Sums of GUE matrices and concentration of hives from correlation decay of eigengaps

Hariharan Narayanan (TIFR)   Scott Sheffield (MIT)
Terence Tao (UCLA)

Combinatorics, Geometry and Representation Theory
ICTS
November 17 – 28, 2025

Hives

$\lambda = (40, 30, 20, 10)$

$\mu = (40, 30, 20, 10)$

$\nu = (65, 55, 45, 35)$

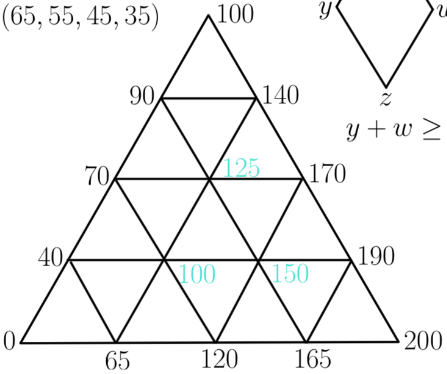

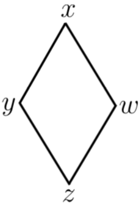

$y + w \geq x + z$

Hives

- Knutson and Tao invented hives and used them to prove the saturation conjecture for Littlewood-Richardson coefficients. Together with the work of Klyachko, this implied the Horn conjecture, which gave a description of the possible spectra of $X + Y$ when the spectrum of the Hermitian matrices X and Y are fixed n-tuples of real numbers.

Hives: relation to honeycombs

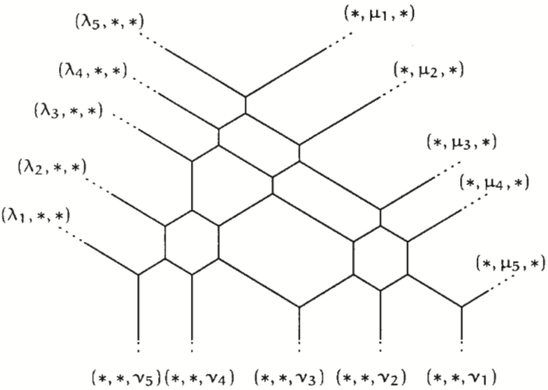

Figure: Knutson and Tao (1999).

Hives: appearance in representation theory

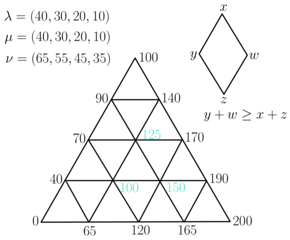

- When λ_n, μ_n and ν_n correspond to partitions, the Littlewood-Richardson coefficient $c^{\nu_n}_{\lambda_n \mu_n}$ is the multiplicity with which the irreducible representation V_{ν_n} of $GL_n(\mathbb{C})$ appears in the direct sum decomposition of the tensor product $V_{\lambda_n} \otimes V_{\mu_n}$. It equals the number of hives with integer coordinates and the appropriate boundary conditions.

Hives: appearance in representation theory

- The saturation conjecture (now a theorem due to Knutson and Tao) states that if for some positive k, $c^{k\nu}_{k\lambda\, k\mu} \geq 1$, then $c^{\nu}_{\lambda\mu} \geq 1$.

Littlewood-Richardson coefficients and square-triangle tilings

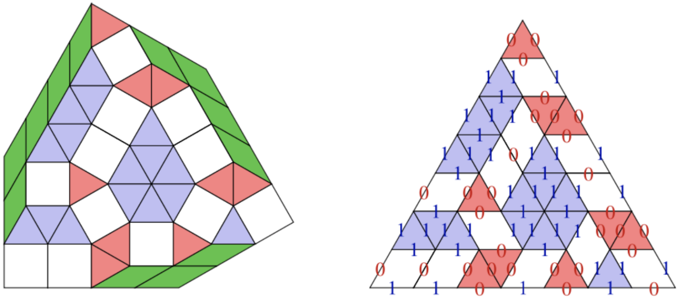

Figure: Image source: Kevin Purbhoo, Puzzles, tableaux, and mosaics

The bijection between honeycombs and Littlewood-Richardson skew tableaux

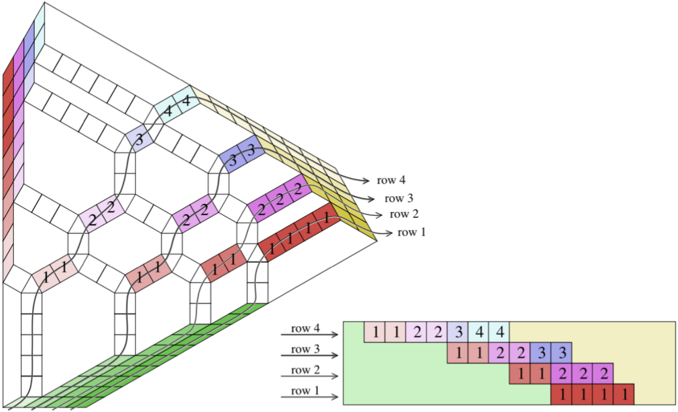

Figure: Image source: Kevin Purbhoo, Puzzles, tableaux, and mosaics

Let us define a relation

$$\lambda \boxplus \mu \to \nu \tag{1}$$

if there exist Hermitian matrices A, B with eigenvalues λ, μ respectively such that $A + B$ has eigenvalues ν.
Weyl asked the question of determining necessary and sufficient conditions on $\lambda, \mu, \nu \in \texttt{Spec}$ for the relation (1) to hold.

As conjectured by Horn and proved by Klyachko, Knutson and Tao, the set

$$\mathtt{HORN}_{\lambda \boxplus \mu} := \{\nu \in \mathtt{Spec} : \lambda \boxplus \mu \to \nu\}$$

of possible ν arising from a given choice of λ, μ forms a polytope (known as the *Horn polytope*), given by the trace condition

$$\sum \lambda + \sum \mu = \sum \nu$$

together with a recursively defined set of linear inequalities known as the *Horn inequalities*.

Introduction

There is a natural probability measure on the Horn polytope $\mathtt{HORN}_{\lambda \boxplus \mu}$, referred to as the *Horn probability measure*, defined as the eigenvalues of $A + B$ when A, B are chosen independently and uniformly from the space of all Hermitian matrices with eigenvalues λ, μ respectively.

This Horn measure turns out to be piecewise polynomial and was computed explicitly by Coquereaux and Zuber to be given by the formula

$$\frac{V(\nu)V(\tau)}{V(\lambda)V(\mu)}|\mathtt{HIVE}_{\lambda\boxplus\mu\to\nu}|\, d\nu$$

for $\lambda, \mu \in \mathtt{Spec}^{\circ}$, where

$$V(\lambda) = V_n(\lambda) := \prod_{1\leq i<j\leq n} (\lambda_i - \lambda_j)$$

is the Vandermonde determinant, and τ is the tuple

$$\tau := (n, n-1, \ldots, 1).$$

Introduction

We introduce the relation

$$\mathrm{diag}(\lambda) \to a$$

for $\lambda \in \mathtt{Spec}$ and $a \in \mathbb{R}^n$ to denote the claim that there exists a Hermitian matrix A with eigenvalues λ and diagonal entries $a_1, \ldots, a_n$.

Introduction

The classical *Schur–Horn theorem* asserts that the relation $\mathrm{diag}(\lambda) \to a$ holds if and only if a is *majorized* by λ in the sense that one has the trace condition

$$\sum a = \sum \lambda$$

and the majorizing inequalities

$$a_{i_1} + \cdots + a_{i_k} \leq \lambda_1 + \cdots + \lambda_k$$

for all $1 \leq i_1 < \cdots < i_k \leq n$; equivalently, a lies in the *permutahedron* formed by the convex hull of the image of λ under the permutation group S_n.

Introduction

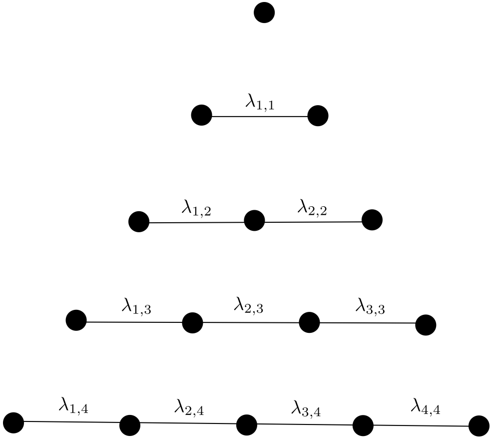

Figure: An $n = 4$ Gelfand–Tsetlin pattern. Each number $\lambda_{i,j}$ in the pattern is greater than or equal to numbers immediately to the northeast or southeast of the pattern; in particular, every row of the pattern is decreasing. Note that such patterns are sometimes depicted as inverted pyramids instead of pyramids in the literature.

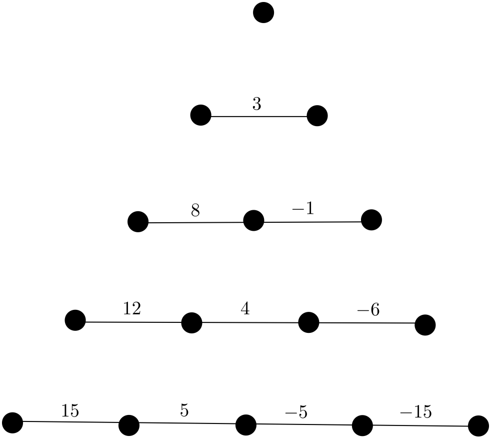

Figure: A Gelfand–Tsetlin pattern with boundary $\mathrm{diag}(15, 5, -5, -15) \to (3, 4, 3, -10)$.

Introduction

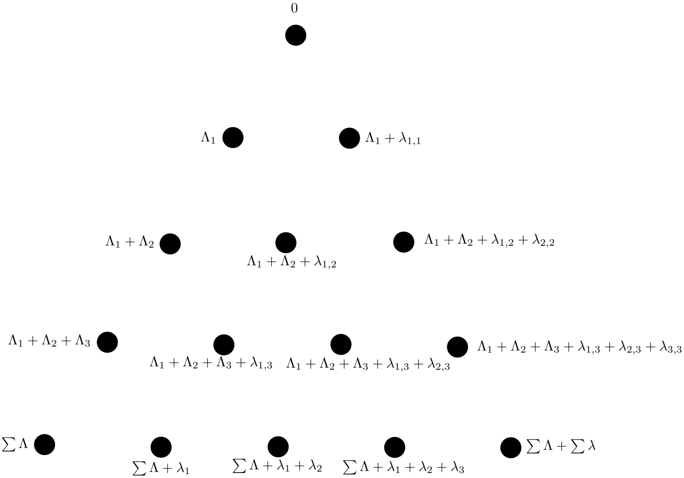

Figure: The hive associated with the Gelfand–Tsetlin pattern in an earlier figure and some large gap tuple Λ.

Introduction

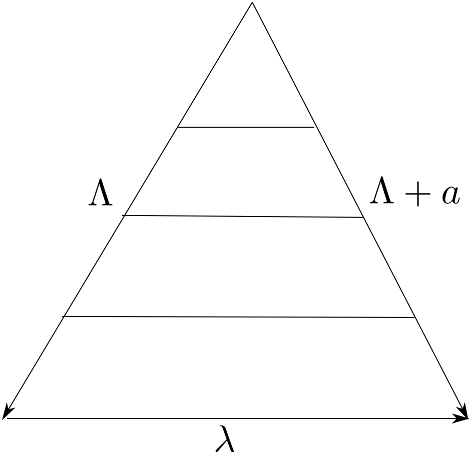

Figure: A schematic depiction of the boundary conditions of the hive in the previous figure. The horizontal “creases” inside the triangle indicate that the rhombus concavity condition is essentially an automatic consequence of the large gaps hypothesis for rhombi that cross these creases.

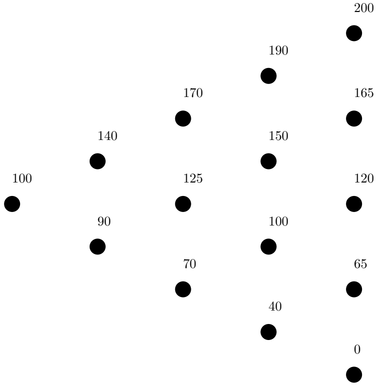

Figure: A hive with boundary condition $(40, 30, 20, 10) \boxplus (40, 30, 20, 10) \to (65, 55, 45, 35)$.

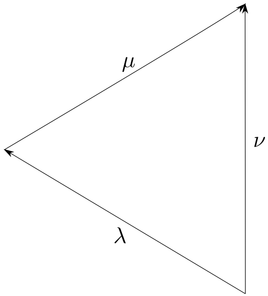

If $\lambda, \mu, \nu \in \texttt{Spec}$, we say that a hive h has boundary condition $\lambda \boxplus \mu \to \nu$ if one has $h(0, i) = \sum_{j=1}^{i} \lambda_j$, and $h(i, n) = \sum \lambda + \sum_{j=1}^{i} \mu_j$, and $h(i, i) = \sum_{j=1}^{i} \nu_j$ for all $0 \leq i \leq n$, and write $\texttt{HIVE}_{\lambda \boxplus \mu \to \nu}$ for the set of all hives with boundary condition $\lambda \boxplus \mu \to \nu$.

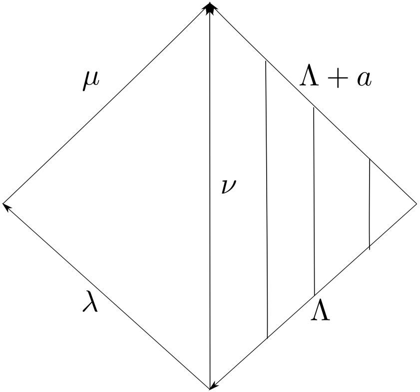

Figure: A schematic depiction of an augmented hive in $\mathtt{AUGHIVE}_{\mathrm{diag}(\lambda \boxplus \mu \to \nu) \to a}$, where we artificially shift by a tuple Λ with large gaps in order to create two hives, instead of a hive and a Gelfand–Tsetlin pattern.

Introduction

It is natural to ask if Lebesgue measure on the polytope $\mathtt{AUGHIVE}_{\mathrm{diag}(\lambda \boxplus \mu \to *) \to *}$ also exhibits concentration. As a first step towards this goal, we are able to establish this for spectra λ, μ that are not deterministic, but are instead drawn from (scalar multiples of) the *GUE ensemble*.

Introduction

To establish normalization conventions, we define a GUE random matrix to be a random Hermitian matrix $M = (\xi_{ij})_{1 \le i,j \le n}$ where $\xi_{ij} = \overline{\xi_{ji}}$ for $i < j$ are independent complex gaussians of mean zero and variance 1, ξ_{ii} are independent real gaussians of mean zero and variance 1, independent of the ξ_{ij} for $i < j$.

As is well known, if $\sigma > 0$ and A is a random matrix with $\frac{A}{\sqrt{\sigma^2 n}}$ drawn from the GUE ensemble, then the eigenvalues $\lambda \in \mathtt{Spec}$ of A are distributed with probability density function

$$C_n \sigma^{-n^2} \exp\left(-\frac{|\lambda|^2}{2\sigma^2 n}\right) V(\lambda)^2$$

for some constant $C_n > 0$ depending only on n.

If $\sigma_\lambda, \sigma_\mu > 0$ are fixed and A, B are independent random matrices with $\frac{A}{\sqrt{\sigma_\lambda^2 n}}, \frac{B}{\sqrt{\sigma_\mu^2 n}}$ drawn from the GUE ensemble, then the the distribution of the eigenvalues of $A + B$ are the pushforward of the measure on the $n(n+1)$-dimensional *augmented hive cone*

$$\mathtt{AUGHIVE}_{\mathrm{diag}(*\boxplus*\to*)\to*} := \bigcup_{\lambda,\mu,\nu,a} \left(\mathtt{HIVE}_{\lambda\boxplus\mu\to\nu} \times \mathtt{GT}_{\mathrm{diag}(\nu)\to a}\right).$$

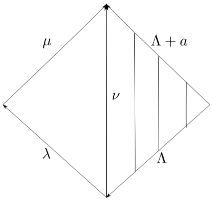

Here the probability density function of this measure is given by

$$C_{n,\sigma_\lambda,\sigma_\mu} \exp\left(-\frac{|\lambda|^2}{2\sigma_\lambda^2 n} - \frac{|\mu|^2}{2\sigma_\mu^2 n}\right) V(\lambda) V(\mu) \qquad (2)$$

on the slices

$$\mathtt{AUGHIVE}_{\mathrm{diag}(\lambda \boxplus \mu \to *) \to *} := \bigcup_{\nu, a} \left(\mathtt{HIVE}_{\lambda \boxplus \mu \to \nu} \times \mathtt{GT}_{\mathrm{diag}(\nu) \to a}\right).$$

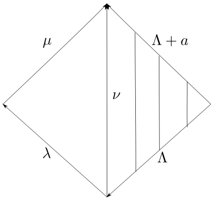

The main result of this talk is the following concentration bound.

Theorem (Concentration of augmented hives)

Let $\sigma_\lambda, \sigma_\mu > 0$ *be fixed, and let* $(h, \gamma) \in \mathtt{AUGHIVE}_{\mathrm{diag}(*\boxplus*\to*)\to*}$ *be a random augmented hive (with GUE boundary data). Then for all* $v \in T$*, we have the variance bound*

$$\operatorname{var} h(v) = o(n^4)$$

as $n \to \infty$*, uniformly in* v*.*

Informally, this theorem asserts that randomly selected hives have an asymptotic limiting profile, at least in a subsequential sense.

The hive part of an augmented hive with GUE boundary conditions for $n = 250$

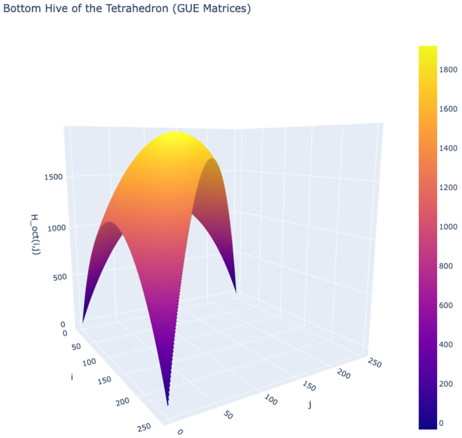

Figure: Simulation by Aalok Gangopadhyay

Relation to previous work

In the work of N.-Sheffield'21, a large deviation principle was proved for (the hive part of) augmented hives in the L^∞ metric with fixed boundaries that were C^2 and strongly concave. The result of this talk does not directly follow; whether uniqueness holds for the rate function minimizer in that large deviation principle is yet unknown.

Methods of proof

The first step is to exploit the *octahedron recurrence*, which has appeared in the enumerative combinatorics literature several times, which was observed by Knutson, Tao and Woodward to have an "associativity" property

$$\bigcup_{\nu} \mathtt{HIVE}_{\lambda \boxplus \mu \to \nu} \times \mathtt{HIVE}_{\gamma \boxplus \nu \to \pi} \equiv \bigcup_{\sigma} \mathtt{HIVE}_{\gamma \boxplus \lambda \to \sigma} \times \mathtt{HIVE}_{\sigma \boxplus \mu \to \pi}$$

on hives related to the trivial associativity

$$(A + B) + C = A + (B + C)$$

of the addition operation on Hermitian matrices A, B, C.

Methods of proof: The octahedron recurrence

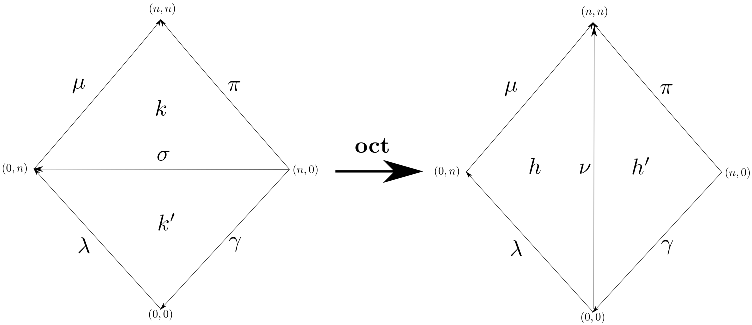

Figure: A schematic depiction of the octahedron recurrence that transforms one pair $(k, k') \in \texttt{HIVE}_{\sigma\boxplus\mu\to\pi} \times \texttt{HIVE}_{\gamma\boxplus\lambda\to\sigma}$ of hives into another $(h, h') \in \texttt{HIVE}_{\lambda\boxplus\mu\to\nu} \times \texttt{HIVE}_{\gamma\boxplus\nu\to\pi}$. The hives h, h', k, k' have been shifted to lie on triangles T, T', U, U' respectively.

Methods of proof

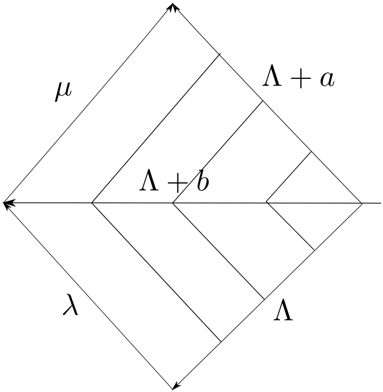

Figure: A schematic depiction of a pair in $\mathrm{GT}_{\mathrm{diag}(\lambda)\to b} \times \mathrm{GT}_{\mathrm{diag}(\mu)\to a-b}$, where an artificial shift by a tuple Λ with large gaps is used to re-interpret this pair as a pair of hives with a common edge.

Methods of proof

In our context (viewing Gelfand–Tsetlin patterns as degenerations of hives), the octahedron recurrence is a piecewise-linear volume-preserving bijection

$$\textbf{oct}: \mathtt{GT}_{\mathrm{diag}(*)\to *} \times \mathtt{GT}_{\mathrm{diag}(*)\to *} \to \mathtt{AUGHIVE}_{\mathrm{diag}(*\boxplus * \to *)\to *}$$

between the two $n(n-1)$-dimensional convex cones $\mathtt{GT}_{\mathrm{diag}(*)\to *} \times \mathtt{GT}_{\mathrm{diag}(*)\to *}$, $\mathtt{AUGHIVE}_{\mathrm{diag}(*\boxplus * \to *)\to *}$.
In fact **oct** is a piecewise-linear volume-preserving bijection

$$\textbf{oct}: \bigcup_b \mathtt{GT}_{\mathrm{diag}(\lambda)\to b} \times \mathtt{GT}_{\mathrm{diag}(\mu)\to a-b} \to \bigcup_\nu \mathtt{AUGHIVE}_{\mathrm{diag}(\lambda\boxplus\mu\to\nu)\to a}$$

for any $\lambda, \mu \in \mathtt{Spec}^\circ$ and $a \in \mathbb{R}^d$.

Theorem (Excavation form of octahedron recurrence)

Let v be an element of the triangle T. Then there is an explicit finite family $\mathcal{W}_v \colon \mathtt{GT}_{\mathrm{diag}(*)\to *} \times \mathtt{GT}_{\mathrm{diag}(*)\to *} \to \mathbb{R}$ of linear functionals on $\mathtt{GT}_{\mathrm{diag}(*)\to *} \times \mathtt{GT}_{\mathrm{diag}(*)\to *}$, such that whenever $(h, g) = \mathbf{oct}(g_1, g_2)$ is the image of the octahedron recurrence for some $g_1, g_2 \in \mathtt{GT}_{\mathrm{diag}(*)\to *}$, then

$$h(v) = \max_{w \in \mathcal{W}} w(g_1, g_2).$$

The linear functionals are given in terms of lozenge tilings of a certain hexagon $\hexagon_v$ associated to v.

The octahedron recurrence

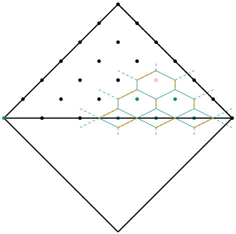

(a) Perfect matching

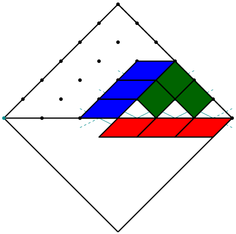

(b) Lozenge (and triangle) tiling

Figure: Correspondence between perfect matchings and lozenge and triangle tilings

The octahedron recurrence

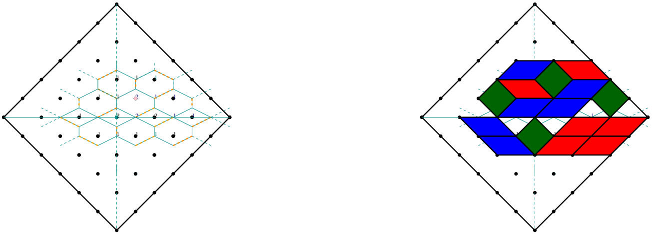

(a) Perfect matching (b) Lozenge (and triangle) tiling

Figure: Correspondence between perfect matchings and lozenge and triangle tilings

The octahedron recurrence

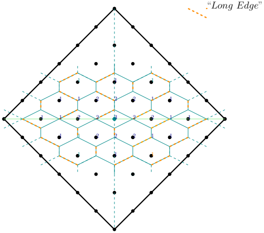

(a) Perfect matching

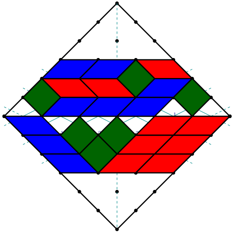

(b) Lozenge (and triangle) tiling

Figure: Correspondence between perfect matchings and lozenge and triangle tilings

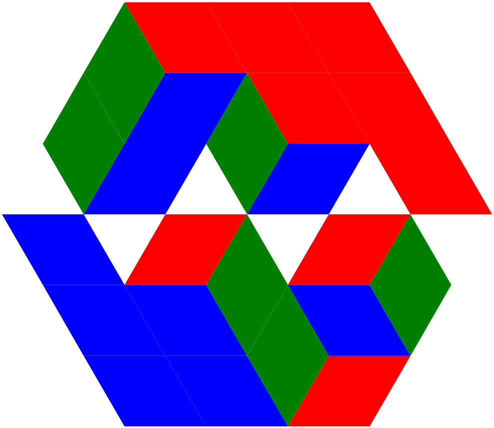

Figure: $n = 6$, Simulation by Aalok Gangopadhyay

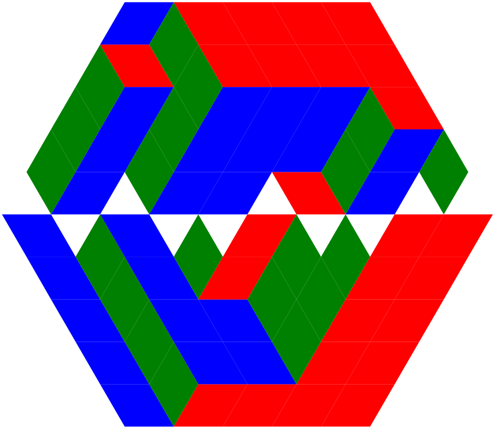

Figure: $n = 10$, Simulation by Aalok Gangopadhyay

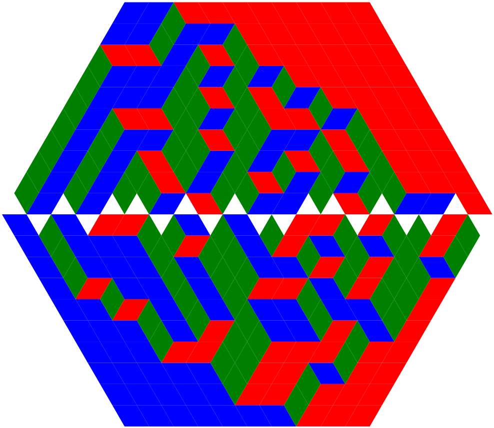

Figure: $n = 20$, Simulation by Aalok Gangopadhyay

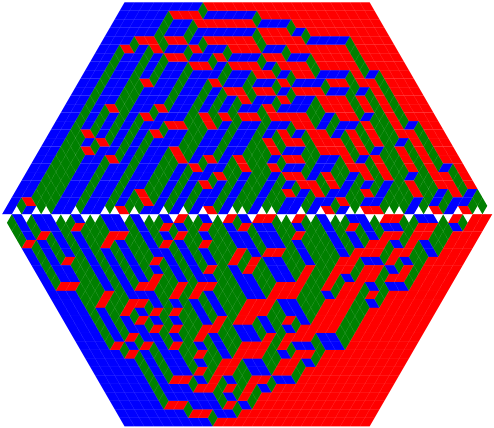

Figure: $n = 50$, Simulation by Aalok Gangopadhyay

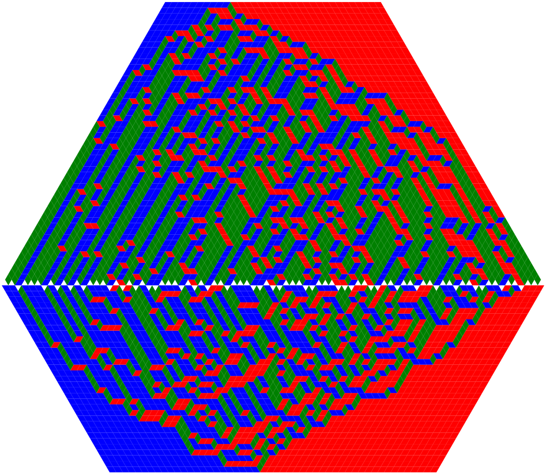

Figure: $n = 100$, Simulation by Aalok Gangopadhyay

Methods of proof: Recent progress on the KLS conjecture

Direct calculation reveals that the density function on $\mathrm{GT}_{\mathrm{diag}(*)\to *} \times \mathrm{GT}_{\mathrm{diag}(*)\to *}$ is log-concave. The supremum in the last Theorem can then be handled by the following tool:

Lemma

Let η be an log-concave probability measure in $\mathbb{R}^d$, and let $\mathcal{W}$ be a family of affine functions $w\colon \mathbb{R}^d \to \mathbb{R}$. Then

$$\operatorname{var}_\eta \left(\sup_{w \in \mathcal{W}_v} w \right) \ll \sup_{w \in \mathcal{W}} (\operatorname{var}_\eta w) \log(2 + d).$$

This lemma is a consequence of Cheeger's inequality and recent work of Klartag on the KLS conjecture.

Methods of proof

In view of this lemma, it would now suffice to establish the variance bound

$$\operatorname{var} w(\gamma_1, \gamma_2) = O(n^{4-c})$$

for all $v \in T$ and $w \in \mathcal{W}_v$ and some constant $c > 0$, where (γ_1, γ_2) was the random variable

$$(\gamma_1, \gamma_2) = \left((\lambda_{j,k})_{1 \le j \le k \le n}, (\mu_{j,k})_{1 \le j \le k \le n}\right);$$

the additional factor of n^{-c} is needed to overcome the logarithmic loss in the preceding Lemma. This is a variance estimate for linear statistics of the GUE minor process.

Methods of proof

As it turns out, certain covariance estimates for eigenvalue gaps of GUE established by Cipolloni, Erdős and Schröder, combined with some further manipulations from the theory of determinantal processes to analyze the minor process, are *almost* enough to obtain this sort of bound; there is however a technical difficulty because these bounds are only established in the bulk of the spectrum and not on the edge. However, the contributions coming from the edge region can be controlled by relatively crude estimates, and after removing these contributions to focus on the bulk contribution the above strategy can be made to work.

Subsequent work: Asymptotic Height Functions and S_v

- Let AHT_v^∞ denote the set of asymptotic height function pairs $f = (f^{\mathrm{up}}, f^{\mathrm{lo}})$ on the limiting hexagon $\hexagon_v^\infty$, satisfying:
  - Lipschitz conditions with $\nabla f \in K$ a.e.,
  - prescribed boundary values coming from the GUE scaling limit,
  - linearity of $f^{\mathrm{up}} + f^{\mathrm{lo}}$ on the overlap strip with slope 1.
- An idea of Sheffield: Using $\sigma_\diamond$ and equator/hexagon corrections, define

$$S_v(f) = S_{v,\diamond}(f) + S_{v,\Delta}(f) + S_{v,\hexagon},$$

where:
  - $S_{v,\diamond}(f)$ is the bulk integral of $\sigma_\diamond(\rho(x), \partial f(x))$,
  - $S_{v,\Delta}(f)$ is the contribution from the equator line,
  - $S_{v,\hexagon}$ is the contribution of the hexagon boundary itself.
- This functional arises as the scaling limit of expected lozenge weights.

Subsequent work

Relying heavily on results of Tao and Fefferman, it is proved in [N'25]
"On the limit of random hives with GUE boundary conditions" that if a sequence of vertices v_n with $v_n/n \to v$ in the interior of the triangle, then the following holds.

Theorem:
The limit

$$\lim_{n\to\infty} \frac{1}{n^2}\,\mathbb{E}\left[h_n(v_n)\right]$$

exists, and equals

$$\sup_{f\in \mathrm{AHT}_v^\infty} S_v(f).$$

- Thus there is a unique continuum hive limit in the GUE setting, described via a variational principle over asymptotic height functions.

Open questions

1. What can be said about the concentration of random real valued augmented hives with general boundary conditions? If they do concentrate, what are the possible subsequential limit shapes? In particular, is the limit unique? In the limit when one of the boundary conditions is more spread out than the other, the limit shape should essentially degenerate to fractional free convolution powers (See Shlyakhtenko and Tao).
2. Do the local statistics of the random augmented GUE hive process converge (either in the bulk or the edge) to a known limit? In the case of the random Gelfand–Tsetlin process, the limit is known to essentially be the Boutillier bead process.
3. Do random integer valued augmented hives with general boundary conditions concentrate? Again, if they do concentrate, what are the possible subsequential limit shapes? In particular, is the limit unique?

Thank you for your attention!

Details of analysis

Poincaré inequalities over log-concave measures

We establish a useful Poincaré inequality over log-concave measures. Namely, we show

Proposition (Poincaré inequality on log-concave measures)

Let η be an log-concave probability measure in $\mathbb{R}^d$ with finite second moments, and define the $d \times d$ inertia matrix M by the formula

$$M := \mathbb{E}_\eta x x^T - (\mathbb{E}_\eta x)(\mathbb{E}_\eta x)^T.$$

Then for any Lipschitz function $f \colon \mathbb{R}^d \to \mathbb{R}$, one has

$$\operatorname{var}_\eta f \ll \left(\mathbb{E}_\eta |\nabla f|^2\right) \|M\|_{\mathrm{op}} \log(2+d).$$

Poincaré inequalities over log-concave measures

Without loss of generality, we may assume that η is an *isotropic measure* in the sense that

$$\mathbb{E}_\eta x = 0; \quad \mathbb{E}_\eta x x^T = I_d.$$

Define the *Cheeger constant* $D_{\mathrm{Che}}(\eta)$ of η (with respect to the Euclidean inner product) by the formula

$$D_{\mathrm{Che}}(\eta) := \inf_{A \subset \mathbb{R}^d} \frac{\int_{\partial A} \rho}{\min(\eta(A), 1 - \eta(A))}$$

where the infimum runs over all open subsets A of $\mathbb{R}^d$ with smooth boundary with $0 < \eta(A) < 1$, and ∂_A is integrated using surface measure.

Poincaré inequalities over log-concave measures

Proof.

By the Cheeger inequality, one has the Poincaré inequality

$$D_{\mathrm{Che}}(\eta)^2 \operatorname{var}_\eta f \ll \mathbb{E}_\eta |\nabla f|^2$$

so the task reduces to (and is in fact equivalent to) the lower bound

$$D_{\mathrm{Che}}(\eta) \gg \frac{1}{\sqrt{\log(2+d)}}$$

on the Cheeger constant of an isotropic log-concave measure.

But this follows from recent work of Klartag (building upon previous advances by Chen, Klartag-Lehec and Jambalapati-Lee-Vempala.) □

Proposition (Weighted Poincaré inequality on log-concave measures)

Let η be an log-concave probability measure in $\mathbb{R}^d$ with finite second moments. Express $\mathbb{R}^d$ as a Cartesian product $\mathbb{R}^{d_1} \times \cdots \times \mathbb{R}^{d_k}$ for some $d_1, \ldots, d_k$ summing to d (so that a vector $x \in \mathbb{R}^d$ is expressed as $(x_1, \ldots, x_k)$ for $x_j \in \mathbb{R}^{d_j}$), and for each $i = 1, \ldots, k$, and define the $d_j \times d_j$ inertia matrix M_j by the formula

$$M_j := \mathbb{E}_\eta x_j x_j^T - (\mathbb{E}_\eta x_j)(\mathbb{E}_\eta x_j)^T.$$

Then for any Lipschitz function $f\colon \mathbb{R}^d \to \mathbb{R}$, and any weights $\alpha_1, \ldots, \alpha_k > 0$, one has

$$\mathrm{var}_\eta f \ll \left(\mathbb{E}_\eta \sum_{j=1}^k \alpha_j |\nabla_j f|^2\right) \left(\sum_{j=1}^k \alpha_j^{-2} \|M_j\|_{\mathrm{op}}^2\right)^{1/2} \log(2+d).$$

By pushing forward η by the map $(x_1, \ldots, x_k) \mapsto (\alpha_1^{1/2} x_1, \ldots, \alpha_k^{1/2} x_k)$ we may normalize $\alpha_j = 1$ for all j, so that $\sum_{j=1}^k \alpha_j |\nabla_j f|^2 = |\nabla f|^2$. By the previous proposition, it thus suffices to establish the bound

$$\|M\|_{\mathrm{op}} \leq \left(\sum_{j=1}^k \|M_j\|_{\mathrm{op}}^2 \right)^{1/2}.$$

Poincaré inequalities over log-concave measures

For any $x = (x_1, \ldots, x_k) \in \mathbb{R}^d$, where $x_j \in \mathbb{R}^{d_j}$ for all j, it follows from the positive semi-definiteness of M and the triangle inequality followed by Cauchy–Schwarz that

$$\begin{aligned}
(x^T M x)^{1/2} &\leq \sum_{j=1}^{k} (x_j^T M_j x_j)^{1/2} \\
&\leq \sum_{j=1}^{k} \|M_j\|_{\mathrm{op}} |x_j| \\
&\leq \left(\sum_{j=1}^{k} \|M_j\|_{\mathrm{op}}^2 \right)^{1/2} |x|
\end{aligned}$$

and the claim follows.

□

The octahedron recurrence

We identify a pair $(k, k') \in \mathtt{HIVE}_{\sigma \boxplus \mu \to \pi} \times \mathtt{HIVE}_{\gamma \boxplus \lambda \to \sigma}$ with a single function $\tilde{k}\colon U \cup U' \to \mathbb{R}$ defined on the square. Even though $T \cup T'$ and $U \cup U'$ are both technically equal to the same set $\{0, \ldots, n\}^2$, it is conceptually better to think of these sets as being distinct (except on the boundary).

The octahedron recurrence

We will view these two copies of $\{0, \ldots, n\}^2$ as the upper and lower faces respectively of a certain tetrahedron `tet`, and the octahedron recurrence **oct** can be constructed by "excavating" that tetrahedron.

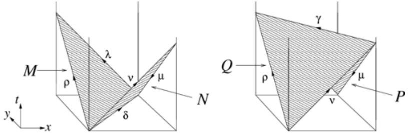

Figure: Lower and upper panels of `tet`.
Image credits, Henriques and Kamnitzer

The octahedron recurrence

For vertices $v = (i, j)$ in the interior or $\{0, \ldots, n\}^2$, the octahedron recurrence specifying $\tilde{h}(i, j)$ was initially defined by recursively “excavating” a real-valued function on a tetrahedron $\{(a, b, c, d) \in \mathbb{Z}^4 : a, b, c, d \geq 0; a + b + c + d = n\}$ with $\tilde{k}$ describing the values on the top two faces, and $\tilde{h}$ the bottom two faces.

The octahedron recurrence

An alternate description was given by Speyer, in terms of perfect matchings of an "excavation graph" associated to (i, j). We will use a modification of Speyer's formula that is more convenient for our purposes, in which the perfect matchings are replaced by *lozenge tilings*.

The octahedron recurrence

Given a lozenge $\diamond = ABCD$ and a function $\tilde{k}\colon \{0,\ldots,n\}^2 \to \mathbb{R}$ defined as before, we define the *weight* $\mathbf{wt}(\diamond) = \mathbf{wt}(\diamond, \tilde{k})$ to be the quantity

$$\mathbf{wt}(\diamond) := \frac{1}{3}(\tilde{k}(A) + \tilde{k}(C) - \tilde{k}(B) - \tilde{k}(D)).$$

Similarly, given a border triangle $\Delta = ABC$, the weight $\mathbf{wt}(\Delta) = \mathbf{wt}(\tau, \tilde{k})$ is defined as

$$\mathbf{wt}(\Delta) := \frac{1}{3}(\tilde{k}(B) - \tilde{k}(A)).$$

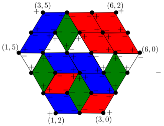

.

Figure: A typical lozenge tiling of $\hexagon_{(3,2)}$, $n = 6$.

The original point $v = (i, j)$ is then the intersection of the diagonals BE and CF. The line AD is called the *equator*; it lies on the border between U and U'.

Definition

The *weight* $\mathbf{wt}(\hexagon_v) = \mathbf{wt}(\hexagon_v, \tilde{k})$ of this hexagon is defined as

$$\mathbf{wt}(\hexagon_v) := \frac{1}{3}(\tilde{k}(B) + \tilde{k}(C) - \tilde{k}(D) + \tilde{k}(E) + \tilde{k}(F)). \qquad (3)$$

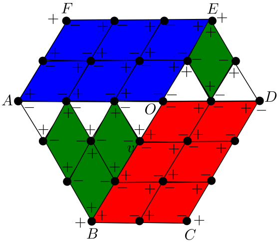

Definition

The *weight* $w_\Xi = w_\Xi(\tilde{k})$ of such a tiling is defined to be the sum of the weights of all the lozenges $\diamond$ and triangles Δ in the tiling, as well as the weight of the entire hexagon $\hexagon_v$:

$$w_\Xi := \sum_{\diamond \in \Xi} \mathbf{wt}(\diamond) + \sum_{\Delta \in \Xi} \mathbf{wt}(\Delta) + \mathbf{wt}(\hexagon_v). \tag{4}$$

Note that the w_Ξ depend linearly on $\tilde{k}$, and hence on k, k'. We then define

$$\tilde{h}(v) := \max_{\Xi \text{ tiles } \hexagon_v} w_\Xi$$

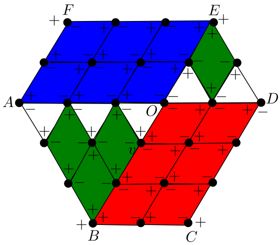

Lemma (Replacing red lozenges with blue)

Let v be an interior point of $\{0, \ldots, n\}^2$, and let Ξ be a lozenge tiling of $\hexagon_v = ABCDEF$. Then one has the identity

$$\sum_{\diamond \in \Xi,\ red} \mathbf{wt}(\diamond) - \sum_{\diamond \in \Xi,\ blue} \mathbf{wt}(\diamond) =$$

$$\frac{1}{3}(-\tilde{k}(A) + \tilde{k}(B) - \tilde{k}(C) + \tilde{k}(D) - \tilde{k}(E) + \tilde{k}(F)).$$

The proof of the main theorem is thus reduced to

Proposition (Reduction to the minor process)

Let $\sigma_\lambda, \sigma_\mu > 0$ be fixed, and let A, B with $\frac{A}{\sqrt{\sigma_\lambda^2 n}}, \frac{B}{\sqrt{\sigma_\mu^2 n}}$ be drawn independently from the GUE ensemble, and let (g, g') be the resulting Gelfand–Tsetlin patterns. Then for any $v \in T$, we have

$$\operatorname{var} \max_{\Xi \text{ tiles } \hexagon_v} 2 \sum_{\diamond \in \Xi, \text{ blue}} \mathbf{wt}(\diamond) + \sum_{\diamond \in \Xi, \text{ green}} \mathbf{wt}(\diamond) + \sum_{\Delta \in \Xi} \mathbf{wt}(\Delta)$$

$$+\mathbf{wt}'(\hexagon_v) = o(n^4)$$

where we identify (g, g') with a pair of hives (k, k') using a large gaps tuple γ as indicated above.

Using eigenvalue rigidity to remove edge contributions

Lemma (Eigenvalue rigidity (Tao-Vu'13))

Let A be a matrix with $A/\sqrt{n}$ having the distribution of GUE. Then for any $1 \leq i \leq n$ we have

$$\mathbb{P}(n^{-1/3}\min(i, n-i+1)^{1/3}|\lambda_i - \sqrt{n}\gamma_i| \geq T) \ll n^{O(1)} exp(-cT^c)$$

for any $T > 0$ and some absolute constant $c > 0$, where the classical location γ_i is the value predicted by the semicircular law:

$$\int_{-\infty}^{\gamma_i} \frac{1}{2\pi}(4-x^2)^{1/2}\, dx = \frac{i}{n}.$$

In particular,

$$\lambda_i, \mathbb{E}\lambda_i = \sqrt{n}\gamma_i + O(n^{1/3}\min(i, n-i+1)^{-1/3}\log^{O(1)} n)$$

with overwhelming probability.

Using eigenvalue rigidity to remove edge contributions

We conclude that for any lozenge tiling Ξ of $\hexagon_v$, we have

$$\sum_{\Delta \in \Xi} \mathbf{wt}(\Delta) = \sum_{\Delta \in \Xi} \mathbb{E}\mathbf{wt}(\Delta) + O(n^{4/3} \log^{O(1)} n)$$

with overwhelming probability. The weight $\mathbf{wt}'(\hexagon_v)$ is a certain linear combination of the eigenvalues λ_i, μ_j with bounded coefficients. Similarly, we conclude that

$$\begin{aligned} \mathbf{wt}'(\hexagon_v) &= \mathbb{E}\mathbf{wt}'(\hexagon_v) + O\left(\sum_{i=1}^{n} n^{1/3} \min(i, n-i+1)^{-1/3} \log^{O(1)} n \right) \\ &= \mathbb{E}\mathbf{wt}'(\hexagon_v) + O(n \log^{O(1)} n). \end{aligned}$$

Again, the contribution of the $O(n \log^{O(1)} n)$ error is acceptable, so we may also replace $\mathbf{wt}'(\hexagon_v)$ by $\mathbb{E}\mathbf{wt}'(\hexagon_v)$.

Using eigenvalue rigidity to remove edge contributions

Let $\epsilon > 0$ be a small parameter, and let U_ϵ denote the portion of U that lies at Euclidean distance at least ϵn from the boundary of U. Define U'_ϵ similarly. Using telescoping sums, we show

$$\sum_{\diamond \not\subset U_\epsilon \cup U'_\epsilon} |\mathbf{wt}(\diamond)| \ll \epsilon^{1/3} n^2$$

with overwhelming probability, where the sum is over all blue or green lozenges in U or U' that are not contained in U_ϵ or U'_ϵ.

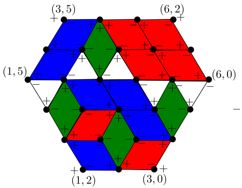

•

Henceforth we fix $\epsilon > 0$ and assume n sufficiently large depending on ϵ. By the triangle inequality, it thus suffices to establish the bound

$$\operatorname{var} \sum_{\diamond \in \Xi, \text{ blue}: \diamond \subset U'_\epsilon} \mathbf{wt}(\diamond) = O(n^{4-c+o(1)}) \tag{5}$$

and similarly with blue replaced by green, or U'_ϵ replaced by U_ϵ, or both.

We focus on establishing (5), as the other three cases are proven similarly. It suffices to establish the bound

$$\operatorname{var} \sum_{(j,k) \in \Omega} \lambda_{j,k+1} - \lambda_{j,k} = O(n^{4-c+o(1)})$$

whenever Ω is a collection of tuples of integers $1 \leq j \leq k \leq n$ with $j, k-j, n-k \gg \epsilon n$.

By the triangle inequality, it suffices to show that

$$\operatorname{var} \sum_{j \in S_k} \lambda_{j,k+1} - \lambda_{j,k} = O(n^{2-c+o(1)})$$

for each $\epsilon n \ll k \leq n-1$, where S_k is some subset of the bulk region $\{1 \leq j \leq k : j, k-j \gg \epsilon n\}$. Since the minor of a GUE matrix is a rescaled version of a GUE matrix, it suffices to establish this claim for the case $k = n-1$, that is to say (after adjusting ϵ slightly) to show that

$$\operatorname{var} X_S = O(n^{2-c+o(1)})$$

for an arbitrary subset S of $\{2\epsilon n \leq j \leq (1-2\epsilon)n\}$, where X_S denotes the random variable

$$X_S := \sum_{j \in S} \lambda_j - \lambda_{j,n-1}.$$

It is convenient to exclude a small exceptional set to keep the eigenvalues λ_j somewhat under control. From a lemma of Tao and Vu, we already know that there is a constant C_0 such that

$$|\lambda_j - \sigma_\lambda \gamma_j n^{1/2}| \leq n^{1/3} \min(j, n-j+1)^{-1/3} \log^{C_0} n \qquad (6)$$

for all $1 \leq j \leq n$ with overwhelming probability. From the Wegner estimate (see Erdős-Schlein-Yau) and enlarging C_0 if needed, we also see that

$$|\lambda_{j+1} - \lambda_j| \geq \exp(-\log^{C_0} n) \qquad (7)$$

with overwhelming probability for all $\epsilon n \leq j \leq (1-\epsilon)n$. Thus, if we let E denote the event that (6), (7) both hold for all $\epsilon n \leq j \leq (1-\epsilon)n$, then E holds with overwhelming probability; for future reference we also note the constraints (6), (7) defining E are restricting λ to a certain convex subset of `Spec`. It suffices to show that

$$\mathrm{var}(X_S|E) = O(n^{2-c+o(1)}).$$

We split this by conditioning on the spectrum λ of A. By the law of total variance (noting that the event E is measurable with respect to λ), it suffices to establish the bounds

$$\mathrm{var}(\mathbb{E}(X_S|\lambda)|E) = O(n^{2-c+o(1)}) \tag{8}$$

and

$$\mathbb{E}(\mathrm{var}(X_S|\lambda)|E) = O(n^{2-c+o(1)}). \tag{9}$$

To prove (9), we expand out the left-hand side as

$$\sum_{i,j \in S} \mathbb{E}(\mathrm{cov}(\lambda_i - \lambda_{i,n-1}, \lambda_j - \lambda_{j,n-1}|\lambda)|E)$$

and use tools from the theory of determinantal processes and complex analysis which we will not describe here.

Proving $\text{var}(\mathbb{E}(X_S|\lambda)|E) = O(n^{2-c+o(1)})$

Since the event E is restricting λ to a convex set in $\mathbb{R}^n$, so the probability distribution function of λ is still log-concave after conditioning to E. Thus Poincaré estimates such as Proposition 1 become available. As it turns out, a direct application of this proposition gives unfavorable estimates, basically because of long-range correlations between λ_i and λ_j make the operator norm of the inertia matrix large, and also because the known correlation decay estimates are currently only available in the bulk. To resolve this, we do not use the standard basis $e_1, \ldots, e_n$ of $\mathbb{R}^n$, but instead the following basis consisting of three groups:

- The vector $e_1 + \cdots + e_n$.
- The vectors $e_{i+1} - e_i$ for i in the bulk region $\texttt{bulk} := \{i : \epsilon n \leq i < (1-\epsilon)n\}$.
- The vectors $e_{i+1} - e_i$ for i in the edge region $\texttt{edge} := \{i : 1 \leq i < \epsilon n \text{ or } (1-\epsilon)n \leq i < n\}$.

Proving $\text{var}(\mathbb{E}(X_S|\lambda)|E) = O(n^{2-c+o(1)})$

The point is that $\mathbb{E}(X_S|\lambda)$ has different behavior in each of the three groups of vectors. In the direction $e_1 + \cdots + e_n$, the function $\mathbb{E}(X_S|\lambda)$ is in fact constant. This is because once one conditions on λ, the random variable $\lambda_{j,n-1}$ has the distribution of the j^{th} largest eigenvalue of the top left $n-1 \times n-1$ minor of a Hermitian matrix chosen uniformly at random amongst all matrices with eigenvalue λ. Moving λ in the direction $e_1 + \cdots + e_n$ then amounts to shifting λ_j and $\lambda_{j,n-1}$ by the same constant, so the expectation $\mathbb{E}(X_S|\lambda)$ remains unchanged.

As it turns out, $\mathbb{E}(X_S|\lambda)$ is significantly more sensitive to the bulk eigenvalue gaps $\lambda_{i+1} - \lambda_i$ than the edge eigenvalue gaps $\lambda_{j+1} - \lambda_j$. To exploit this, we apply the weighted log-concave Poincaré inequality with suitable choices of weights (sending the weight on the basis vector $e_1 + \cdots + e_n$ to infinity) to conclude that

$$\mathrm{var}(\mathbb{E}(X_S|\lambda)|E) \ll \mathbb{E}\left(|\nabla_{\mathtt{bulk}}\mathbb{E}(X_S|\lambda)|^2 + n|\nabla_{\mathtt{edge}}\mathbb{E}(X_S|\lambda)|^2|E\right) \times \left(\|M_{\mathtt{bulk}}\|_{\mathrm{op}} + n^{-1}\|M_{\mathtt{edge}}\|_{\mathrm{op}}\right)\log n \tag{10}$$

where for $\Omega = \mathtt{bulk}, \mathtt{edge}$ one has

$$|\nabla_\Omega \mathbb{E}(X_S|\lambda)|^2 := \sum_{i\in\Omega} |(\partial_{\lambda_{i+1}} - \partial_{\lambda_i})\mathbb{E}(X_S|\lambda)|^2$$

and M_Ω is the covariance matrix with entries

$$\mathrm{cov}(\lambda_{i+1} - \lambda_i, \lambda_{j+1} - \lambda_j|E)$$

for $i, j \in \Omega$.

Proving $\mathrm{var}(\mathbb{E}(X_S|\lambda)|E) = O(n^{2-c+o(1)})$

To prove $\mathrm{var}(\mathbb{E}(X_S|\lambda)|E) = O(n^{2-c+o(1)})$, it suffices (modulo controlling $\|M_\Omega\|_{\mathrm{op}}$ suitably, which follows from Cipolloni-Erdős-Schröder), to establish the bound

$$\mathbb{E}\left(|\nabla_{\mathtt{bulk}}\mathbb{E}(X_S|\lambda)|^2 + n|\nabla_{\mathtt{edge}}\mathbb{E}(X_S|\lambda)|^2|E\right) \ll n^{1+o(1)}. \quad (11)$$

We establish the bound

$$\mathbb{E}(|\partial_{\lambda_i}\mathbb{E}(\lambda_j - \lambda_{j,n-1}|\lambda)|^2|E) \ll n^{o(1)}(1 + n|\gamma_i - \gamma_j|)^{-4} \quad (12)$$

using determinantal kernels whenever $1 \leq i \leq n$ and $j \in \mathtt{bulk}$.

Proving $\text{var}(\mathbb{E}(X_S|\lambda)|E) = O(n^{2-c+o(1)})$

Taking square roots and summing over $j \in S$ using the triangle inequality, one obtains

$$\mathbb{E}(|\partial_{\lambda_i}\mathbb{E}(X_S|\lambda)|^2|E) \ll n^{-2+o(1)}$$

for $i \in \texttt{edge}$, and

$$\mathbb{E}(|\partial_{\lambda_i}\mathbb{E}(X_S|\lambda)|^2|E) \ll n^{o(1)}$$

for $i \in \texttt{bulk}$. Summing in i, one obtains

$$\mathbb{E}\left(|\nabla_{\texttt{bulk}}\mathbb{E}(X_S|\lambda)|^2 + n|\nabla_{\texttt{edge}}\mathbb{E}(X_S|\lambda)|^2|E\right) \ll n^{1+o(1)}.$$

and thus that

$$\text{var}(\mathbb{E}(X_S|\lambda)|E) = O(n^{2-c+o(1)}).$$

□

Open questions

1. What can be said about the concentration of random real valued augmented hives with general boundary conditions? If they do concentrate, what are the possible subsequential limit shapes? In particular, is the limit unique? In the limit when one of the boundary conditions is more spread out than the other, the limit shape should essentially degenerate to fractional free convolution powers (See Shlyakhtenko and Tao).
2. Do the local statistics of the random augmented GUE hive process converge (either in the bulk or the edge) to a known limit? In the case of the random Gelfand–Tsetlin process, the limit is known to essentially be the Boutillier bead process.
3. Do random integer valued augmented hives with general boundary conditions concentrate? Again, if they do concentrate, what are the possible subsequential limit shapes? In particular, is the limit unique?

Thank you for your attention!