the mininum ε for which a cut exists?
- As a function of ε what proportion of spanning trees are cuttable?
- As a function of ε what proportion of edges in a given tree are cuttable?
- What is the fastest way to sample uniformly from $k-1$ balanced cut edges?

MORAL:

Computational Redistricting is NOT a solved problem!

The End

Thanks!