Several Problems in Extremal and Probabilistic Combinatorics

Corrine Yap

February 28, 2023

Rutgers University

Thanks to

Photo credit: Jozef Skokan

my advisor Bhargav Narayanan,

Thanks to

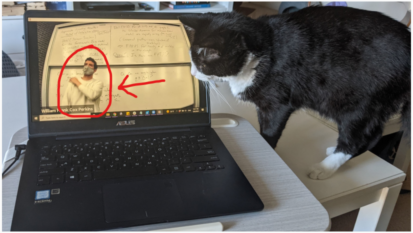

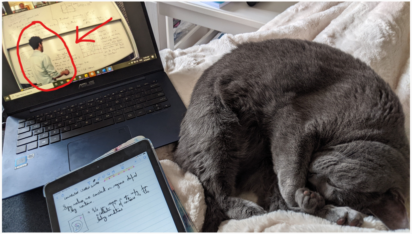

my external member Will Perkins,

Thanks to

my external member Will Perkins,

Thanks to

my committee members Jeff Kahn and Swee Hong Chan.

This thesis is dedicated to sarah-marie belcastro,

who convinced me to continue on even when I did not think I could.



Table of Contents

Table of Contents

Table of Contents

Jason Long (Squarepoint Capital)

A Topological Turán Problem

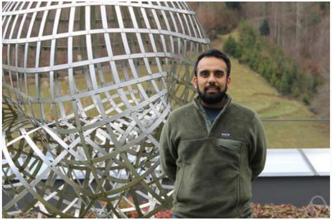

Background

Throughout the talk we will be dealing with graphs. A graph $G = (V, E)$ is a set of vertices V and a collection of edges E where an edge is a pair of vertices.

Examples

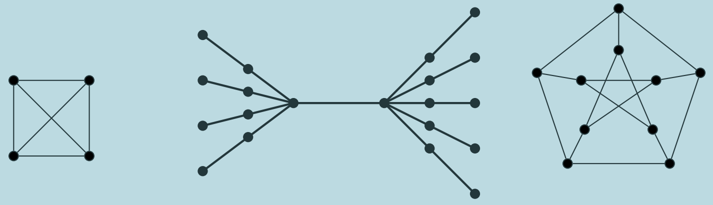

Background

Turán Problem

Given a graph H, how many edges does an n-vertex graph G need to guarantee it contains an isomorphic copy of H?

Background

Turán Problem

Given a graph H, how many edges does an n-vertex graph G need to guarantee it contains an isomorphic copy of H?

E.g. Suppose $H = \triangle$ and $n = 6$.

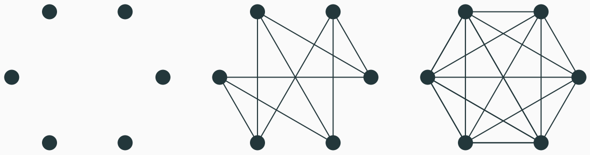

Background

Turán Problem

Given a graph H, how many edges does an n-vertex graph G need to guarantee it contains an isomorphic copy of H?

E.g. Suppose $H = \triangle$ and $n = 6$.

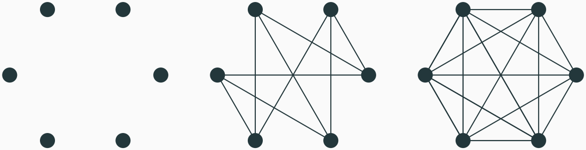

Q1: What if we replace "graph" with "r-uniform hypergraph?"
(r-graph)

Background

Hypergraph Turán Problem

Given a **hyper**graph H, how many **hyper**edges does an n-vertex **hyper**graph G need to guarantee it contains an isomorphic copy of H?

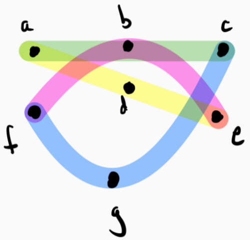

a 3-uniform hypergraph

Background

Hypergraph Turán Problem

Given a **hyper**graph *H*, how many **hyper**edges does an *n*-vertex **hyper**graph *G* need to guarantee it contains an isomorphic copy of *H*?

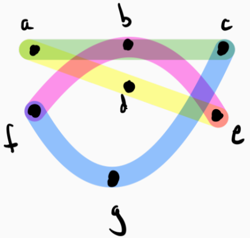

a 3-uniform hypergraph

Q2: What if we replace “isomorphic” with “homeomorphic?”

Topological POV

We can view r-graphs H and G as geometric/topological structures.

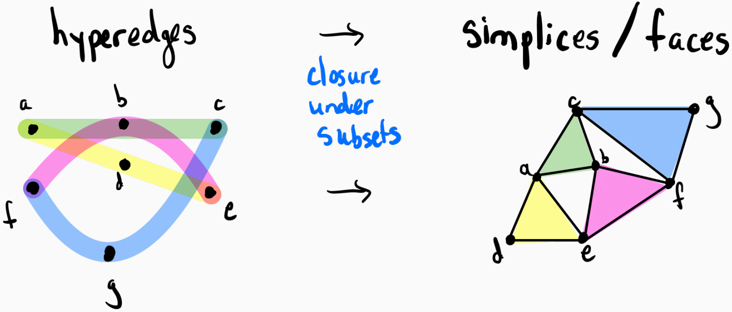

Topological POV

We can view r-graphs H and G as geometric/topological structures.

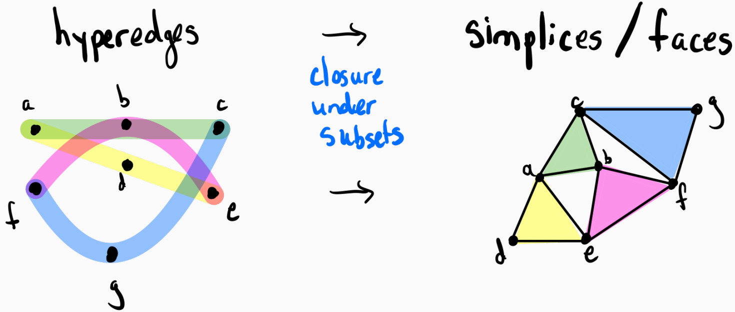

3-uniform hypergraph → 2-dimensional homogeneous simplicial complex

Topological POV

We can view r-graphs H and G as geometric/topological structures.

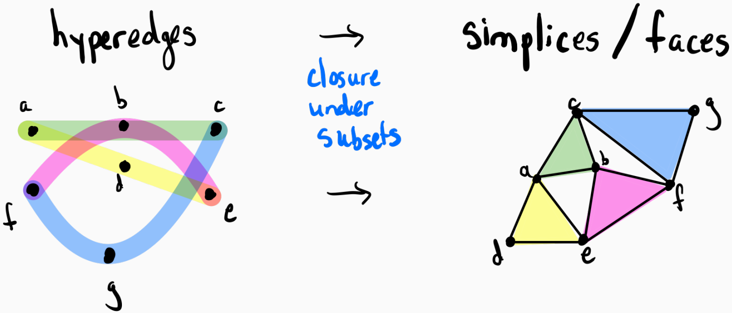

3-graph $\rightarrow$ 2-complex

Topological Turán Problem

Linial '07

Given a k-**complex** $\mathcal{S}$, how many k-**dimensional simplices (facets)** does an n-vertex k-**complex** need to guarantee it contains a **homeomorphic** copy of $\mathcal{S}$?

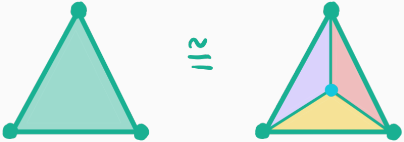

Topological Turán Problem

Linial '07

Given a k-**complex** $\mathcal{S}$, how many k-**dimensional simplices (facets)** does an n-vertex k-**complex** need to guarantee it contains a **homeomorphic** copy of $\mathcal{S}$?

Folklore: For a k-complex $\mathcal{S}$, there exists $\lambda(\mathcal{S}) > 0$ such that every k-complex on $cn^{k+1-\lambda(S)}$ facets contains a homeomorph of $\mathcal{S}$.

Why n^{k+1}? The maximum number of facets in a k-complex on n vertices is $\binom{n}{k+1} \sim n^{k+1}$.

Question

Is there some $\lambda_k > 0$ such that $\lambda(\mathcal{S}) \geq \lambda_k$ for all k-dimensional $\mathcal{S}$?

Topological Turán Problem

Linial '07

Given a k-**complex** $\mathcal{S}$, how many k-**dimensional simplices (facets)** does an n-vertex k-**complex** need to guarantee it contains a **homeomorphic** copy of $\mathcal{S}$?

Folklore: For a k-complex $\mathcal{S}$, there exists $\lambda(\mathcal{S}) > 0$ such that every k-complex on $cn^{k+1-\lambda(S)}$ facets contains a homeomorph of $\mathcal{S}$.

Why n^{k+1}? The maximum number of facets in a k-complex on n vertices is $\binom{n}{k+1} \sim n^{k+1}$.

Question

Is there some $\lambda_k > 0$ such that $\lambda(\mathcal{S}) \geq \lambda_k$ for all k-dimensional $\mathcal{S}$?

Topological Turán Problem

Linial '07

Given a k-**complex** $\mathcal{S}$, how many k-**dimensional simplices (facets)** does an n-vertex k-**complex** need to guarantee it contains a **homeomorphic** copy of $\mathcal{S}$?

Folklore: For a k-complex $\mathcal{S}$, there exists $\lambda(\mathcal{S}) > 0$ such that every k-complex on $cn^{k+1-\lambda(S)}$ facets contains a homeomorph of $\mathcal{S}$.

Why n^{k+1}? The maximum number of facets in a k-complex on n vertices is $\binom{n}{k+1} \sim n^{k+1}$.

Question

Is there some $\lambda_k > 0$ such that $\lambda(\mathcal{S}) \geq \lambda_k$ for all k-dimensional $\mathcal{S}$?

History: Dimensions 1 and 2

- Mader '67: $\lambda_1 = 1$.
- Keevash–Long–Narayanan–Scott '20: $\lambda_2 \geq \frac{1}{5}$
- Brown–Erdős–Sós '73: $\lambda(\mathcal{S}) = \frac{1}{2}$ when $\mathcal{S}$ is the 2-sphere.
- Kupavskii–Polyanskii–Tomon–Zakharov '20: $\lambda(\mathcal{S}) = \frac{1}{2}$ if $\mathcal{S}$ is a triangulation of a closed orientable surface.
- Sankar '22: extended KPTZ to non-orientable surfaces.

Conjecture

$$\lambda_2 = \frac{1}{2}$$

3-complexes: ??? General k-complexes: ???

History: Dimensions 1 and 2

- Mader '67: $\lambda_1 = 1$.
- Keevash–Long–Narayanan–Scott '20: $\lambda_2 \geq \frac{1}{5}$
- Brown–Erdős–Sós '73: $\lambda(\mathcal{S}) = \frac{1}{2}$ when $\mathcal{S}$ is the 2-sphere.
- Kupavskii–Polyanskii–Tomon–Zakharov '20: $\lambda(\mathcal{S}) = \frac{1}{2}$ if $\mathcal{S}$ is a triangulation of a closed orientable surface.
- Sankar '22: extended KPTZ to non-orientable surfaces.

Conjecture

$$\lambda_2 = \frac{1}{2}$$

3-complexes: ??? General k-complexes: ???

History: Dimensions 1 and 2

- Mader '67: $\lambda_1 = 1$.
- Keevash–Long–Narayanan–Scott '20: $\lambda_2 \geq \frac{1}{5}$
- Brown–Erdős–Sós '73: $\lambda(\mathcal{S}) = \frac{1}{2}$ when $\mathcal{S}$ is the 2-sphere.
- Kupavskii–Polyanskii–Tomon–Zakharov '20: $\lambda(\mathcal{S}) = \frac{1}{2}$ if $\mathcal{S}$ is a triangulation of a closed orientable surface.
- Sankar '22: extended KPTZ to non-orientable surfaces.

Conjecture

$$\lambda_2 = \frac{1}{2}$$

3-complexes: ??? General k-complexes: ???

Main Theorem

Theorem (Long–Narayanan–Y. '20)

$\lambda_k \geq k^{-2k^2}$ for all $k \in \mathbb{N}$.

Far from tight: e.g. $\lambda_2 \geq 2^{-8}$

But the proof is purely combinatorial, compared to the dimension-specific arguments in previous results.

Our combinatorial/probabilistic tools:

1. Trace-bounded hypergraphs
2. Dependent random choice

Main Theorem

Theorem (Long–Narayanan–Y. ’20)

$\lambda_k \geq k^{-2k^2}$ for all $k \in \mathbb{N}$.

Far from tight: e.g. $\lambda_2 \geq 2^{-8}$

But the proof is purely combinatorial, compared to the dimension-specific arguments in previous results.

Our combinatorial/probabilistic tools:

1. Trace-bounded hypergraphs
2. Dependent random choice

Trace-Bounded Hypergraphs

Key observation: Every k-complex has a homeomorph whose associated $(k+1)$-graph is *trace-bounded*.

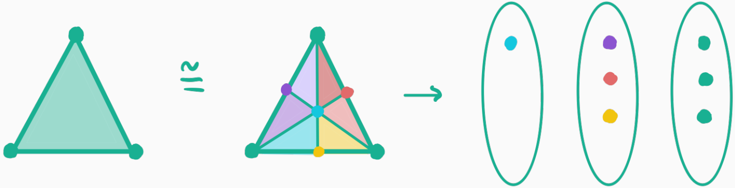

Trace-Bounded Hypergraphs

Key observation: Every k-complex has a homeomorph whose associated $(k+1)$-graph is *trace-bounded*.

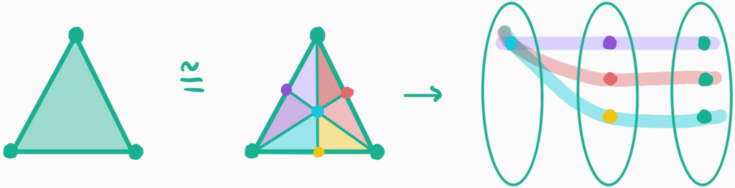

Trace-Bounded Hypergraphs

Key observation: Every k-complex has a homeomorph whose associated $(k+1)$-graph is *trace-bounded*.

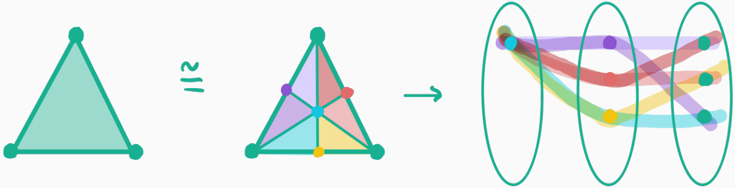

Trace-Bounded Hypergraphs

Key observation: Every k-complex has a homeomorph whose associated $(k+1)$-graph is *trace-bounded*.

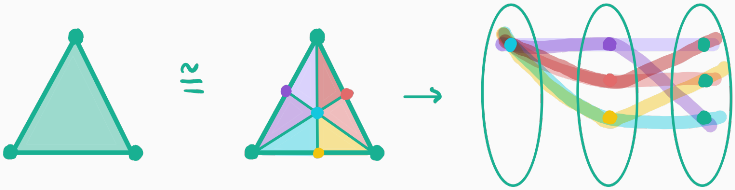

trace-bounded: all the "forward" degrees are bounded.

Trace-Bounded Hypergraphs

Key observation: Every k-complex has a homeomorph whose associated $(k+1)$-graph is *trace-bounded*.

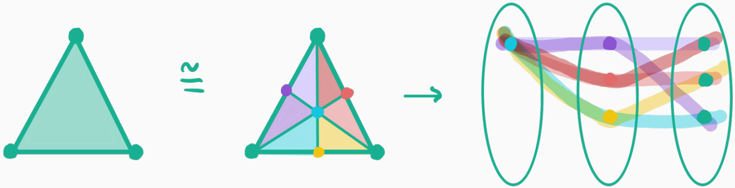

trace-bounded: all the "forward" degrees are bounded.

So to find a homeomorph of $\mathcal{S}$ (in a dense k-complex), it's enough to find a copy of every trace-bounded $(k+1)$-graph (in a dense $(k+1)$-graph).

Main Combinatorial Result

Theorem 2 (LNY '20)

$\exists\ \alpha_{r,d} \geq (5rd)^{1-r}$ such that for any d-**trace-bounded** r-partite r-graph H, any r-graph with at least $n^{r-\alpha_{r,d}}$ edges contains an (isomorphic) copy of H.

Proof idea: Use a variation of *dependent random choice* to embed H one part at a time.

Main Combinatorial Result

Theorem 2 (LNY '20)

$\exists\, \alpha_{r,d} \geq (5rd)^{1-r}$ such that for any d-**trace-bounded** r-partite r-graph H, any r-graph with at least $n^{r-\alpha_{r,d}}$ edges contains an (isomorphic) copy of H.

Proof idea: Use a variation of *dependent random choice* to embed H one part at a time.

Further Directions

- Optimal values for λ_k (in particular, $\lambda_2 = \frac{1}{2}$)?
- $\lambda(S^k)$?
- Optimal values for $\alpha_{r,d}$?
- Other settings where trace-boundedness is a natural restriction?

Hypergraph Ramsey Theory

Quentin Dubroff

(Rutgers)

António Girão

(Oxford)

Eoin Hurley

(Amsterdam)

Background

Definition

The complete k-uniform hypergraph on n vertices is the collection of all subsets of size k and is denoted $K_n^{(k)}$.

Definition

The Ramsey number $r_k(t)$ is the minimum n such that any 2-coloring of (the edges of) $K_n^{(k)}$ contains a monochromatic $K_t^{(k)}$.

Background

Definition

The complete k-uniform hypergraph on n vertices is the collection of all subsets of size k and is denoted $K_n^{(k)}$.

Definition

The Ramsey number $r_k(t)$ is the minimum n such that any 2-coloring of (the edges of) $K_n^{(k)}$ contains a monochromatic $K_t^{(k)}$.

Example

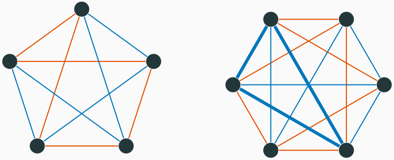

$$r_2(3) = 6$$

History

A central open question: what is the behavior of $r_3(t)$?

Erdős–Hajnal–Rado (1964)

There exist constants $c, c' > 0$ such that

$$2^{ct^2} \leq r_3(t) \leq 2^{2^{c't}}$$

History

A central open question: what is the behavior of $r_3(t)$?

Erdős–Hajnal–Rado (1964)

There exist constants $c, c' > 0$ such that

$$2^{ct^2} \leq r_3(t) \leq 2^{2^{c't}}$$

Why is $k = 3$ so important?

Stepping-Up Lemma (Erdős–Hajnal, 1972)

$$r_{k+1}(2t + k - 4) \geq 2^{r_k(t)}.$$

History

Erdős–Hajnal–Rado (1964)

There exist constants $c, c' > 0$ such that

$$2^{ct^2} \leq r_3(t) \leq 2^{2^{c't}}$$

Why is $k = 3$ so important?

Stepping-Up Lemma (Erdős–Hajnal, 1972)

$$r_{k+1}(2t + k - 4) \geq 2^{r_k(t)}.$$

This implies $r_k(t) \geq T_{k-1}(ct^2)$ where $T_1(x) = x, T_i(x) = 2^{T_{i-1}(x)}$.

EHR also showed $r_k(t) \leq T_k(ct)$. Showing $r_3(t)$ is double-exponential would tell us $r_k(t) = \Theta(T_k(t))$ for all $k > 3$.

Many Colors

One approach: find out what can happen when we use more than 2 colors.

Definition

The Ramsey number $r_k(t; q)$ is the minimum n such that any q-coloring of $K_n^{(k)}$ contains a monochromatic $K_t^{(k)}$.

Many Colors

One approach: find out what can happen when we use more than 2 colors.

Definition

The Ramsey number $r_k(t; q)$ is the minimum n such that any q-coloring of $K_n^{(k)}$ contains a monochromatic $K_t^{(k)}$.

- Erdős and Hajnal (1972): $r_3(t; 4) \geq 2^{2^{ct}}$.
- Conlon, Fox, and Sudakov ('10): $r_3(t; 3) \geq 2^{t^{c \log t}}$.

Erdős's Conjecture

$$r_3(t; 2) = \Theta(2^{2^t})$$

Many Colors

What if we relax the clique requirement to a different type of graph? Can the 2-color Ramsey number be "that different" from the 4-color, or in general q-color, Ramsey number?

Definition

The Ramsey number $r_k(G; q)$ is the minimum n such that any q-coloring of $K_n^{(k)}$ contains a monochromatic copy of G.

Hedgehogs

Definition

The Ramsey number $r_k(G; q)$ is the minimum n such that any q-coloring of $K_n^{(k)}$ contains a monochromatic copy of G.

Conlon–Fox–Rödl, '17

There is a family of k-uniform hypergraphs called *hedgehogs* such that $r_k(G; 2)$ is bounded above by a polynomial in the size of G but $r_k(G; 4)$ is at least double-exponential.

Hedgehogs

Definition

The Ramsey number $r_k(G; q)$ is the minimum n such that any q-coloring of $K_n^{(k)}$ contains a monochromatic copy of G.

Conlon–Fox–Rödl, ’17

There is a family of k-uniform hypergraphs called *hedgehogs* such that $r_k(G; 2)$ is bounded above by a polynomial in the size of G but $r_k(G; 4)$ is at least double-exponential.

Question (CFR)

Is there a family of hypergraphs whose 2-color and q-color Ramsey numbers differ by an arbitrarily large tower?

Main Theorem

Let $\hat{H}_t$ be the *balanced hedgehog*.

Theorem (Dubroff–Girão–Hurley–Y. '22)

There exist $c > 0$ and $q : \mathbb{N} \to \mathbb{N}$ such that for all $k \in \mathbb{N}$ and sufficiently large t, we have

(1) $r_{2k+1}(\hat{H}_t; 2) \leq t^{k+3}$, and

(2) $r_{2k+1}(\hat{H}_t; q(k)) \geq T_{\lfloor c \log_2 \log_2 k \rfloor}(t)$.

Main Theorem

Let $\hat{H}_t$ be the *balanced hedgehog*.

Theorem (Dubroff–Girão–Hurley–Y. '22)

There exist $c > 0$ and $q : \mathbb{N} \to \mathbb{N}$ such that for all $k \in \mathbb{N}$ and sufficiently large t, we have

(1) $r_{2k+1}(\hat{H}_t; 2) \leq t^{k+3}$, and

(2) $r_{2k+1}(\hat{H}_t; q(k)) \geq T_{\lfloor c \log_2 \log_2 k \rfloor}(t)$.

3-uniform hedgehog H_3

Main Theorem

Let $\hat{H}_t$ be the *balanced hedgehog*.

Theorem (Dubroff–Girão–Hurley–Y. '22)

There exist $c > 0$ and $q : \mathbb{N} \to \mathbb{N}$ such that for all $k \in \mathbb{N}$ and sufficiently large t, we have

(1) $r_{2k+1}(\hat{H}_t; 2) \le t^{k+3}$, and

(2) $r_{2k+1}(\hat{H}_t; q(k)) \ge T_{\lfloor c \log_2 \log_2 k \rfloor}(t)$.

3-uniform hedgehog H_3

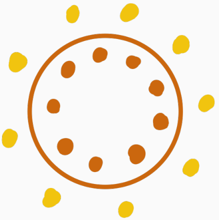

Main Theorem

Let $\hat{H}_t$ be the *balanced hedgehog*.

Theorem (Dubroff–Girão–Hurley–Y. '22)

There exist $c > 0$ and $q : \mathbb{N} \to \mathbb{N}$ such that for all $k \in \mathbb{N}$ and sufficiently large t, we have

(1) $r_{2k+1}(\hat{H}_t; 2) \leq t^{k+3}$, and

(2) $r_{2k+1}(\hat{H}_t; q(k)) \geq T_{\lfloor c \log_2 \log_2 k \rfloor}(t)$.

3-uniform hedgehog H_3

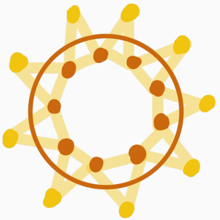

Main Theorem

Let $\hat{H}_t$ be the *balanced hedgehog*.

Theorem (Dubroff–Girão–Hurley–Y. '22)

There exist $c > 0$ and $q : \mathbb{N} \to \mathbb{N}$ such that for all $k \in \mathbb{N}$ and sufficiently large t, we have

(1) $r_{2k+1}(\hat{H}_t; 2) \leq t^{k+3}$, and

(2) $r_{2k+1}(\hat{H}_t; q(k)) \geq T_{\lfloor c \log_2 \log_2 k \rfloor}(t)$.

3-uniform hedgehog H_3

Main Theorem

Let $\hat{H}_t$ be the *balanced hedgehog*.

Theorem (Dubroff–Girão–Hurley–Y. '22)

There exist $c > 0$ and $q : \mathbb{N} \to \mathbb{N}$ such that for all $k \in \mathbb{N}$ and sufficiently large t, we have

(1) $r_{2k+1}(\hat{H}_t; 2) \leq t^{k+3}$, and

(2) $r_{2k+1}(\hat{H}_t; q(k)) \geq T_{\lfloor c \log_2 \log_2 k \rfloor}(t)$.

3-uniform hedgehog H_3

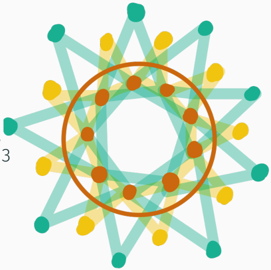

Main Theorem

Let $\hat{H}_t$ be the *balanced hedgehog*.

Theorem (Dubroff–Girão–Hurley–Y. '22)

There exist $c > 0$ and $q : \mathbb{N} \to \mathbb{N}$ such that for all $k \in \mathbb{N}$ and sufficiently large t, we have

(1) $r_{2k+1}(\hat{H}_t; 2) \leq t^{k+3}$, and

(2) $r_{2k+1}(\hat{H}_t; q(k)) \geq T_{\lfloor c \log_2 \log_2 k \rfloor}(t)$.

3-uniform hedgehog H_3

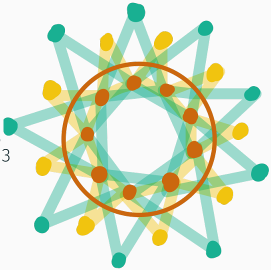

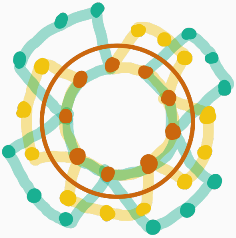

6-uniform balanced hedgehog $\hat{H}_6$

Example: Stepping up from 2- to 3-uniform

Given: “good” coloring of K_m. Want: “good” coloring of $K^{(3)}_{2^m}$.

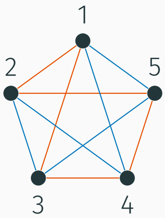

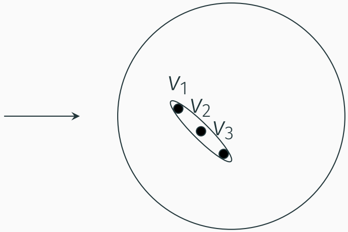

V' = binary strings of length 5

Example: Stepping up from 2- to 3-uniform

Given: "good" coloring of K_m. Want: "good" coloring of $K^{(3)}_{2^m}$.

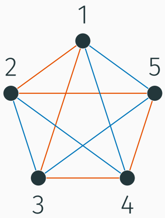

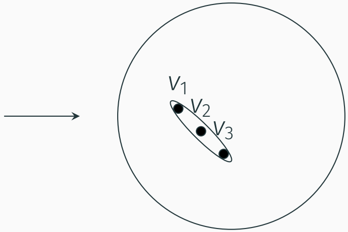

$V' =$ binary strings of length 5

For $u, v \in V'$, define $\delta(u, v) = \{\max i : v_i \neq u_i\}$.

$$\left.\begin{aligned} v_1 &= (0,0,0,0,1) \\ v_2 &= (0,1,1,0,1) \\ v_3 &= (1,0,1,1,0) \end{aligned}\right\} \quad \begin{aligned} \delta_1 &= 3 \\ \delta_2 &= 5 \end{aligned}$$

Example: Stepping up from 2- to 3-uniform

Given: “good” coloring of K_m. Want: “good” coloring of $K^{(3)}_{2^m}$.

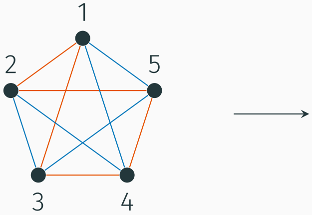

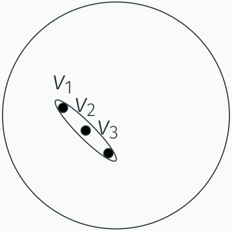

$V' =$ binary strings of length 5

For $u, v \in V'$, define $\delta(u,v) = \{\max i : v_i \neq u_i\}$.

$$\left.\begin{aligned} v_1 &= (0,0,0,0,1) \\ v_2 &= (0,1,1,0,1) \\ v_3 &= (1,0,1,1,0) \end{aligned}\right\} \quad \begin{aligned} \delta_1 &= 3 \\ \delta_2 &= 5 \end{aligned}$$

Edge $(3,5)$ is blue, and (δ_1, δ_2) is increasing, which tells us what color to assign (v_1, v_2, v_3).

Lower Bound

We show an “Erdős-Hajnal”-type result for a certain sequence property: the δ-sequence either contains a long monotone subsequence or some “forbidden” subsequence.

By stepping-up, we construct colorings where cliques span many colors. From this, we produce colorings with no monochromatic $\hat{H}_t$.

Lower Bound

We show an “Erdős-Hajnal”-type result for a certain sequence property: the δ-sequence either contains a long monotone subsequence or some “forbidden” subsequence.

By stepping-up, we construct colorings where cliques span many colors. From this, we produce colorings with no monochromatic $\hat{H}_t$.

Further Questions

- Can we say more about Ramsey numbers when we replace "monochromatic" with "spanning fewer than p colors" (denoted $r_k(t; q, p)$)?
- In particular, Conlon-Fox-Rödl asked: is there an integer q such that $r_3(t; q, 3) \geq 2^{2^{ct}}$? Our methods don't apply here because of limited ranges for the stepping-up parameters.

Reconstructing Random Pictures

Background

Reconstruction Problem

Given a discrete structure, can we uniquely reconstruct it from the list of its substructures of a fixed size?

Most famous example: graphs—Vertex and Edge Reconstruction Conjectures (Kelly, Ulam 1957, Harary 1964)

Mossel–Ross ’18

What about “shotgun assembly?” (motivated by shotgun sequencing of DNA)

Background

Reconstruction Problem

Given a discrete structure, can we uniquely reconstruct it from the list of its substructures of a fixed size?

Most famous example: graphs—Vertex and Edge Reconstruction Conjectures (Kelly, Ulam 1957, Harary 1964)

Mossel–Ross ’18

What about “shotgun assembly?” (motivated by shotgun sequencing of DNA)

Background

Reconstruction Problem

Given a discrete structure, can we uniquely reconstruct it from the list of its substructures of a fixed size?

Most famous example: graphs—Vertex and Edge Reconstruction Conjectures (Kelly, Ulam 1957, Harary 1964)

Mossel–Ross '18

What about "shotgun assembly?" (motivated by shotgun sequencing of DNA)

Random Pictures

Today: Let P_n be a random picture, i.e. an $n \times n$ grid with $\{0, 1\}$ entries chosen uniformly at random. Let $\mathcal{D}$ be the deck of its $k \times k$ subgrids.

Question

For what $k = k(n)$ is P_n reconstructible from $\mathcal{D}$ with high probability?

Example

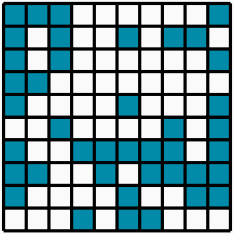

A 10 × 10 picture

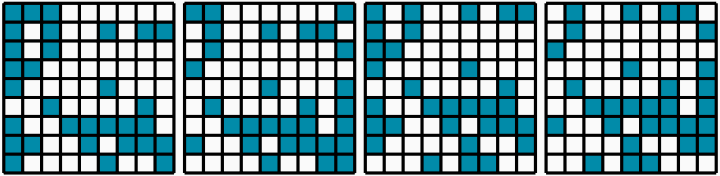

Deck of 9 × 9 subgrids

Example

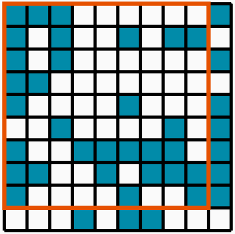

A 10 × 10 picture

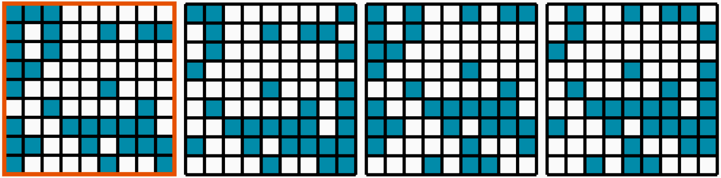

Deck of 9 × 9 subgrids

Example

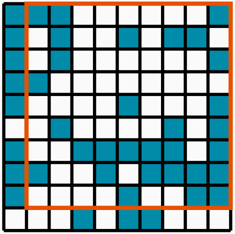

A 10 $\times$ 10 picture

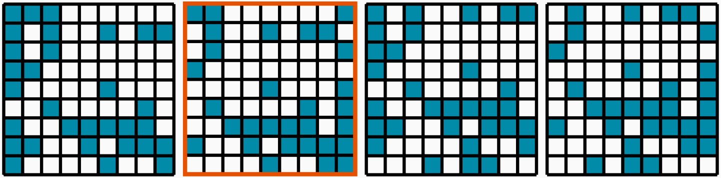

Deck of 9 $\times$ 9 subgrids

Example

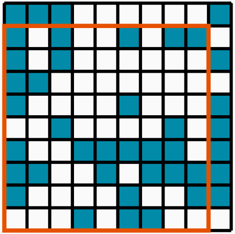

A 10 × 10 picture

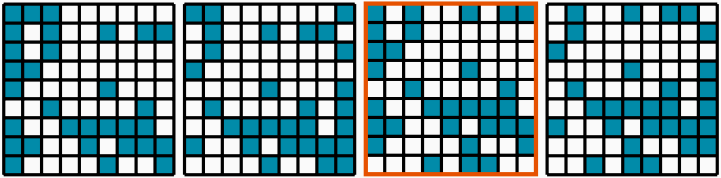

Deck of 9 × 9 subgrids

Example

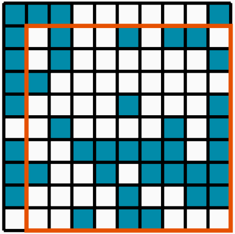

A 10×10 picture

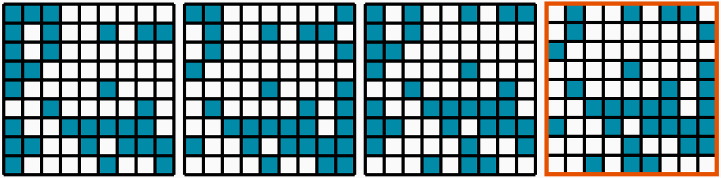

Deck of 9×9 subgrids

Example

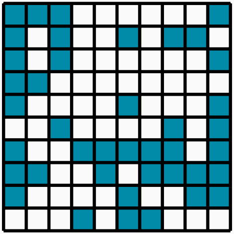

A 10×10 picture

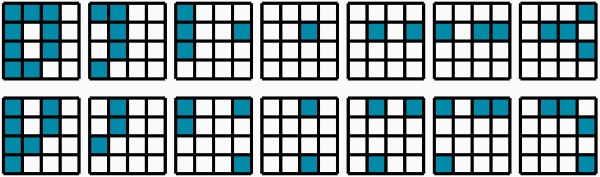

...

Deck of 4×4 subgrids

Example

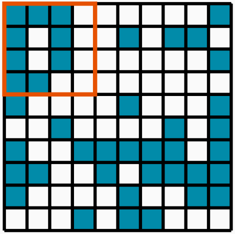

A 10 × 10 picture

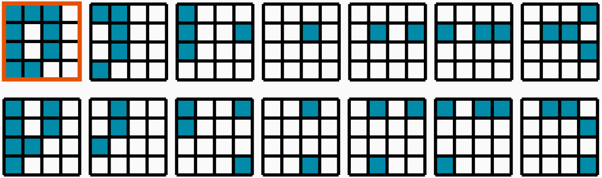

. . .

Deck of 4 × 4 subgrids

Example

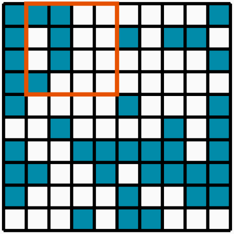

A 10 × 10 picture

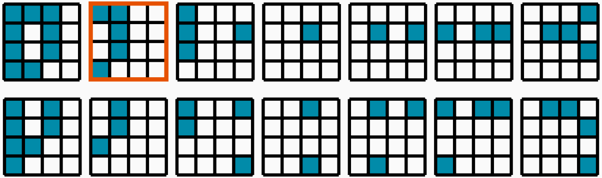

. . .

Deck of 4 × 4 subgrids

Example

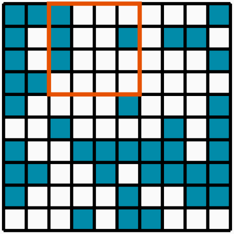

A 10 × 10 picture

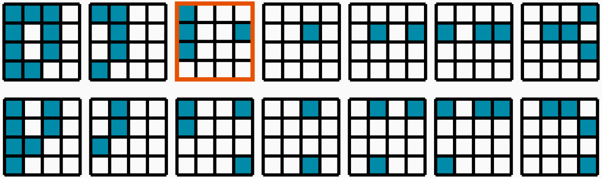

. . .

Deck of 4 × 4 subgrids

Example

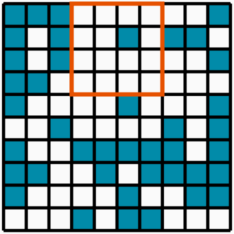

A 10×10 picture

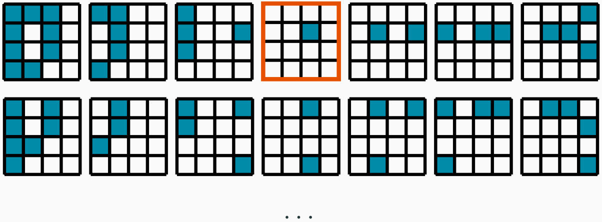

Deck of 4×4 subgrids

Example

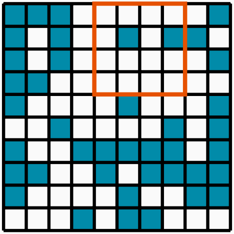

A 10 × 10 picture

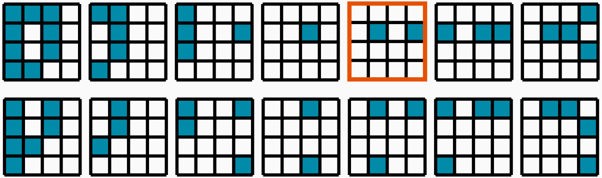

. . .

Deck of 4 × 4 subgrids

Example

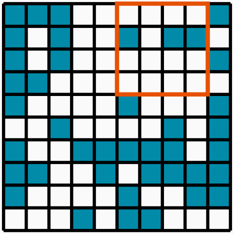

A 10×10 picture

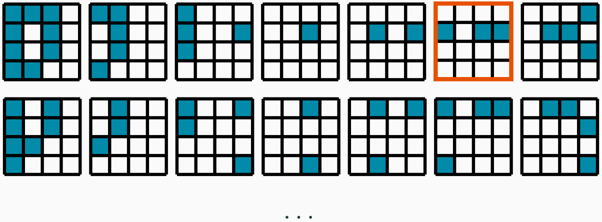

Deck of 4×4 subgrids

Example

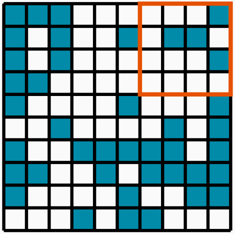

A 10 × 10 picture

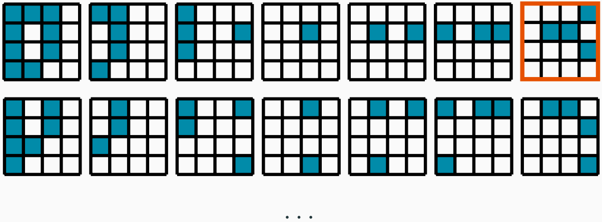

Deck of 4 × 4 subgrids

Example

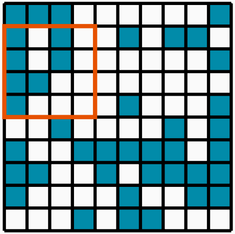

A 10 × 10 picture

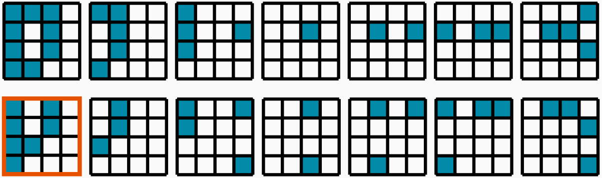

...

Deck of 4 × 4 subgrids

Example

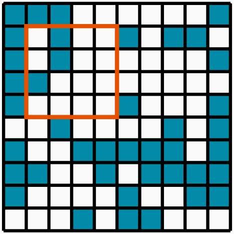

A 10 × 10 picture

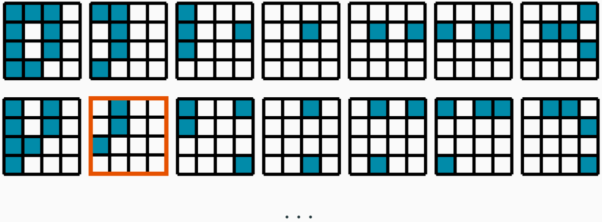

Deck of 4 × 4 subgrids

Example

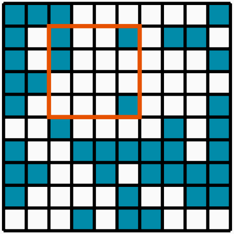

A 10×10 picture

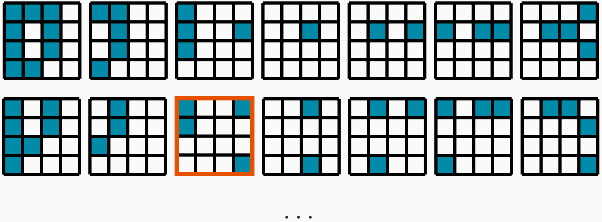

Deck of 4×4 subgrids

Example

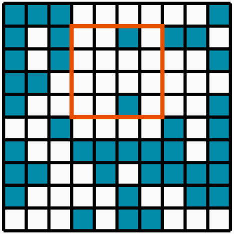

A 10×10 picture

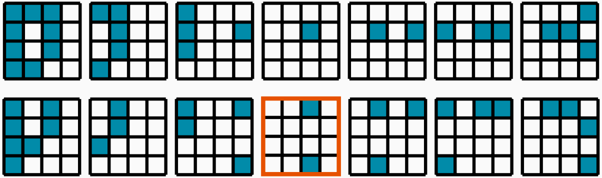

$\cdots$

Deck of 4×4 subgrids

Example

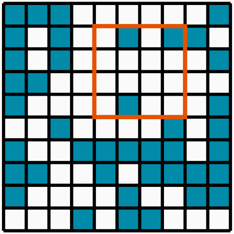

A 10×10 picture

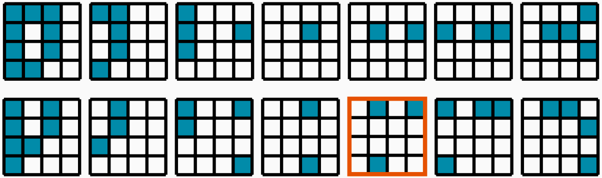

. . .

Deck of 4×4 subgrids

Example

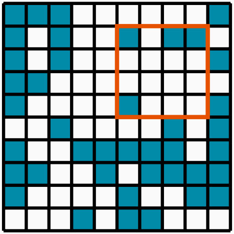

A 10 × 10 picture

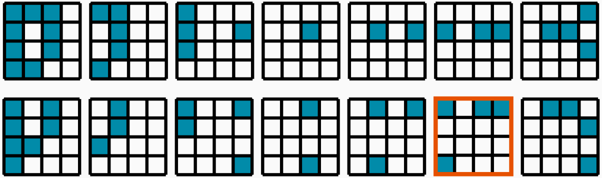

. . .

Deck of 4 × 4 subgrids

Example

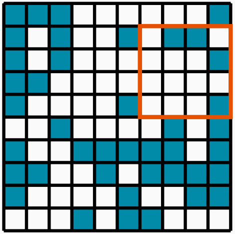

A 10 × 10 picture

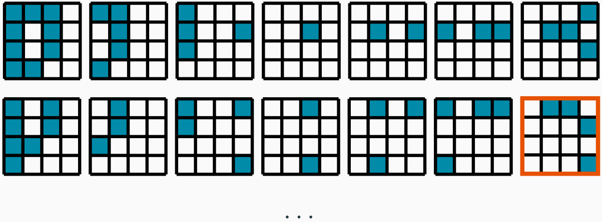

Deck of 4 × 4 subgrids

Main Theorem

Let $R(n,k)$ be the event that P_n is reconstructible from its k-deck.

Narayanan-Y. '22+

There exists $k_c(n)$ such that as $n \to \infty$,

$$\mathrm{Prob}[R(n,k)] \to \begin{cases} 0 & \text{if } k < k_c(n) \\ 1 & \text{if } k > k_c(n) \end{cases}$$

Moreover, $k_c(n)$ takes one of two values: $\lfloor \sqrt{2 \log_2 n} \rfloor, \lceil \sqrt{2 \log_2 n} \rceil$.

Main Theorem

Let $R(n,k)$ be the event that P_n is reconstructible from its k-deck.

Narayanan-Y. '22+

There exists $k_c(n)$ such that as $n \to \infty$,

$$\mathrm{Prob}[R(n,k)] \to \begin{cases} 0 & \text{if } k < k_c(n) \\ 1 & \text{if } k > k_c(n) \end{cases}$$

Moreover, $k_c(n)$ takes one of two values: $\lfloor \sqrt{2\log_2 n} \rfloor, \lceil \sqrt{2\log_2 n} \rceil$.

Proof of the 0-Statement: If $k < k_c(n)$, then $n^2 2^{-k^2} \to \infty$ as $n \to \infty$.

Counting argument; bound the number of reconstructible pictures by the number of k-decks.

Main Theorem

Let $R(n,k)$ be the event that P_n is reconstructible from its k-deck.

Narayanan-Y. '22+

There exists $k_c(n)$ such that as $n \to \infty$,

$$\mathrm{Prob}[R(n,k)] \to \begin{cases} 0 & \text{if } k < k_c(n) \\ 1 & \text{if } k > k_c(n) \end{cases}$$

Moreover, $k_c(n)$ takes one of two values: $\lfloor \sqrt{2\log_2 n} \rfloor, \lceil \sqrt{2\log_2 n} \rceil$.

Proof of the 1-Statement: If $k > k_c(n)$, then $n^2 k 2^{-k^2+k} \to 0$. Our goal is to give an algorithm for reconstructing P_n from its deck and prove that the probability of failure tends to 0.

Reconstruction Algorithm

Step 0: Arbitrarily order the deck $\mathcal{D}$ and begin with the first deck element.

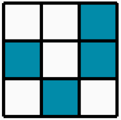

Reconstruction Algorithm

Step 1: Extend downward to $3k$ rows by placing the first deck element that fits.

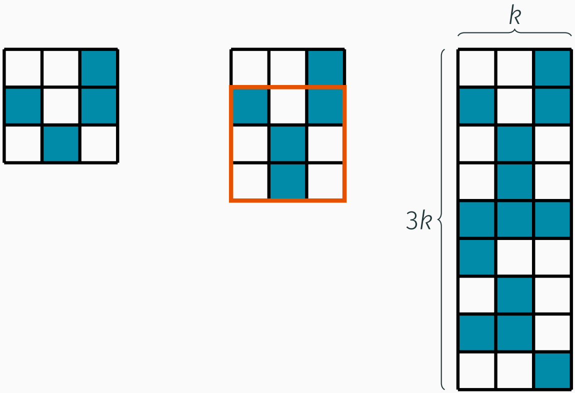

Reconstruction Algorithm

Step 2: Extend to the right one column at a time, first at each of the corners

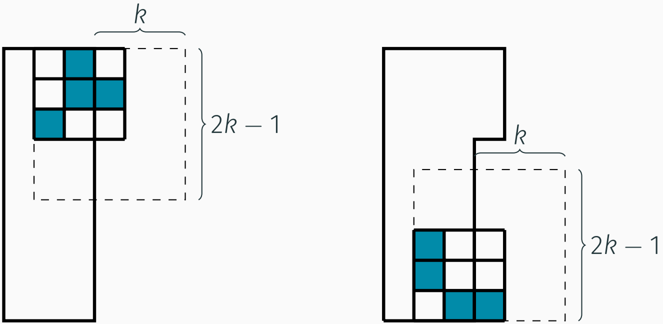

Reconstruction Algorithm

Step 2: Extend to the right one column at a time, first at each of the corners then internally. Repeat to the right and left until n columns.

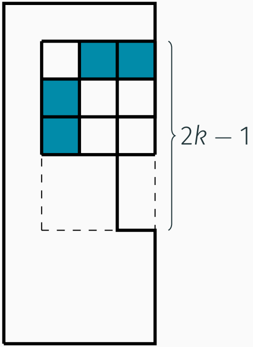

Reconstruction Algorithm

Step 3: Extend upward one row at a time, then downward until n rows.

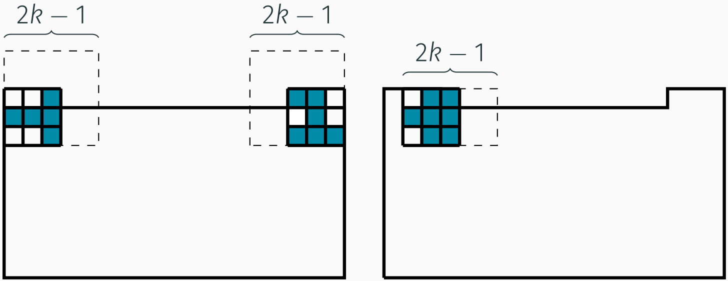

Naive Extensions

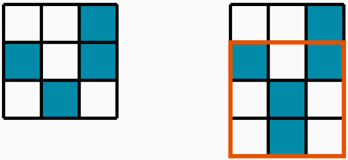

Observe that for each naive extension,

$$\mathrm{Prob}[\text{mistake}] \leq n^2 2^{-k^2+k}$$

So by union bound,

$$\mathrm{Prob}[\text{there is a mistake in the first step}] \leq 3kn^2 2^{-k^2+k}$$

which tends to 0 by our assumption. However, we cannot afford to do naive extensions for the entire grid. This is why we introduce the corner and internal extensions.

Naive Extensions

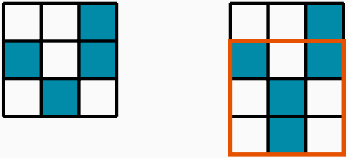

Observe that for each naive extension,

$$\mathrm{Prob}[\text{mistake}] \leq n^2 2^{-k^2+k}$$

So by union bound,

$$\mathrm{Prob}[\text{there is a mistake in the first step}] \leq 3kn^2 2^{-k^2+k}$$

which tends to 0 by our assumption. However, we cannot afford to do naive extensions for the entire grid. This is why we introduce the corner and internal extensions.

Corner Extensions

Suppose we have reconstructed S and are extending to the right. Before placing a corner subgrid T, we check to see if it can be extended to a $(2k-1) \times (2k-1)$ subgrid S' using deck elements.

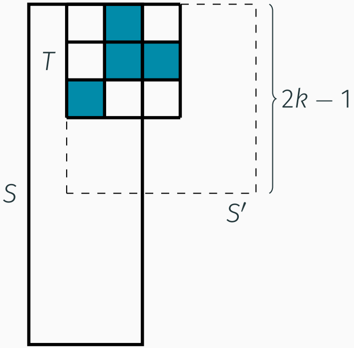

Corner Extensions

Suppose we have reconstructed S and are extending to the right. Before placing a corner subgrid T, we check to see if it can be extended to a $(2k-1) \times (2k-1)$ subgrid S' using deck elements.

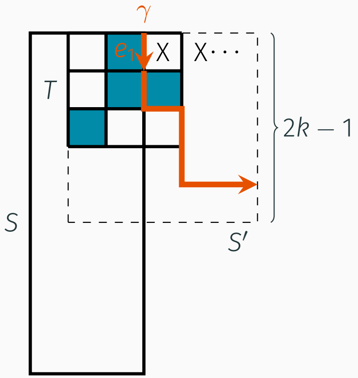

A k-grid is *bad* if it is incorrect with respect to P_n. We mark the upper-right corner with an X.

An interface path is a path separating the good and bad entries.

Interface Paths

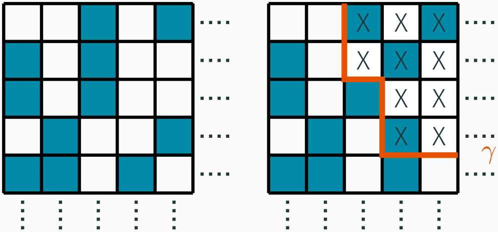

We compute probabilities associated with the interface paths. For example,

$$\mathrm{Prob}[\text{first step}] \leq n^2 2^{-k^2+k}$$

but

$$\mathrm{Prob}[\text{second step} \mid \text{first step}] \leq n^2 2^{-k^2+1} + 2(4k^2)(2^{-k+1})$$

Further Directions

- Sharp threshold?
- Higher dimensions
- Variants: p-biased, multicolor, noisy, correlated...

Charlie Carlson
(CU Boulder → UCSB)

Ewan Davies
(Colorado State)

Algorithms for the Potts Model

Nicolas Fraiman
(UNC Chapel Hill)

Alexandra Kolla
(UCSC)

Aditya Potukuchi
(York)

Background

In statistical physics, we want to determine macroscopic properties of physical systems of particles by studying their microscopic interactions.

- macroscopic: phase transitions, long-range correlations
- microscopic: nearest neighbor interactions

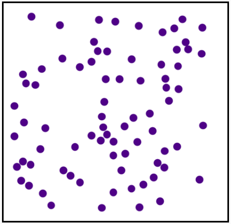

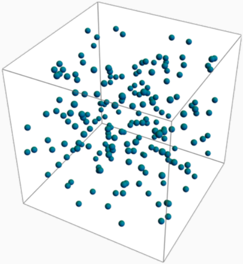

Background

In statistical physics, we want to determine macroscopic properties of physical systems of particles by studying their microscopic interactions.

- macroscopic: phase transitions, long-range correlations
- microscopic: nearest neighbor interactions

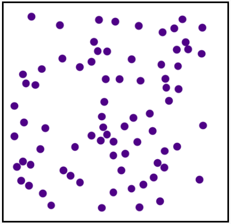

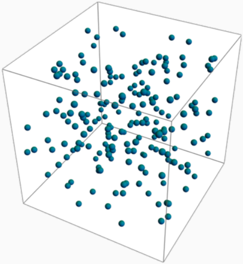

Background

In statistical physics, we want to determine macroscopic properties of physical systems of particles by studying their microscopic interactions.

- macroscopic: phase transitions, long-range correlations
- microscopic: nearest neighbor interactions

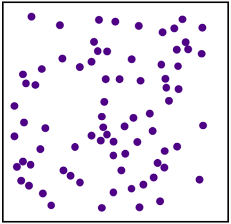

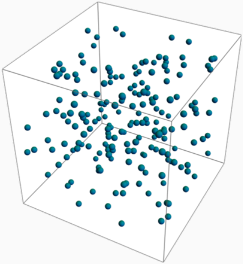

Background

In statistical physics, we want to determine macroscopic properties of physical systems of particles by studying their microscopic interactions.

- macroscopic: phase transitions, long-range correlations
- microscopic: nearest neighbor interactions

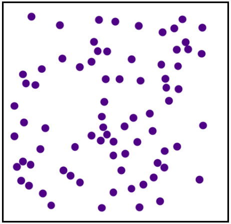

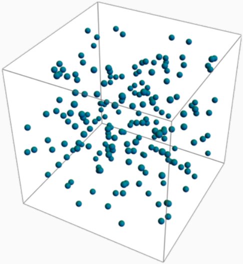

How can we study these systems mathematically? Graphs!

We begin with

- a graph $G = (V, E)$
- a finite set of spins $\{1, 2, \ldots, q\}$, and
- spin configurations $\sigma : V \to \{1, \ldots, q\}$

Example

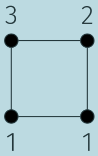

An example spin configuration σ.

We begin with

- a graph $G = (V, E)$
- a finite set of spins $\{1, 2, \ldots, q\}$, and
- spin configurations $\sigma : V \to \{1, \ldots, q\}$

Example

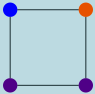

An example spin configuration σ.

We begin with

- a graph $G = (V, E)$
- a finite set of spins $\{1, 2, \ldots, q\}$, and
- spin configurations $\sigma : V \to \{1, \ldots, q\}$ aka vertex-colorings

Example

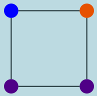

An example spin configuration σ.

We begin with

- a graph $G = (V, E)$
- a finite set of spins $\{1, 2, \ldots, q\}$, and
- spin configurations $\sigma : V \to \{1, \ldots, q\}$ aka vertex-colorings

Example

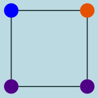

An example spin configuration σ.

- a parameter $\beta > 0$ called the inverse temperature
- $m(G, \sigma)$ = number of monochromatic edges in σ

We begin with

- a graph $G = (V, E)$
- a finite set of spins $\{1, 2, \ldots, q\}$, and
- spin configurations $\sigma : V \to \{1, \ldots, q\}$ aka vertex-colorings
- a parameter $\beta > 0$ called the inverse temperature
- $m(G, \sigma)$ = number of monochromatic edges in σ

Definition

The **Potts model** is defined by a probability distribution μ on q-vertex-colorings. For a coloring σ, we set $\mu_G(\sigma) \propto e^{\beta m(G,\sigma)}$, i.e.

$$\mu_G(\sigma) = \frac{e^{\beta m(G,\sigma)}}{\sum_{\text{colorings}\chi} e^{\beta m(G,\chi)}} =: \frac{e^{\beta m(G,\sigma)}}{Z_G(q, \beta)}$$

The denominator $Z_G(q, \beta)$ is called the *partition function*.

Understanding the Distribution

For a coloring σ, we set $\mu_G(\sigma) \propto e^{\beta m(G,\sigma)}$

Getting Some Intuition

- As $\beta \to 0$, μ_G approaches the *uniform distribution over all q-colorings.*
- As $\beta \to \infty$, μ_G approaches the *uniform distribution over all monochromatic states.*

Low β (high temperature)

Understanding the Distribution

For a coloring σ, we set $\mu_G(\sigma) \propto e^{\beta m(G,\sigma)}$

Getting Some Intuition

- As $\beta \to 0$, μ_G approaches the *uniform distribution over all q-colorings.*
- As $\beta \to \infty$, μ_G approaches the *uniform distribution over all monochromatic states.*

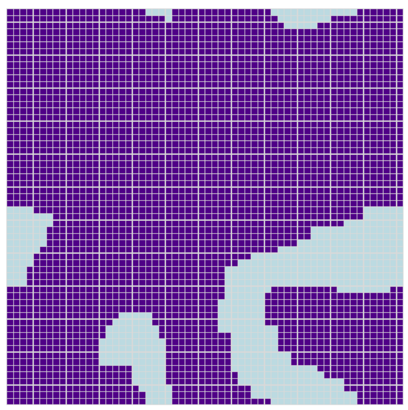

High β (low temperature)

The Potts Model

Low β

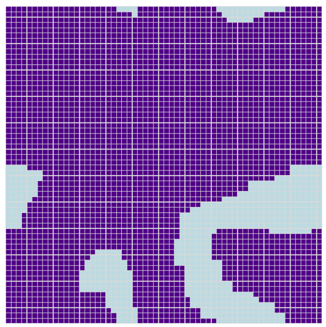

High β

Physical intuition: high temperature is “gas-like” and low temperature is “solid-like.” On what graphs is this the truth?

The Potts Model—Structural Results

Structural Idea (Carlson-Davies-Fraiman-Kolla-Potukuchi-Y. '22)

For d large enough, and $q \geq d^{27}$, the Potts model on d-regular 2-expander graphs exhibits a "structural phase transition."

Similar results were previously known for:

- random d-regular graphs (Helmuth–Jenssen–Perkins '21+),
- $\mathbb{Z}^d$ (Borgs–Chayes–Helmuth–Perkins–Tetali '20),
- d-regular $\Omega(d)$-expander graphs when $q \geq d^{\Omega(d)}$ (Jenssen–Keevash–Perkins '20)

The Potts Model—Structural Results

Structural Idea (Carlson-Davies-Fraiman-Kolla-Potukuchi-Y. '22)

For d large enough, and $q \geq d^{27}$, the Potts model on d-regular 2-expander graphs exhibits a "structural phase transition."

Similar results were previously known for:

- random d-regular graphs (Helmuth–Jenssen–Perkins '21+),
- $\mathbb{Z}^d$ (Borgs–Chayes–Helmuth–Perkins–Tetali '20),
- d-regular $\Omega(d)$-expander graphs when $q \geq d^{\Omega(d)}$ (Jenssen–Keevash–Perkins '20)

The Potts Model—Structural Results

Structural Theorem (CDFKPY '22)

Fix $\epsilon > 0$. Let $d \geq d_0(\epsilon)$, $q \geq d^{27}$, and G be a d-regular n-vertex graph with edge expansion at least 2.

Then the following hold for the q-color Potts model on G with high probability:

When $\beta < (1-\epsilon)\beta_0$, each color class has size $< (1+o(1))\frac{n}{q}$.

When $\beta > (1+\epsilon)\beta_0$, there is a color class with $> (1-o(1))n$ vertices.

The Potts Model—Structural Results

Structural Theorem (CDFKPY '22)

Fix $\epsilon > 0$. Let $d \geq d_0(\epsilon)$, $q \geq d^{27}$, and G be a d-regular n-vertex graph with edge expansion at least 2.

Then the following hold for the q-color Potts model on G with high probability:

When $\beta < (1-\epsilon)\beta_0$, each color class has size $< (1+o(1))\frac{n}{q}$.

When $\beta > (1+\epsilon)\beta_0$, there is a color class with $> (1-o(1))n$ vertices.

*G is an η-expander if $A \subset V$, $|A| \leq \frac{n}{2}$ implies $|\nabla(A)| \geq \eta|A|$.

The Potts Model—Structural Results

Structural Theorem (CDFKPY '22)

Fix $\epsilon > 0$. Let $d \geq d_0(\epsilon)$, $q \geq d^{27}$, and G be a d-regular n-vertex graph with edge expansion at least 2.

Then the following hold for the q-color Potts model on G with high probability:

When $\beta < (1-\epsilon)\beta_0$, each color class has size $< (1+o(1))\frac{n}{q}$.

When $\beta > (1+\epsilon)\beta_0$, there is a color class with $> (1-o(1))n$ vertices.

*G is an η-expander if $A \subset V$, $|A| \leq \frac{n}{2}$ implies $|\nabla(A)| \geq \eta|A|$.

β_0 is the *order-disorder threshold* of the Potts model on the random d-regular graph ($\beta_0 \sim \frac{2\ln q}{d}$)

Proof: A Combinatorial Tool

We want to say that most of the weight in the partition function comes from mostly-monochromatic colorings. To do this, we provide an upper bound on the weight coming from colorings that are far from monochromatic.

A Coloring Lemma

The number of q-colorings with exactly ℓ non-monochromatic edges is at most $\binom{n}{2\ell/d} q^{2\ell/d}$.

Proof technique: an adaptation of Karger's randomized algorithm for computing *min-cuts* of a graph.

Proof: A Combinatorial Tool

We want to say that most of the weight in the partition function comes from mostly-monochromatic colorings. To do this, we provide an upper bound on the weight coming from colorings that are far from monochromatic.

A Coloring Lemma

The number of q-colorings with exactly ℓ non-monochromatic edges is at most $\binom{n}{2\ell/d} q^{2\ell/d}$.

Proof technique: an adaptation of Karger's randomized algorithm for computing *min-cuts* of a graph.

Algorithms for Counting and Sampling

Big Question

Does there exist an efficient approximate sampling algorithm (FPTAS) for a given statistical physics model?

For many models, this is equivalent to finding efficient algorithms for approximating $Z_G(q, \beta)$.

The low-temperature Potts model lies in a class called #BIS (Goldberg–Jerrum '12, Galanis–Štefankovič–Vigoda–Yang '16). What are the hard instances?

The Potts Model—Algorithmic Results

Theorem (CDFKPY '22)

For every ϵ, there is d large enough such that for $q \geq d^{27}$, there are polynomial-time approximation algorithms for the Potts model on d-regular 2-expander graphs, for $\beta \leq (1 - \epsilon)\beta_0$ and $\beta \geq (1 + \epsilon)\beta_0$.

Moreover, the high-temperature result does not require expansion and only requires q large in terms of ϵ.

(Compare to: Jenssen–Keevash–Perkins '20 which had overlapping temperature ranges but required $q \geq d^{\Omega(d)}$ and $\Omega(d)$-expansion in our low-temperature range.)

The Potts Model—Algorithmic Results

Theorem (CDFKPY '22)

For every ϵ, there is d large enough such that for $q \geq d^{27}$, there are polynomial-time approximation algorithms for the Potts model on d-regular 2-expander graphs, for $\beta \leq (1-\epsilon)\beta_0$ and $\beta \geq (1+\epsilon)\beta_0$.

Moreover, the high-temperature result does not require expansion and only requires q large in terms of ϵ.

(Compare to: Jenssen–Keevash–Perkins '20 which had overlapping temperature ranges but required $q \geq d^{\Omega(d)}$ and $\Omega(d)$-expansion in our low-temperature range.)

The Potts Model—Algorithmic Results

Theorem (CDFKPY '22)

For every ϵ, there is d large enough such that for $q \geq d^{27}$, there are polynomial-time approximation algorithms for the Potts model on d-regular 2-expander graphs, for $\beta \leq (1-\epsilon)\beta_0$ and $\beta \geq (1+\epsilon)\beta_0$.

Moreover, the high-temperature result does not require expansion and only requires q large in terms of ϵ.

(Compare to: Jenssen–Keevash–Perkins '20 which had overlapping temperature ranges but required $q \geq d^{\Omega(d)}$ and $\Omega(d)$-expansion in our low-temperature range.)

Proof: A Statistical Physics Tool

Our second tool is called an *abstract polymer model*.

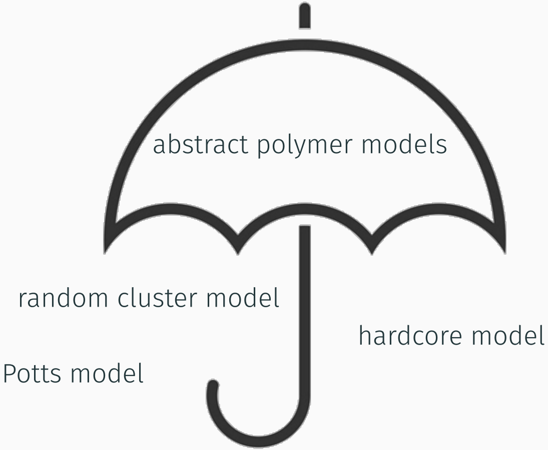

Proof: A Statistical Physics Tool

"I know what an analogy is! It's like a thought with another thought's hat on." (Britta Perry, *Community*)

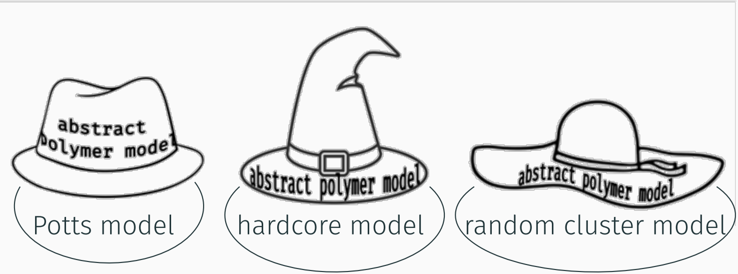

Algorithmic Proof—Low Temp

1. Characterize colorings by their “defects” from the monochromatic states.
2. Define an abstract polymer model (polymer = defect) with the same partition function Ξ as $Z_G(q, \beta)$.
3. Approximate Ξ efficiently using the *cluster expansion*.

Helmuth–Perkins–Regts ’19

There are efficient algorithms for abstract polymer models that satisfy certain conditions, including *convergence of the cluster expansion*.

The cluster expansion is simply a Taylor series expansion of the partition function!

Algorithmic Proof—Low Temp

1. Characterize colorings by their "defects" from the monochromatic states.
2. Define an abstract polymer model (polymer = defect) with the same partition function Ξ as $Z_G(q, \beta)$.
3. Approximate Ξ efficiently using the *cluster expansion*.

Helmuth–Perkins–Regts '19

There are efficient algorithms for abstract polymer models that satisfy certain conditions, including *convergence of the cluster expansion*.

The cluster expansion is simply a Taylor series expansion of the partition function!

Algorithmic Proof—Low Temp

1. Characterize colorings by their "defects" from the monochromatic states.
2. Define an abstract polymer model (polymer = defect) with the same partition function Ξ as $Z_G(q, \beta)$.
3. Approximate Ξ efficiently using the *cluster expansion*.

Helmuth–Perkins–Regts '19

There are efficient algorithms for abstract polymer models that satisfy certain conditions, including *convergence of the cluster expansion*.

The cluster expansion is simply a Taylor series expansion of the partition function!

Algorithmic Proof—Low Temp

1. Characterize colorings by their "defects" from the monochromatic states.
2. Define an abstract polymer model (polymer = defect) with the same partition function Ξ as $Z_G(q, \beta)$.
3. Approximate Ξ efficiently using the *cluster expansion*.

Helmuth–Perkins–Regts '19

There are efficient algorithms for abstract polymer models that satisfy certain conditions, including *convergence of the cluster expansion.*

The cluster expansion is simply a Taylor series expansion of the partition function!

Algorithmic Proof—Low Temp

1. Characterize colorings by their "defects" from the monochromatic states.
2. Define an abstract polymer model (polymer = defect) with the same partition function Ξ as $Z_G(q, \beta)$.
3. Approximate Ξ efficiently using the *cluster expansion*.

Helmuth–Perkins–Regts '19

There are efficient algorithms for abstract polymer models that satisfy certain conditions, including *convergence of the cluster expansion.*

The cluster expansion is simply a Taylor series expansion of the partition function!

Further Directions

- Extend arguments to the ϵ gap around β_0 and to the random cluster model.
- Constant q?
- Weaker expansion

plus future projects in the intersections of statistical physics and combinatorics as I continue on to a postdoc at Georgia Tech!

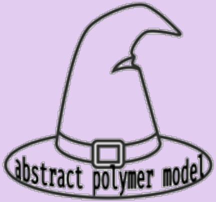

Thank you!