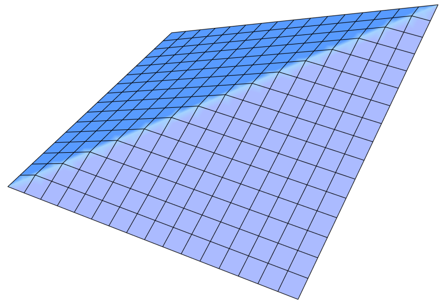

MAT 1128H: Topics in Probability: Random Planar Geometry and Random Matrices

Prof. Bálint Virág • Winter 2021 • University of Toronto
W 10:00AM-12:00PM, F 10:00-11:00AM (Online)

Transcribed, LaTeXed, and edited by Rob Zimmerman

Note: These are personal lecture notes, not official course materials. They may contain errors (which by default you should attribute to me, Rob Zimmerman, rather than the course instructor). The permanent home of these — and all of my other lecture notes — is `https://rob-zimmerman.github.io/coursenotes`.

Contents

Editorial Note Most Friday lectures were presented by students in the course on topics arranged in advance with the professor.

Lecture 1 — Wednesday, January 13, 2021

1. Random Growth and First-Passage Percolation

Random Plane Geometry and Random Matrices

Random planar geometry is a comparatively young subject, developed substantially over the past three decades and with many contributions from Toronto researchers. Its questions are geometric in the classical sense: they concern random shapes, distances, and paths. The theory draws on classical probability, particularly Brownian motion, but many of its central results are too recent to have reached standard textbooks.

As a first model, consider bond percolation on the square lattice. Each edge is retained independently with probability $p > 1/2$, perhaps with $p = 0.99$. The graph distance between two vertices is then random; it is finite precisely when the vertices lie in the same connected component. This construction is illustrated in Figure 1. This is a random geometry because the percolation config-

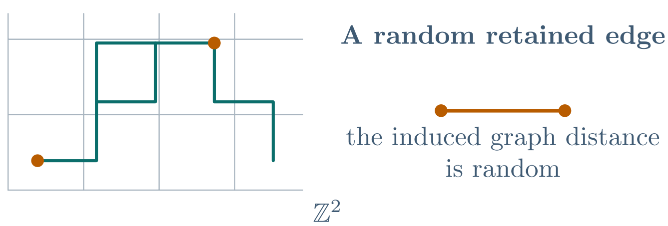

Figure 1

uration induces a random notion of distance. For a fixed origin, the ball of radius r, consisting of the vertices at graph distance at most r, is therefore a random set. After rescaling, large balls often approach a deterministic **limit shape** while retaining random boundary fluctuations, as suggested by Figure 2. Under broad hypotheses, a deterministic limit shape exists, although its exact form is generally unknown.

We may also begin with a random collection of points in a region $A \subseteq \mathbb{R}^2$, connect nearby points, and use shortest path length as the distance. In this model, shown in Figure 3, every pair of points can be arranged to have finite distance. A familiar nonrandom analogue is the travel time supplied by a navigation service. As Figure 4 emphasizes, traffic, weather, and construction make this travel time random and time dependent. Moreover, the travel time from a to b need not equal

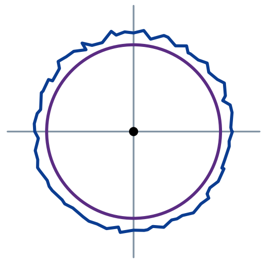

FIGURE 2

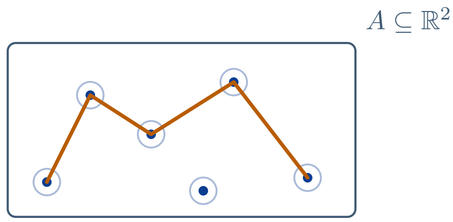

FIGURE 3

the travel time from b to a. The triangle inequality survives, but symmetry does not.

FIGURE 4

Definition 1.1 Let S be a set. A **directed metric** on S is a function

$$d\colon S \times S \longrightarrow \mathbb{R} \cup \{\infty\}$$

such that the following conditions hold:

1. $d(x, x) = 0$ for every $x \in S$.
2. $d(x, z) \leq d(x, y) + d(y, z)$ for every $x, y, z \in S$.

No symmetry condition is imposed. Thus $d(x, y)$ and $d(y, x)$ may be different.

Example 1.2 The shortest path distance on a directed graph is a directed metric. In the graph in Figure 5, the shortest directed path from x to y uses two edges, so $d(x, y) = 2$.

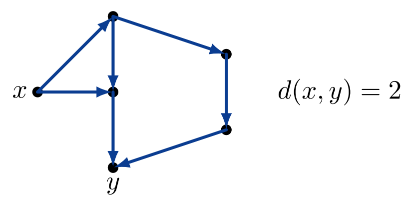

FIGURE 5

One of the principal goals is to understand **universality**: very different random models can exhibit the same limiting fluctuations. The central limit theorem provides the classical example. For independent and identically distributed summands, one expects

$$S_n = cn + c'\sqrt{n}\,Z + o(\sqrt{n}) \qquad \text{and} \qquad Z \sim \mathcal{N}(0, 1).$$

In many models of random geometry, a passage time instead has the characteristic form

$$d\big((0,0),(n,n)\big) \approx cn + c'n^{1/3}\,\mathcal{T}_2,$$

where c and c' are constants and $\mathcal{T}_2$ has the parameter-2 Tracy–Widom distribution. The exponents and limiting law differ from those in the central limit theorem, but the underlying universality principle is analogous.

Lecture 2 — Wednesday, January 20, 2021

The definition of a directed metric includes several familiar examples:

1. The shortest path distance on a directed graph is a directed metric.
2. Every ordinary metric is a directed metric.
3. The shortest path distance on a directed graph with nonnegative edge or vertex weights is a directed metric.

For the third example, assign a nonnegative weight $w(e)$ to each directed edge e. For a directed path $\pi = (v_0, v_1, \ldots, v_k)$, write $e_i = (v_{i-1}, v_i)$ for its ith edge. The length of π is

$$\text{len}(\pi) = \sum_{i=1}^{k} w(e_i).$$

The path shown in Figure 6 has length $w(e_1) + w(e_2) + w(e_3)$. The Seppäläinen–Johansson model

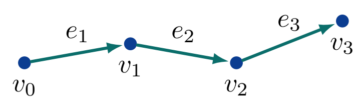

FIGURE 6

is an example of an **exactly solvable model**, meaning that many quantities admit explicit formulas. On the directed lattice in Figure 7, upward edges have a fixed nonnegative weight, whereas rightward edges have independent Bernoulli weights: The rightward edge weights are

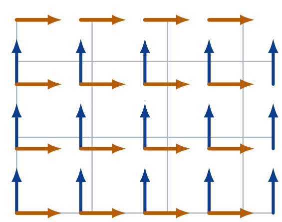

FIGURE 7

mutually independent, and

$$w(\uparrow) = c \geq 0 \quad \text{deterministically} \qquad \text{and} \qquad w(\rightarrow) \sim \text{Bernoulli}(p).$$

Another important model is **vertex-weighted last-passage percolation**. Assign independent and identically distributed nonnegative weights $w(v)$ to the vertices of a directed graph G. For a directed path π, set $w(\pi) = \sum_{v \in \pi} w(v)$, and define the directed distance with negative sign

$$d(e, f) = \sup_{\substack{\text{directed paths } \pi \\ \pi\colon e \to f}} w(\pi).$$

Exactly solvable versions often take the vertex weights to be independent geometric random variables. The maximizing path in Figure 8 has weight 7. If the underlying graph is not directed, arbitrarily long cycles may make the supremum infinite.

Several elementary operations preserve directed metrics. If d_1 and d_2 are directed metrics on the same set, then the following are also directed metrics:

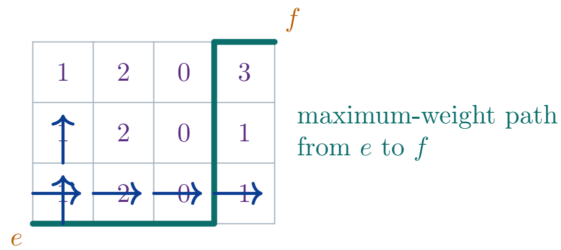

FIGURE 8

1. The sum $d_1 + d_2$ is a directed metric.
2. The scalar multiple ad_1 is a directed metric for every $a \geq 0$.
3. The pointwise maximum $d_1 \vee d_2$ is a directed metric.

For example, on $\mathbb{R}^2$, define

$$d_1(\boldsymbol{x}, \boldsymbol{y}) = |x_1 - y_1| \qquad \text{and} \qquad d_2(\boldsymbol{x}, \boldsymbol{y}) = |x_2 - y_2|.$$

Then $(d_1 + d_2)(\boldsymbol{x}, \boldsymbol{y}) = \|\boldsymbol{x} - \boldsymbol{y}\|_1$. To verify the third construction, set $d = d_1 \vee d_2$. For $i \in \{1, 2\}$,

$$d_i(x, y) \leq d_i(x, z) + d_i(z, y) \leq d(x, z) + d(z, y).$$

Taking the maximum over i proves the triangle inequality for d. The same argument applies to the pointwise supremum of any family of directed metrics, provided that the supremum is well defined.

Every function $g \colon S \to \mathbb{R}$ also determines a directed metric by

$$d_g(x, y) = g(y) - g(x).$$

The triangle inequality is an equality for d_g, so this is called a **trivial directed metric**.

Definition 2.1 Let $f \colon S^2 \to \mathbb{R} \cup \{\infty\}$, and suppose that at least one directed metric d satisfies $d \leq f$. The **directed metric induced by $\boldsymbol{f}$** is the pointwise supremum of all such directed metrics.

This supremum is again a directed metric: its diagonal values are zero, and

$$\sup_d d(x, z) \leq \sup_d \big(d(x, y) + d(y, z)\big) \leq \sup_d d(x, y) + \sup_d d(y, z).$$

Example 2.2 Let $G = (V, E)$, and define $f: V^2 \to \mathbb{R} \cup \{\infty\}$ by

$$f(x, y) = \begin{cases} 1, & xy \in E, \\ +\infty, & xy \notin E. \end{cases}$$

Exercise 2.3 Prove that the directed metric induced by the function f in Example 2.2 is the graph distance.

Solution. Let ρ denote the graph distance. This is a directed metric, and $\rho \le f$: if x and y are adjacent, then $\rho(x, y) = 1 = f(x, y)$, while otherwise $f(x, y) = +\infty$. Thus ρ is one of the directed metrics appearing in the supremum that defines the directed metric induced by f.

Conversely, let d be any directed metric satisfying $d \le f$, and let $x = v_0, v_1, \ldots, v_k = y$ be a path in G. Repeated use of the triangle inequality gives

$$d(x, y) \le \sum_{i=1}^{k} d(v_{i-1}, v_i) \le \sum_{i=1}^{k} f(v_{i-1}, v_i) = k.$$

Taking the infimum over all paths from x to y shows that $d(x, y) \le \rho(x, y)$. Hence every directed metric dominated by f is also dominated by ρ. It follows that their pointwise supremum is exactly ρ, the graph distance. □

Directed Norms and Scaling Limits

Definition 2.4 Let S be a real vector space. A **directed norm with positive sign** is a function

$$\|\cdot\|: S \longrightarrow \mathbb{R} \cup \{\infty\}$$

such that the following conditions hold:

1. The triangle inequality $\|\boldsymbol{x} + \boldsymbol{y}\| \le \|\boldsymbol{x}\| + \|\boldsymbol{y}\|$ holds for every $\boldsymbol{x}, \boldsymbol{y} \in S$.
2. One has $\|\mathbf{0}\| = 0$, and positive homogeneity, $\|a\boldsymbol{x}\| = a\|\boldsymbol{x}\|$, holds for every $\boldsymbol{x} \in S$ and every $a > 0$.

Reversing the inequality in (1) and replacing the codomain by $\mathbb{R} \cup \{-\infty\}$ gives a **directed norm with negative sign**; the zero and positive-homogeneity conditions are unchanged. Restricting homogeneity to $a > 0$ avoids the undefined products $0 \cdot (+\infty)$ and $0 \cdot (-\infty)$. A directed norm produces a translation-invariant directed metric of the same sign by setting $d(\boldsymbol{x}, \boldsymbol{y}) = \|\boldsymbol{y} - \boldsymbol{x}\|$.

Example 2.5 Fix $p \in (0, 1]$. For $\boldsymbol{x} = (x_1, \ldots, x_d) \in \mathbb{R}^d$, define

$$\|\boldsymbol{x}\|_p = \begin{cases} \left(\displaystyle\sum_{i=1}^{d} x_i^p\right)^{1/p}, & x_i \geq 0 \text{ for every } i, \\ -\infty, & \text{otherwise.} \end{cases}$$

This is a directed norm with negative sign: its homogeneity is immediate, and its reverse triangle inequality is

$$\|\boldsymbol{x} + \boldsymbol{y}\|_p \geq \|\boldsymbol{x}\|_p + \|\boldsymbol{y}\|_p.$$

Consequently, $d(\boldsymbol{x}, \boldsymbol{y}) = \|\boldsymbol{y} - \boldsymbol{x}\|_p$ is a directed metric with negative sign.

Consider $p = 1/2$ and $d = 2$. On the nonnegative quadrant,

$$\|(x, y)\|_{1/2} = (\sqrt{x} + \sqrt{y})^2.$$

The unit circle satisfies

$$\sqrt{x} + \sqrt{y} = 1 \iff y = (1 - \sqrt{x})^2 = 1 + x - 2\sqrt{x},$$

and is shown in Figure 9.

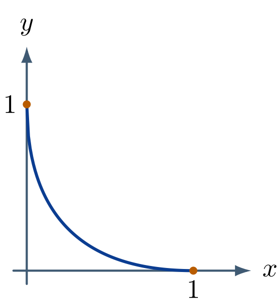

FIGURE 9

For a norm with negative sign, the ball of radius 1 centred at the origin is the superlevel set

$$B_1(0) = \{(x,y) : \|(x,y)\|_{1/2} \geq 1\}.$$

As Figure 10 shows, this ball does not contain its center.

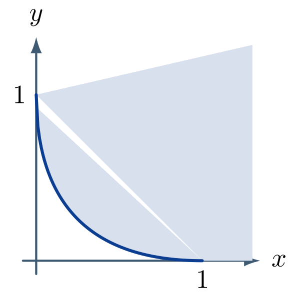

FIGURE 10

Exercise 2.6 Prove that balls for directed norms are convex.

Solution. First suppose that N is a directed norm with positive sign, and define its ball of radius r by

$$B_r = \{\boldsymbol{x} : N(\boldsymbol{x}) \leq r\}.$$

If $\boldsymbol{x}, \boldsymbol{y} \in B_r$ and $\theta \in [0,1]$, then positive homogeneity and the triangle inequality give

$$N\big(\theta\boldsymbol{x} + (1-\theta)\boldsymbol{y}\big) \leq \theta N(\boldsymbol{x}) + (1-\theta)N(\boldsymbol{y}) \leq r.$$

Thus B_r is convex.

For a directed norm with negative sign, the ball is the superlevel set

$$B_r = \{\boldsymbol{x} : N(\boldsymbol{x}) \geq r\}.$$

The reverse triangle inequality yields

$$N\big(\theta\boldsymbol{x} + (1-\theta)\boldsymbol{y}\big) \geq \theta N(\boldsymbol{x}) + (1-\theta)N(\boldsymbol{y}) \geq r,$$

so this ball is also convex. Translating either ball preserves convexity, which proves the result for balls centred at arbitrary points. □

Remark 2.7 Given a closed convex set $A \subseteq \mathbb{R}^d$ containing the origin, we may ask whether A is the unit ball of a directed norm. For the linear subspace

$$A = \{(x, 0) : x \in \mathbb{R}\},$$

the function

$$\|(x,y)\| = \begin{cases} 0, & y = 0, \\ \infty, & y \neq 0 \end{cases}$$

provides such a norm. The analogous construction that assigns 0 on an arbitrary convex set A and ∞ outside A is not generally homogeneous. It works only when $aA = A$ for every $a > 0$.

The central example is the **stretched Euclidean norm** on $\mathbb{R}^2$, defined by

$$\|(x,y)\|_{\uparrow} \coloneqq \begin{cases} \dfrac{x^2}{y}, & y > 0, \\ 0, & x = y = 0, \\ +\infty, & \text{otherwise.} \end{cases}$$

Its unit ball is the epigraph of the parabola $y = x^2$, as shown in Figure 11. The name arises

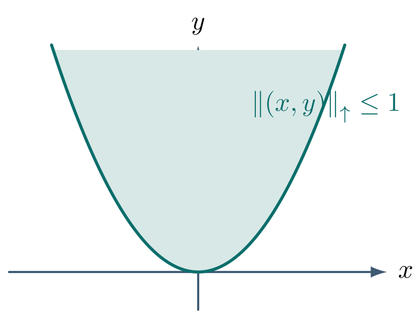

Figure 11

by stretching Euclidean space in the vertical direction. For a large parameter a, the rectangle in Figure 12 represents a region with horizontal scale 1 and vertical scale a. For $y > 0$, define the **centred Euclidean norm**

$$\|(x,y)\|_a \coloneqq \|(x, ay)\|_2 - ay.$$

Since $\|(x, ay)\|_2 \sim ay$, a Taylor expansion of the square root about 1 gives

$$\|(x,y)\|_a = \sqrt{x^2 + a^2y^2} - ay$$

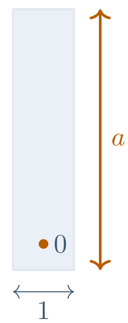

FIGURE 12

$$
\begin{aligned}
&= ay\left(\sqrt{1+\frac{x^2}{a^2y^2}}-1\right)\\
&\sim ay\left(\frac{x^2}{2a^2y^2}\right)=\frac{x^2}{2ay}.
\end{aligned}
$$

Therefore,

$$2a\|(x,y)\|_a \xrightarrow[a\to\infty]{} \|(x,y)\|_{\uparrow}.$$

This convergence can be expressed entirely in terms of directed metrics. For $\boldsymbol{p}=(p_1,p_2)$, let $\boldsymbol{p}_a=(p_1,ap_2)$, and define

$$d_a(\boldsymbol{p},\boldsymbol{q}) := 2a\left(d_{\text{Euclidean}}(\boldsymbol{p}_a,\boldsymbol{q}_a)+a(p_2-q_2)\right).$$

The Euclidean term is a metric, the term $a(p_2-q_2)$ is a trivial directed metric, and nonnegative scalar multiples and sums preserve directed metrics. Hence every d_a is a directed metric. If $q_2>p_2$, then

$$d_a(\boldsymbol{p},\boldsymbol{q}) \xrightarrow[a\to\infty]{} \|\boldsymbol{q}-\boldsymbol{p}\|_{\uparrow}.$$

In this sense, the stretched Euclidean geometry is a scaling limit of Euclidean metric geometry. The transition from isotropic Euclidean geometry to a geometry in which direction matters is summarized in Figure 13.

Remark 2.8 What changes if the Euclidean L^2-norm is replaced by an L^p-norm?

The stretched Euclidean norm is universal in the following sense: after suitable stretching, centring, and scaling, any sufficiently regular norm has the same directional quadratic limit.

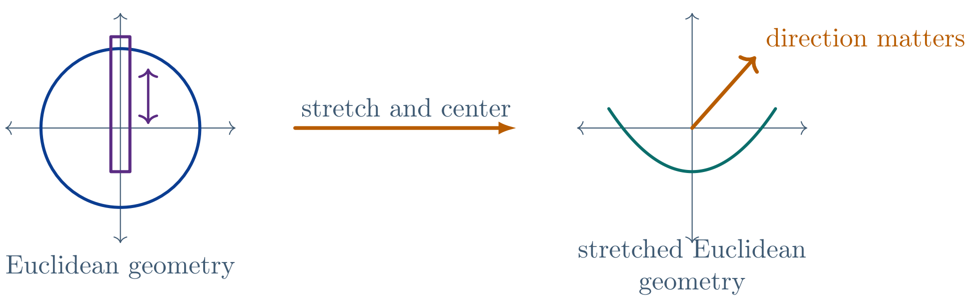

Figure 13

Exercise 2.9 Let $p > 1$, and let $\boldsymbol{u}, \boldsymbol{v}$ be perpendicular unit vectors in $\mathbb{R}^2$ such that $\boldsymbol{u} \neq (0, \pm 1)$ and $\boldsymbol{v} \neq (\pm 1, 0)$. Determine the asymptotic behaviour of $\|x\boldsymbol{v} + ay\boldsymbol{u}\|_p$ as $a \to \infty$.

Solution. Set $U = \|\boldsymbol{u}\|_p$ and define

$$F(s) \coloneqq \|\boldsymbol{u} + s\boldsymbol{v}\|_p.$$

The assumptions on $\boldsymbol{u}$ ensure that both of its coordinates are nonzero, so F is twice continuously differentiable near 0. For fixed x and $y > 0$,

$$\|x\boldsymbol{v} + ay\boldsymbol{u}\|_p = ayF\left(\frac{x}{ay}\right) = ayU + xF'(0) + \frac{x^2}{2ay}F''(0) + o(a^{-1}).$$

For $y < 0$, the same calculation and the symmetry of the norm give

$$\|x\boldsymbol{v} + ay\boldsymbol{u}\|_p = a|y|U + \operatorname{sgn}(y)xF'(0) + \frac{x^2}{2a|y|}F''(0) + o(a^{-1}).$$

When $y = 0$, the expression is simply $|x|\|\boldsymbol{v}\|_p$. In particular, whenever $y \neq 0$, its leading-order behaviour is

$$\|x\boldsymbol{v} + ay\boldsymbol{u}\|_p \sim a|y|\|\boldsymbol{u}\|_p.$$

□

After normalizing $\|\boldsymbol{u}\|_p = 1$ and restricting to $y > 0$, the expansion contains both a term linear in x and a quadratic correction. Thus there are constants $c > 0$ and $c' \in \mathbb{R}$, depending on $\boldsymbol{u}$, $\boldsymbol{v}$, and p, such that

$$a\left(\|x\boldsymbol{v} + ay\boldsymbol{u}\|_p + c'x - ay\right) \longrightarrow c\|(x, y)\|_{\uparrow}. \tag{2.1}$$

Exercise 2.10 Prove the convergence in (2.1) and determine c and c'. The constants may depend on $\boldsymbol{u}$, $\boldsymbol{v}$, and p, but not on x or y.

Solution. Write $\boldsymbol{u} = (u_1, u_2)$ and $\boldsymbol{v} = (v_1, v_2)$, and set

$$\alpha := \sum_{i=1}^{2} |u_i|^{p-2} u_i v_i \qquad \text{and} \qquad \beta := \sum_{i=1}^{2} |u_i|^{p-2} v_i^2.$$

Set $U = \|\boldsymbol{u}\|_p$ and $F(s) = \|\boldsymbol{u} + s\boldsymbol{v}\|_p$. Direct differentiation gives

$$F'(0) = U^{1-p}\alpha \qquad \text{and} \qquad F''(0) = (p-1)\left(U^{1-p}\beta - U^{1-2p}\alpha^2\right).$$

Under the normalization $U = 1$, the expansion for $y > 0$ becomes

$$\|x\boldsymbol{v} + ay\boldsymbol{u}\|_p = ay + \alpha x + \frac{p-1}{2a}\left(\beta - \alpha^2\right)\frac{x^2}{y} + o(a^{-1}).$$

Consequently,

$$c' = -\alpha \qquad \text{and} \qquad c = \frac{p-1}{2}\left(\beta - \alpha^2\right).$$

The weighted Cauchy–Schwarz inequality shows that $\beta - \alpha^2 \geq 0$. Equality would force $\boldsymbol{v}$ to be parallel to $\boldsymbol{u}$, contrary to perpendicularity, so $c > 0$. Since $\|(x, y)\|_\uparrow = x^2/y$ for $y > 0$, this proves (2.1). $\square$

The rescaling has converted Euclidean geometry on a long, thin region into the stretched Euclidean geometry. The geometric change is illustrated schematically in Figure 14. For points $\boldsymbol{p} = (p_1, p_2)$

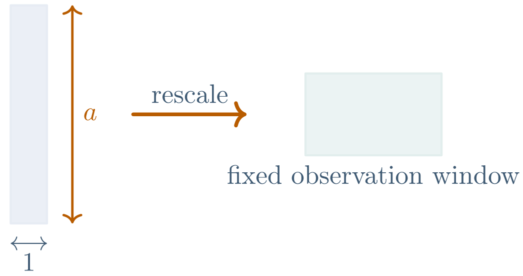

Figure 14

and $\boldsymbol{q} = (q_1, q_2)$ with $q_2 > p_2$, the directed length of the line segment from $\boldsymbol{p}$ to $\boldsymbol{q}$ is $\|\boldsymbol{q} - \boldsymbol{p}\|_\uparrow$, as shown in Figure 15. For a piecewise linear path through $\boldsymbol{p}_0, \boldsymbol{p}_1, \ldots, \boldsymbol{p}_k$, the length is the sum of

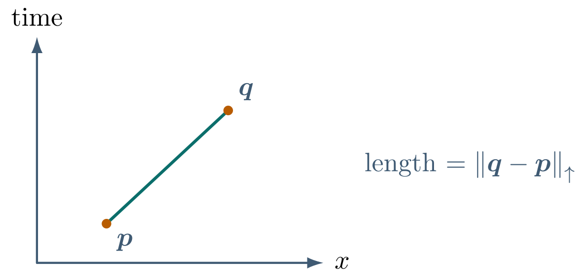

FIGURE 15

the directed lengths of its segments:

$$\sum_{i=1}^{k} \|\boldsymbol{p}_i - \boldsymbol{p}_{i-1}\|_{\uparrow}.$$

This additivity is illustrated in Figure 16. Now let $\gamma\colon [0,1] \to \mathbb{R}^2$ be continuous. For a partition

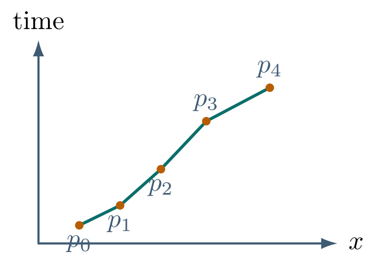

FIGURE 16

$$t_0 = 0 < t_1 < \cdots < t_k = 1,$$

define

$$|\gamma|_{\Pi} = \sum_{i=1}^{k} \|\gamma(t_i) - \gamma(t_{i-1})\|_{\uparrow}.$$

The length of γ is

$$|\gamma| = \sup_{\Pi} |\gamma|_{\Pi}.$$

Because the stretched Euclidean norm is infinite in backward time directions, any path that backtracks in its time coordinate has infinite length.

Exercise 2.11 Suppose that $\gamma(t) = (f(t), t)$. Express $|\gamma|$ in terms of f, and identify the resulting familiar functional.

Solution. For a partition $\Pi = \{0 = t_0 < t_1 < \cdots < t_k = 1\}$, the definition gives

$$|\gamma|_\Pi = \sum_{i=1}^{k} \frac{\big(f(t_i) - f(t_{i-1})\big)^2}{t_i - t_{i-1}}.$$

Therefore,

$$|\gamma| = \sup_{\Pi} \sum_{i=1}^{k} \frac{\big(f(t_i) - f(t_{i-1})\big)^2}{t_i - t_{i-1}}.$$

If f is absolutely continuous and $f' \in L^2([0,1])$, the Cauchy–Schwarz inequality on each interval of the partition gives

$$\frac{\big(f(t_i) - f(t_{i-1})\big)^2}{t_i - t_{i-1}} \le \int_{t_{i-1}}^{t_i} \big(f'(t)\big)^2 \,\mathrm{d}t.$$

Approximating f' by its averages on increasingly fine partitions gives the reverse inequality in the limit. Hence

$$|\gamma| = \int_0^1 \big(f'(t)\big)^2 \,\mathrm{d}t,$$

the **Dirichlet energy** of f. For a continuous function outside the Sobolev space $H^1([0,1])$, the displayed supremum is infinite. $\square$

Lecture 3 — Friday, January 22, 2021

Theorem 3.1 (Birkhoff ergodic theorem) Let $(\Omega, \mathcal{A}, \mu)$ be a probability space, let $T\colon \Omega \to \Omega$ be measure preserving, and let $f \in L^1(\mu)$. Thus

$$\mu(T^{-1}(A)) = \mu(A)$$

for every $A \in \mathcal{A}$. Define

$$f_n(x) = \sum_{i=0}^{n-1} f(T^i x).$$

Then the limit

$$\bar{f}(x) := \lim_{n\to\infty} \frac{f_n(x)}{n}$$

exists almost surely. Moreover, $\bar{f} \in L^1(\mu)$, the function $\bar{f}$ is T-invariant, and

$$\bar{f} = \mathbb{E}\left[f \mid \mathcal{B}\right],$$

where $\mathcal{B}$ is the σ-algebra of T-invariant sets. If T is ergodic, then $\bar{f} = \mathbb{E}[f]$ almost surely.

In particular, Theorem 3.1 implies the strong law of large numbers by taking T to be the left shift on a sequence of independent and identically distributed random variables.

Theorem 3.2 (Subadditive ergodic theorem) Let $(\Omega, \mathcal{A}, \mu)$ be a probability space, let $T\colon \Omega \to \Omega$ be measure preserving, and suppose that $f_n \in L^1(\mu)$ for every $n \geq 1$. Assume that

$$f_{m+n} \leq f_m + f_n \circ T^m$$

for all $m, n \geq 1$. Then

$$\bar{f}(x) := \lim_{n\to\infty} \frac{f_n(x)}{n}$$

exists almost surely as an extended real random variable and is T-invariant. If T is ergodic, then $\bar{f}$ is almost surely constant.

First-Passage Percolation

The **nearest-neighbour graph** on $\mathbb{Z}^d$ has vertex set $\mathbb{Z}^d$. Two vertices $\boldsymbol{x}$ and $\boldsymbol{y}$ are joined by an undirected edge precisely when

$$\|\boldsymbol{x} - \boldsymbol{y}\|_1 = 1.$$

Equivalently, the neighbours of $\boldsymbol{x}$ are $\boldsymbol{x} \pm \boldsymbol{e}_i$ for $i \in \{1, \dots, d\}$, where $\boldsymbol{e}_1, \dots, \boldsymbol{e}_d$ are the standard basis vectors. Thus every vertex has degree $2d$. The two-dimensional case is shown in Figure 17; the highlighted edges are exactly the four edges incident to the origin.

Proposition 3.3 Consider the nearest-neighbour graph on $\mathbb{Z}^d$. Assign independent and identically distributed nonnegative weights w_e to its edges, and assume $\mathbb{E}[w_e] < \infty$. For $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{Z}^d$, define the **first-passage time**

$$L(\boldsymbol{x}, \boldsymbol{y}) = \inf_{\pi\colon \boldsymbol{x}\to\boldsymbol{y}} \sum_{e\in\pi} w_e.$$

Fix a nonzero $\boldsymbol{v} \in \mathbb{Z}^d$. Then $L(\boldsymbol{0}, n\boldsymbol{v})/n$ converges almost surely and in L^1 to a finite random variable.

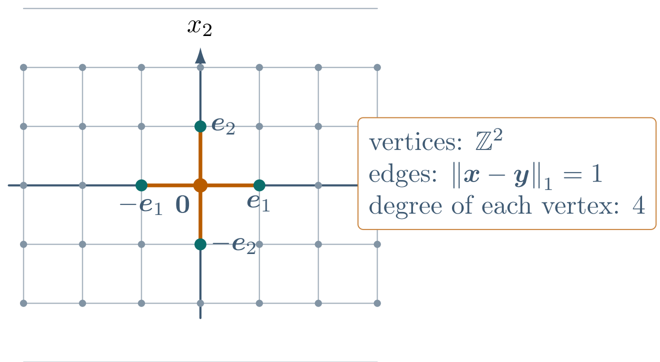

FIGURE 17

Proof. Let ν denote the common law of the edge weights, and equip

$$\Omega = \mathbb{R}_{\geq 0}^{E}$$

with the product measure $\mu = \prod_{e\in E} \nu$. An element $\omega = (t_e)_{e\in E}$ records one realization of all edge weights. Let $T\colon \Omega \to \Omega$ translate the environment by $\boldsymbol{v}$:

$$T\big((t_e)_{e\in E}\big) = (t_{e+\boldsymbol{v}})_{e\in E}.$$

Translation invariance of the product measure makes T measure preserving. Define

$$f_n\big((t_e)_{e\in E}\big) = \inf_{\pi\colon \mathbf{0}\to n\boldsymbol{v}} \sum_{e\in\pi} t_e.$$

The random variable f_n is integrable. Indeed, for any fixed path $\pi\colon \mathbf{0} \to n\boldsymbol{v}$,

$$\mathbb{E}[f_n] \leq \sum_{e\in\pi} \mathbb{E}[w_e] = |\pi|\,\mathbb{E}[w_e] < \infty.$$

For $m, n \geq 1$, translation of the environment gives

$$\begin{aligned} f_n \circ T^m\big((t_e)_{e\in E}\big) &= f_n\big((t_{e+m\boldsymbol{v}})_{e\in E}\big) \\ &= \inf_{\pi\colon \mathbf{0}\to n\boldsymbol{v}} \sum_{e\in\pi} t_{e+m\boldsymbol{v}} \\ &= \inf_{\pi\colon m\boldsymbol{v}\to(m+n)\boldsymbol{v}} \sum_{e\in\pi} t_e. \end{aligned}$$

Concatenating arbitrarily near-minimizing paths from $\mathbf{0}$ to $m\boldsymbol{v}$ and from $m\boldsymbol{v}$ to $(m+n)\boldsymbol{v}$ yields

$$f_{m+n} = \inf_{\pi\colon \mathbf{0}\to(m+n)\boldsymbol{v}} \sum_{e\in\pi} t_e$$

$$\begin{aligned}&\leq \inf_{\pi\colon \mathbf{0}\to m\boldsymbol{v}} \sum_{e\in\pi} t_e + \inf_{\pi\colon m\boldsymbol{v}\to(m+n)\boldsymbol{v}} \sum_{e\in\pi} t_e \\ &= f_m + f_n \circ T^m.\end{aligned}$$

The subadditive ergodic theorem therefore shows that

$$X \coloneqq \lim_{n\to\infty} \frac{L(\mathbf{0}, n\boldsymbol{v})}{n}$$

exists almost surely and is nonnegative. In fact, the convergence also holds in L^1. Repeated subadditivity gives

$$0 \leq \frac{f_n}{n} \leq \frac{1}{n}\sum_{i=0}^{n-1} f_1 \circ T^i.$$

The averages on the right are uniformly integrable: by the de la Vallée–Poussin criterion, choose an increasing convex function Φ with $\Phi(u)/u \to \infty$ and $\mathbb{E}[\Phi(f_1)] < \infty$, and then apply Jensen's inequality. Thus $(f_n/n)_{n\geq 1}$ is uniformly integrable, so its almost sure convergence implies convergence in L^1. Moreover, subadditivity of $(\mathbb{E}[f_n])_{n\geq 1}$ and Fekete's lemma give

$$\mathbb{E}[X] = \lim_{n\to\infty} \frac{\mathbb{E}[f_n]}{n} = \inf_{n\geq 1} \frac{\mathbb{E}[f_n]}{n}.$$

□

Proposition 3.4 The limiting random variable X is almost surely constant.

Proof. Enumerate the edges as $E = \{e_1, e_2, \dots\}$. A path may be made simple by deleting loops, without increasing its weight. Therefore, changing the weights of a finite set F of edges changes every passage-time infimum by at most the finite quantity $\sum_{e\in F} |w'_e - w_e|$, independently of n. After division by n, this change vanishes. Consequently, X is measurable with respect to the tail σ-algebra of $w_{e_1}, w_{e_2}, \dots$. The Kolmogorov zero–one law then implies that X is almost surely constant. □

Proof of Theorem 3.2. The proof reduces the general case to a nonpositive subadditive process. Define

$$f'_n(x) \coloneqq f_n(x) - \sum_{i=0}^{n-1} f_1(T^i x).$$

Subadditivity implies that $f'_n \leq 0$ and that the sequence (f'_n) remains subadditive. By Theorem 3.1, the normalized sum that was subtracted has an almost-sure limit. It therefore suffices to prove the theorem under the additional assumption that $f_n \le 0$ for every n.

Set
$$f(x) \coloneqq \liminf_{n\to\infty} \frac{f_n(x)}{n}.$$
Using subadditivity with $m = 1$, we have
$$\frac{f_{n+1}(x)}{n} \le \frac{f_1(x)}{n} + \frac{f_n(Tx)}{n}.$$
Taking lower limits gives $f(x) \le f(Tx)$. Since T is measure preserving, the level sets of these two functions agree up to null sets, and hence $f = f \circ T$ almost surely.

Fix $M > 1$, and let $F_M = \max\{-M, f\}$. This truncation is also T-invariant. We claim that, for every $\varepsilon \in (0,1)$,
$$\limsup_{n\to\infty} \frac{f_n(x)}{n} \le F_M(x) + \varepsilon \quad \text{almost surely.}$$
To prove the claim, fix $N \in \mathbb{N}$ and define
$$B_N = \left\{ x : \frac{f_\ell(x)}{\ell} > F_M(x) + \varepsilon \text{ for every } 1 \le \ell \le N \right\}.$$

Fix $x \in \Omega$, and partition the integer interval $[0, n)$ inductively. If $k \le n - N$ is the first unassigned integer and $T^k x \notin B_N$, choose $1 \le \ell \le N$ such that
$$f_\ell(T^k x) \le \ell\big(F_M(T^k x) + \varepsilon\big) = \ell\big(F_M(x) + \varepsilon\big),$$
and colour $[k, k+\ell)$ blue. If $T^k x \in B_N$, colour $[k, k+1)$ red. Finally, colour the remaining interval of length at most N green.

Let the blue intervals be $[t_j, t_j + \ell_j)$, let $b = \sum_j \ell_j$ be their total length, and let r and g be the total red and green lengths. Then
$$r \le \sum_{k=0}^{n-1} \mathbb{1}_{B_N}(T^k x), \qquad g \le N \qquad \text{and} \qquad b = n - r - g.$$

Since $f_1 \le 0$, repeated subadditivity yields
$$\begin{aligned} f_n(x) &\le \sum_j f_{\ell_j}(T^{t_j} x) + \sum_{\text{red } k} f_1(T^k x) + \sum_{\text{green } k} f_1(T^k x) \\ &\le \big(F_M(x) + \varepsilon\big) \sum_j \ell_j \end{aligned}$$

$$= \big(F_M(x) + \varepsilon\big)b.$$

If $F_M(x) + \varepsilon > 0$, the claimed bound follows immediately from $f_n \le 0$. On the remaining set, the coefficient is nonpositive. The Birkhoff ergodic theorem gives

$$\liminf_{n\to\infty} \frac{b}{n} \ge 1 - \mathbb{E}\left[\mathbb{1}_{B_N} \mid \mathcal{B}\right] \quad \text{almost surely.}$$

It follows that

$$\limsup_{n\to\infty} \frac{f_n(x)}{n} \le \big(F_M(x) + \varepsilon\big)\left(1 - \mathbb{E}\left[\mathbb{1}_{B_N} \mid \mathcal{B}\right]\right).$$

By the definition of f, the sets B_N decrease to the empty set as $N \to \infty$. Conditional dominated convergence therefore gives

$$\mathbb{E}\left[\mathbb{1}_{B_N} \mid \mathcal{B}\right] \longrightarrow 0.$$

Letting $N \to \infty$ proves the claim. Next let $M \to \infty$, so that $F_M \to f$, and then let $\varepsilon \downarrow 0$. We obtain

$$\limsup_{n\to\infty} \frac{f_n(x)}{n} \le f(x) = \liminf_{n\to\infty} \frac{f_n(x)}{n}.$$

Thus $f_n(x)/n$ converges almost surely. Its limit is T-invariant by the invariance already proved for f, and ergodicity makes every T-invariant extended real random variable almost surely constant. $\square$

Lecture 4 — Wednesday, January 27, 2021

The asymptotic shape in first-passage percolation depends on the law of the edge weights. Two familiar possibilities are the unit balls of the ℓ^1- and ℓ^2-norms, shown in Figure 18.

The contrast is geometrically significant: geodesics between two points are unique in the ℓ^2-norm, whereas they need not be unique in the ℓ^1-norm.

Remark 4.1 It is an open problem to determine whether some distribution of edge weights produces the ℓ^2-norm as its limiting norm.

Remark 4.2 More generally, we may ask when the limiting ball is strictly convex, or when it is strictly convex in a prescribed direction. No general criterion is known, although several

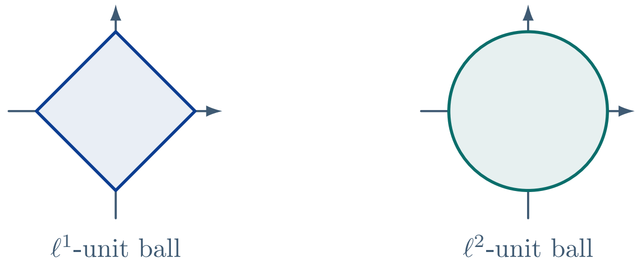

FIGURE 18

special cases can be analyzed.

Growth Models and Geometry

We begin with a deterministic graph model. Let $G = (V, E)$ be an undirected graph and let $S \subseteq V$ be a set of seed vertices. For an integer $r \geq 0$, define the **graph distance ball**

$$B_r(S) := \left\{ v \in V : \begin{array}{l} \text{there is a path } v_0, \ldots, v_k \text{ with } v_0 \in S, \\ v_k = v, \text{ and } k \leq r \end{array} \right\}.$$

The first update is illustrated in Figure 19: the original seed vertices form $B_0(S) = S$, and every neighbouring vertex is added to obtain $B_1(S)$. The sequence

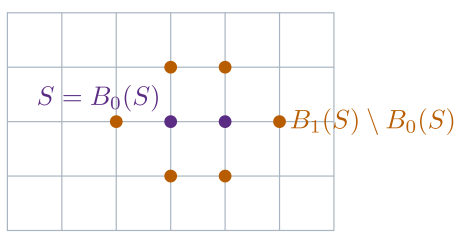

FIGURE 19

$$B_0(S) \subseteq B_1(S) \subseteq B_2(S) \subseteq \cdots$$

is a deterministic growth model. Its recursive rule is simple: to obtain $B_{r+1}(S)$, adjoin every neighbour of $B_r(S)$.

There is a dual construction on a directed acyclic graph. Define

$$B_r(S) := \left\{ v \in V : \begin{array}{l} \text{there is a directed path } v_0, \ldots, v_k \text{ with } v_0 \in S, \\ v_k = v, \text{ and } k \geq r \end{array} \right\}.$$

In the oriented square lattice in Figure 20, every edge points upward or to the right. Passing

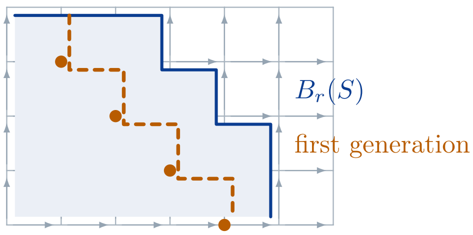

Figure 20

from $B_r(S)$ to $B_{r+1}(S)$ removes the first-generation vertices of $B_r(S)$, namely those vertices all of whose parents lie outside $B_r(S)$. Thus

$$B_0(S) \supseteq B_1(S) \supseteq B_2(S) \supseteq \cdots,$$

and this evolution is called a **shrink model**.

The same geometric idea produces a growth model on functions. Let $h\colon \mathbb{R} \to \mathbb{R} \cup \{-\infty\}$ be upper semicontinuous. For intuition, we may first take h to be continuous. Its **hypograph** is

$$\operatorname{hypo}(h) := \{(x, y) \in \mathbb{R}^2 : y \leq h(x)\}.$$

This region is shaded in Figure 21. Let

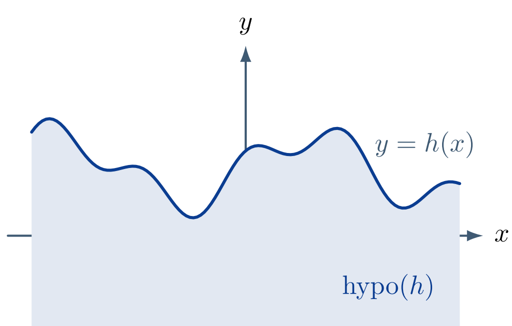

Figure 21

$$d(\boldsymbol{p}, \boldsymbol{q}) := \|\boldsymbol{q} - \boldsymbol{p}\|_{\uparrow}$$

be the directed metric induced by the stretched Euclidean norm. For a point $\boldsymbol{q}$, the backward ball $\{\boldsymbol{p} : d(\boldsymbol{p}, \boldsymbol{q}) \leq r\}$ is bounded above by a downward parabola, as shown in Figure 22. Define the

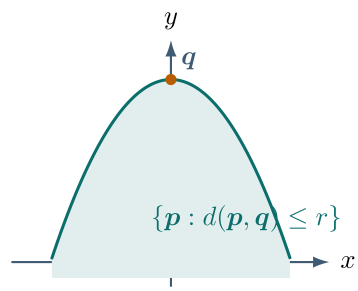

FIGURE 22

***r*-neighbourhood** of the graph of h by

$$B_r(h) \coloneqq \left\{ \boldsymbol{p} \in \mathbb{R}^2 : \inf_{\boldsymbol{q} \in \operatorname{graph}(h)} d(\boldsymbol{p}, \boldsymbol{q}) \leq r \right\}.$$

This set is again a hypograph, so it determines a function h_r through

$$B_r(h) = \operatorname{hypo}(h_r).$$

At time $r = 0$, one has $h_0 = h$, and the family $(h_r)_{r \geq 0}$ is a growth model of functions.

Exercise 4.3 Show, at least formally for a smooth initial function, that the evolution satisfies the **Hamilton–Jacobi equation**

$$\partial_r h_r = \frac{1}{4} (\partial_x h_r)^2.$$

Equivalently, after the time change $\widetilde{h}_s = h_{4s}$, show that $\partial_s \widetilde{h}_s = (\partial_x \widetilde{h}_s)^2$.

Solution. Write $\boldsymbol{p} = (x, z)$ and $\boldsymbol{q} = (y, h(y))$. The inequality $d(\boldsymbol{p}, \boldsymbol{q}) \leq r$ is equivalent to

$$z \leq h(y) - \frac{(x-y)^2}{r}.$$

Taking the upper boundary of the union of these backward balls gives the Hopf–Lax formula

$$h_r(x) = \sup_{y \in \mathbb{R}} \left\{ h(y) - \frac{(x-y)^2}{r} \right\}.$$

Assume formally that the supremum has a unique smooth maximizer $y_r(x)$. The envelope theorem then gives

$$\partial_r h_r(x) = \frac{(x - y_r(x))^2}{r^2} \qquad \text{and} \qquad \partial_x h_r(x) = -\frac{2(x - y_r(x))}{r}.$$

Consequently,

$$\partial_r h_r(x) = \frac{1}{4}\big(\partial_x h_r(x)\big)^2.$$

Finally, the chain rule for $\widetilde{h}_s = h_{4s}$ yields

$$\partial_s \widetilde{h}_s = 4\, \partial_r h_r|_{r=4s} = \big(\partial_x \widetilde{h}_s\big)^2,$$

which is the equivalent rescaled form. □

Hamilton–Jacobi equations originated in Hamilton's reformulation of classical mechanics. In that formulation, the evolution of geodesics is encoded by a partial differential equation rather than by Newton's equations of motion.

Random Growth Models

Fix $p \in (0, 1]$, and let W have the geometric distribution on $\{1, 2, 3, \dots\}$. Thus W is the number of Bernoulli trials required to observe the first success. If $q = 1 - p$, then

$$\mathbb{P}(W > x) = q^x, \qquad x \in \{0, 1, 2, \dots\}.$$

Consequently, for nonnegative integers x and y,

$$\mathbb{P}(W > x + y \mid W > x) = \frac{q^{x+y}}{q^x} = q^y = \mathbb{P}(W > y).$$

This is the **memoryless property** of the geometric distribution.

Let $G = (V, E)$ be an undirected graph, and assign independent vertex weights

$$(W_v)_{v \in V} \quad \text{with} \quad W_v \sim \text{Geometric}(p).$$

For a path $\pi = (v_0, \dots, v_k)$, set

$$W_\pi \coloneqq \sum_{j=0}^{k} W_{v_j}.$$

Given a seed set $S \subseteq V$, define the **passage time**

$$d(v) \coloneqq \min_{\substack{\pi=(v_0,\dots,v_k)\\ v_0\in S,\ v_k=v}} W_\pi$$

and the **random ball**

$$B_r(S) \coloneqq \{v \in V : d(v) \le r\}.$$

Because all vertex weights are positive integers, $B_0(S) = \varnothing$. Write ∂A for the **external vertex boundary** of A, and define the **active boundary**

$$Q_r(S) \coloneqq \big(\partial B_r(S) \cup S\big) \setminus B_r(S).$$

Proposition 4.4 Conditional on $B_0(S), \dots, B_r(S)$, the set $B_{r+1}(S)$ is obtained from $B_r(S)$ by adding each vertex of $Q_r(S)$, independently, with probability p.

Proof. The observed history has two consequences:

1. For each $v \in B_r(S)$, the history determines $d(v)$, namely the first time at which v entered the growing set.
2. For each $v \in Q_r(S)$, the history implies $d(v) \ge r+1$; otherwise, v would already belong to $B_r(S)$.

Only vertices in $Q_r(S)$ can be added at the next step. Fix such a vertex v. By (1), the quantity

$$m_v \coloneqq \min\left(\{d(u) : u \in B_r(S),\ u \sim v\} \cup \begin{cases}\{0\}, & v \in S,\\ \varnothing, & v \notin S\end{cases}\right)$$

is known from the history. The extra candidate 0 accounts for a path that starts at v when v is a seed. In either case,

$$d(v) = W_v + m_v.$$

Condition (2) says precisely that W_v has survived to the threshold $r - m_v$. By the memoryless property, the conditional law of $d(v) - r$ is again Geometric(p). Hence v enters $B_{r+1}(S)$ exactly when $d(v) - r = 1$, an event of conditional probability

$$\mathbb{P}(d(v) = r+1 \mid B_0(S), \dots, B_r(S)) = p.$$

The vertex weights are independent, so these events are conditionally independent over $v \in Q_r(S)$.
$\square$

This Markov chain is the **discrete-time Eden model**. It can represent, for example, the spread of an infection or the growth of a bacterial colony. The geometric waiting time also explains the continuous-time limit. If $W \sim \text{Geometric}(p)$, then

$$pW \xrightarrow[p\downarrow 0]{\mathrm{d}} \mathrm{Exp}(1).$$

Thus, when each discrete step is assigned duration p, the Eden chain converges formally to a continuous-time Markov chain.

A directed version produces a shrink model. Let G now be a directed acyclic graph, retain independent Geometric(p) vertex weights, and set

$$B_r(S) := \left\{ v \in V : \begin{array}{l} \text{there is a directed path } \pi \text{ from } S \text{ to } v \\ \text{such that } W_\pi \geq r \end{array} \right\}.$$

The same memoryless argument as in Proposition 4.4 shows that $B_{r+1}(S)$ is obtained from $B_r(S)$ by independently removing each first-generation vertex with probability p.

When the boundary of $B_r(S)$ is encoded as a lattice path and then rotated, each downward step becomes a particle and each upward step becomes a vacant site. Removing a corner of the boundary is therefore equivalent to moving a particle one site to the right, as illustrated in Figure 23. Under

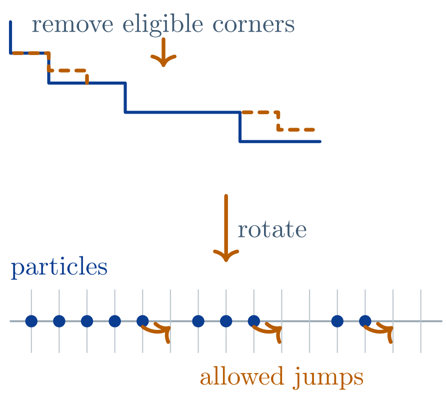

Figure 23

parallel discrete-time dynamics, every particle with a vacant site immediately to its right jumps into that site with probability p, independently of the other eligible particles. This is the **discrete-time totally asymmetric simple exclusion process**. The exclusion rule forbids two particles from occupying the same site.

The process was introduced in part to model one-dimensional transport. In a biological interpretation, particles represent ribosomes moving along messenger ribonucleic acid while synthesizing proteins. In a traffic interpretation, they represent vehicles that move forward whenever the next site is vacant. Both interpretations emphasize the nearest-neighbour exclusion rule.

There is also a passage time representation with Bernoulli edge weights. The two-row example in Figure 24 has horizontal edge weights in $\{0, 1\}$, vertical edges of weight 0, and minimum directed passage time $d(y) = 2$ from the indicated source to the target y. As the target y moves along the

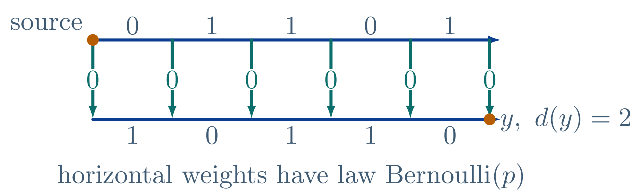

FIGURE 24

boundary, the passage time changes only when an additional unit of weight becomes unavoidable. A typical sample profile is therefore a nondecreasing step function, as depicted in Figure 25.

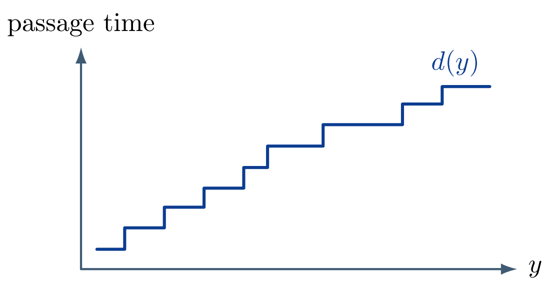

FIGURE 25

Lecture 5 — Friday, January 29, 2021

Directed First-Passage Percolation

Consider the nonnegative quadrant of the square lattice. Every directed path uses only rightward and upward steps. Horizontal edges have weight 0, while the vertical edge from (i,j) to $(i,j+1)$ has a nonnegative random weight $Y_{i,j}$. The geometry and the placement of the random weights are shown in Figure 26. For a path π, let $W(\pi)$ be the sum of its vertical edge weights, and define

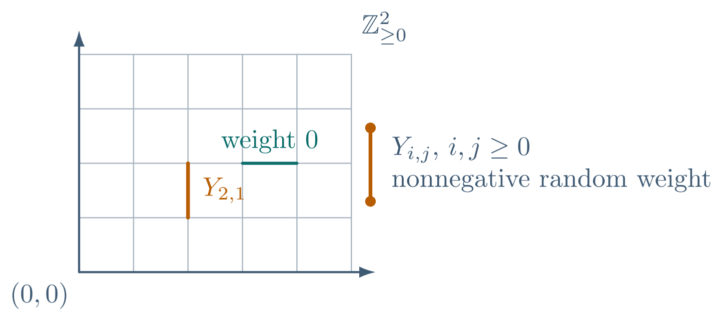

FIGURE 26

$$L(m,n) \coloneqq \min_{\pi:(0,0)\to(m,n)} W(\pi).$$

This model is a simplified form of directed first-passage percolation studied by Seppäläinen. Its large-scale behaviour follows from the subadditive ergodic theorem.

Theorem 5.1 Suppose that the variables $Y_{i,j}$, for $i,j \geq 0$, are independent and identically distributed and satisfy $\mathbb{E}[Y_{0,0}] < \infty$. Then there is a deterministic function $\Psi\colon [0,\infty)^2 \to [0,\infty)$ such that, almost surely and simultaneously for all $x, y \geq 0$,

$$\Psi(x,y) = \lim_{n\to\infty} \frac{L(\lfloor nx \rfloor, \lfloor ny \rfloor)}{n} = \lim_{n\to\infty} \frac{\mathbb{E}[L(\lfloor nx \rfloor, \lfloor ny \rfloor)]}{n}.$$

Moreover, Ψ has the following properties:

1. The function Ψ is nonincreasing in its first coordinate and nondecreasing in its second coordinate.
2. It is positively homogeneous:

$$\Psi(cx, cy) = c\Psi(x,y)$$

for all $c, x, y \geq 0$.

3. It is continuous on $[0,\infty)^2$.
4. It is convex.
5. It is subadditive:

$$\Psi(x_1+x_2, y_1+y_2) \le \Psi(x_1,y_1) + \Psi(x_2,y_2)$$

for all $x_1, x_2, y_1, y_2 \ge 0$.

Proof. Fix a lattice direction $z = (a,b) \in \mathbb{Z}^2_{\ge 0}$. Passage times between the points mz and nz form a stationary subadditive process: concatenating a path from mz to kz with one from kz to nz gives the required subadditivity, while translations of the independent environment give stationarity and ergodicity. The integrability assumption provides the required L^1-bound. The subadditive ergodic theorem therefore yields a deterministic limit

$$\lim_{n\to\infty} \frac{L(na,nb)}{n} = \inf_{n\ge 1} \frac{\mathbb{E}[L(na,nb)]}{n}$$

almost surely and in L^1.

Scaling integer directions gives the homogeneity in (2) first for nonnegative rational arguments. Concatenation of optimal paths similarly yields (5) on rational points. Combining rational homogeneity and subadditivity gives rational convexity:

$$\Psi\big(tu + (1-t)v\big) \le t\Psi(u) + (1-t)\Psi(v), \qquad t \in [0,1]\cap\mathbb{Q}.$$

Because horizontal edges have weight 0, appending horizontal steps shows that $L(m,n)$ is nonincreasing in m. Since vertical edge weights are nonnegative, reaching a higher level cannot lower the passage time. These observations give (1). They also allow arbitrary real points to be squeezed between rational points. On the axes, the strong law of large numbers supplies the required bounds. The rational limit thus extends continuously to the closed quadrant, proving (3); the same approximation extends rational convexity and subadditivity to all real points. This proves (4) and (5), as well as the simultaneous almost-sure convergence. □

The exponential model admits an explicit **shape function**.

Theorem 5.2 If the vertical edge weights have the exponential distribution with rate $\lambda > 0$, then

$$\Psi(x,y) = \frac{(\sqrt{x+y} - \sqrt{x})^2}{\lambda}.$$

To understand the geometry of this formula, fix $v = \Psi(x,y)$. Rearranging the equation gives

$$\sqrt{x+y} - \sqrt{x} = \sqrt{\lambda v}, \qquad x + y = x + 2\sqrt{\lambda v x} + \lambda v \qquad \text{and} \qquad y = 2\sqrt{\lambda v x} + \lambda v.$$

Thus each level set of Ψ is a square-root curve with vertical intercept λv, as shown in Figure 27. The proof of Theorem 5.2 begins with a local recursion. An optimal path to (i,j) must make

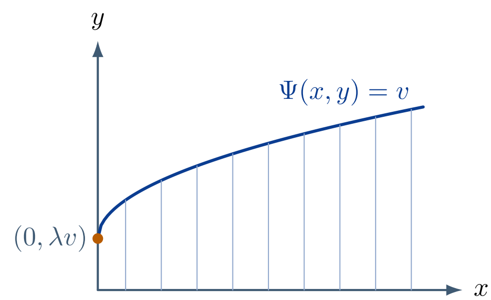

FIGURE 27

its final step either horizontally from $(i-1,j)$ or vertically from $(i,j-1)$. The two global route families are represented in Figure 29, and their final steps are isolated in Figure 28. It follows

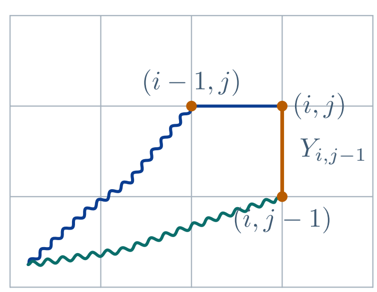

FIGURE 28

that

$$L(i,j) = \min\{L(i-1,j), L(i,j-1) + Y_{i,j-1}\}.$$

The passage times themselves are difficult to analyze directly, but their increments obey a tractable recursion. Define

$$X_{i,j} := L(i,j+1) - L(i,j) \geq 0 \qquad \text{and} \qquad Z_{i,j} := L(i,j) - L(i+1,j) \geq 0.$$

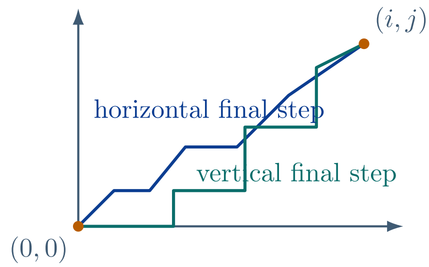

FIGURE 29

The first quantity is the vertical increment, and the second is the decrease in passage time gained by moving one unit to the right. Substituting the passage time recursion twice gives

$$X_{i,j} = \min\{X_{i-1,j} + Z_{i-1,j}, Y_{i,j}\}, \qquad i \geq 1,\ j \geq 0, \tag{5.1}$$

$$Z_{i,j} = \max\{0, X_{i,j-1} + Z_{i,j-1} - Y_{i+1,j-1}\}, \qquad i \geq 0,\ j \geq 1. \tag{5.2}$$

For the original model, $X_{0,j} = Y_{0,j}$ and $Z_{i,0} = 0$. These boundary increments are highlighted in Figure 30.

FIGURE 30

Theorem 5.3 Fix $\alpha, \lambda > 0$. Suppose that the boundary and bulk variables are independent and have laws

$$X_{0,j} \sim \mathrm{Exp}(\lambda + \alpha), \tag{5.3}$$

$$Z_{i,0} \sim \frac{\alpha}{\alpha+\lambda}\delta_0 + \frac{\lambda}{\alpha+\lambda}\mathrm{Exp}(\alpha), \tag{5.4}$$

$$Y_{i,j} \sim \mathrm{Exp}(\lambda). \tag{5.5}$$

Extend X and Z to the quadrant using the increment recursions (5.1) and (5.2). Then the increment field is stationary. In particular,

$$X_{i,j} \sim \mathrm{Exp}(\lambda+\alpha) \qquad \text{and} \qquad Z_{i,j} \sim \frac{\alpha}{\alpha+\lambda}\delta_0 + \frac{\lambda}{\alpha+\lambda}\mathrm{Exp}(\alpha)$$

for all $i, j \geq 0$.

Proof. The stationarity is a consequence of a local exponential identity. Let X, Z, and Y be independent with the three laws in (5.3)–(5.5), and put $S = X + Z$. For $s \geq 0$, the Laplace transform of S is

$$\mathbb{E}[e^{-sS}] = \frac{\lambda+\alpha}{\lambda+\alpha+s}\left(\frac{\alpha}{\alpha+\lambda} + \frac{\lambda}{\alpha+\lambda}\frac{\alpha}{\alpha+s}\right) = \frac{\alpha}{\alpha+s}.$$

Thus $S \sim \mathrm{Exp}(\alpha)$, independently of $Y \sim \mathrm{Exp}(\lambda)$. Define

$$X' := \min\{S, Y\} \qquad \text{and} \qquad Z' := (S-Y)^+.$$

The minimum of the two exponential clocks has rate $\lambda+\alpha$. Moreover, $Z' = 0$ when $S \leq Y$, an event of probability $\alpha/(\alpha+\lambda)$. Conditional on $S > Y$, the memoryless property gives $Z' \sim \mathrm{Exp}(\alpha)$. The exponential race decomposition also shows that X' and Z' are independent. Hence (X', Z') has the same product law as (X, Z).

Each lattice cell applies exactly this local transformation: its incoming increments and bulk weight determine the two outgoing increments. Propagating the product law cell by cell proves that every translated down-right boundary has the same increment law, which is the asserted stationarity. $\square$

We now use the stationary field from Theorem 5.3 to identify the shape function. Modify the edge weights as follows:

1. The vertical edge from $(0, j)$ to $(0, j+1)$ has weight $X_{0,j}$.
2. The vertical edge from (i, j) to $(i, j+1)$, for $i \geq 1$, has weight $Y_{i,j}$.
3. The horizontal edge from $(i, 0)$ to $(i+1, 0)$ has weight $-Z_{i,0}$.
4. Every other horizontal edge has weight 0.

The four types of edges are summarized in Figure 31. Let $L_\alpha(m, n)$ be the minimum weight of a

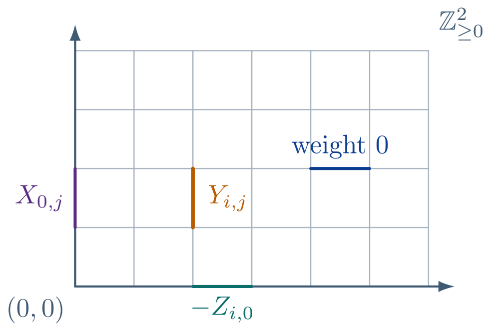

FIGURE 31

directed path from $(0,0)$ to (m,n) in this modified environment. By stationarity,

$$\begin{aligned} L_\alpha(m,n+1) - L_\alpha(m,n) &= X_{m,n} \sim \operatorname{Exp}(\lambda+\alpha), \\ L_\alpha(m,n) - L_\alpha(m+1,n) &= Z_{m,n} \\ &\qquad\qquad\qquad \sim \frac{\alpha}{\alpha+\lambda}\delta_0 + \frac{\lambda}{\alpha+\lambda}\operatorname{Exp}(\alpha). \end{aligned}$$

Consequently,

$$\mathbb{E}[X_{m,n}] = \frac{1}{\lambda+\alpha} \qquad \text{and} \qquad \mathbb{E}[Z_{m,n}] = \frac{\lambda}{\alpha(\lambda+\alpha)}.$$

Telescoping first in the vertical direction and then along the bottom boundary gives

$$\begin{aligned} \mathbb{E}[L_\alpha(m,n)] &= \sum_{j=0}^{n-1} \mathbb{E}\big[L_\alpha(m,j+1) - L_\alpha(m,j)\big] + \mathbb{E}[L_\alpha(m,0)] \\ &= n\,\mathbb{E}[X_{m,0}] - \sum_{i=0}^{m-1} \mathbb{E}[Z_{i,0}] \\ &= \frac{n}{\lambda+\alpha} - \frac{m\lambda}{\alpha(\lambda+\alpha)}. \end{aligned}$$

It follows that

$$\lim_{N\to\infty} \frac{\mathbb{E}[L_\alpha(\lfloor Nx\rfloor, \lfloor Ny\rfloor)]}{N} = \frac{y}{\lambda+\alpha} - \frac{\lambda x}{\alpha(\lambda+\alpha)} =: \varphi_\alpha(x,y).$$

Every path in the modified environment must enter the unmodified bulk either from the bottom boundary or from the left boundary. These two alternatives are illustrated in Figure 32. More precisely, decomposition at the first bulk vertex gives the distributional identity

$$L_\alpha(m,n) \overset{\mathrm{d}}{=} \min\left\{ \min_{0\le k\le m}\left(-\sum_{i=0}^{k-1} Z_{i,0} + L(m-k,n)\right),\right.$$

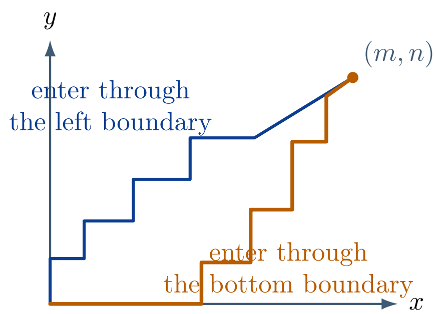

FIGURE 32

$$\min_{0\le k\le n}\left(\sum_{j=0}^{k-1} X_{0,j} + L(m,n-k)\right)\Bigg\},$$

where an empty sum is 0. Divide by N, use the strong law of large numbers for the boundary sums, and apply Theorem 5.1 to the bulk passage times. The result is the variational formula

$$\varphi_\alpha(x,y) = \min\Bigg\{\inf_{0\le t\le x}\left(-\frac{\lambda t}{\alpha(\alpha+\lambda)} + \Psi(x-t,y)\right), \inf_{0\le t\le y}\left(\frac{t}{\alpha+\lambda} + \Psi(x,y-t)\right)\Bigg\}. \tag{5.6}$$

Proof of Theorem 5.2. For fixed $\alpha > 0$ and $y \ge 0$, consider an auxiliary horizontal coordinate $X > 0$. The quantity

$$\varphi_\alpha(X,y) = \frac{y}{\lambda+\alpha} - \frac{\lambda X}{\alpha(\lambda+\alpha)}$$

is negative exactly when $\alpha y < \lambda X$. Choose X large enough that this inequality holds. The second infimum in (5.6) is nonnegative, so the first infimum must be the active branch. Writing

$$c_\alpha \coloneqq \frac{\lambda}{\alpha(\alpha+\lambda)}$$

and substituting $s = X - t$ gives

$$\inf_{0\le s\le X}\big(c_\alpha s + \Psi(s,y)\big) = \varphi_\alpha(X,y) + c_\alpha X = \frac{y}{\lambda+\alpha}.$$

The equality holds for every sufficiently large X, and therefore

$$\inf_{s\ge 0}\big(c_\alpha s + \Psi(s,y)\big) = \frac{y}{\lambda+\alpha}.$$

As α ranges over $(0,\infty)$, the slope c_α ranges over $(0,\infty)$. Since $x \mapsto \Psi(x,y)$ is convex, continuous, and nonincreasing by Theorem 5.1, convex duality recovers it from these infimal transforms:

$$\Psi(x,y) = \sup_{\alpha>0} \left\{ \frac{y}{\lambda+\alpha} - \frac{\lambda x}{\alpha(\lambda+\alpha)} \right\}.$$

For $x, y > 0$, differentiation shows that the maximum occurs at

$$\alpha_\star = \frac{\lambda\big(x + \sqrt{x(x+y)}\big)}{y}.$$

Substitution and simplification yield

$$\Psi(x,y) = \frac{x+y+x-2\sqrt{x(x+y)}}{\lambda} = \frac{\big(\sqrt{x+y}-\sqrt{x}\big)^2}{\lambda}.$$

When $x = 0$ or $y = 0$, the same formula follows by continuity, or directly from the corresponding one-dimensional passage times. $\square$

Lecture 6 — Wednesday, February 03, 2021

2. Pitman Transforms, the Robinson–Schensted–Knuth Correspondence, and Random Matrices

Pitman Transforms and the Robinson–Schensted–Knuth Correspondence

We begin with two horizontal rows of edge weights, $(w_{1,j})_{j\geq 1}$ and $(w_{2,j})_{j\geq 1}$. Vertical edges have weight 0, and directed paths move to the right or upward. A representative last-passage path is highlighted in Figure 33. Define the **cumulative weight functions**

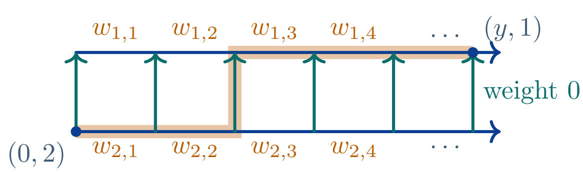

FIGURE 33

$$f_i(y) \coloneqq \sum_{j=1}^{y} w_{i,j}, \qquad i \in \{1,2\}.$$

With the negative sign convention for the induced directed metric, the last-passage value from $(0,2)$ to $(y,1)$ is

$$d_f\big((0,2),(y,1)\big) = \max_{0\le x\le y}\big(f_2(x) + f_1(y) - f_1(x)\big) = f_1(y) + \max_{0\le x\le y}\big(f_2(x) - f_1(x)\big).$$

This motivates the running gap

$$G_f(y) \coloneqq \max_{0\le x\le y}\big(f_2(x) - f_1(x)\big)$$

and the **two-line Pitman transform**

$$(Wf)_1(y) \coloneqq f_1(y) + G_f(y) \qquad \text{and} \qquad (Wf)_2(y) \coloneqq f_2(y) - G_f(y).$$

In particular,

$$(Wf)_1 + (Wf)_2 = f_1 + f_2.$$

The first transformed coordinate is the last-passage value under the negative sign convention $d_f\big((0,2),(y,1)\big)$. The second can equivalently be interpreted using the directed distance with positive sign from $(0,1)$ to $(y,2)$.

If f_1 and f_2 are independent Bernoulli random walks with parameters $p_1 > p_2$, then the transformed pair has the law of the two walks conditioned never to cross after time zero. When $p_1 = p_2$, the same phrase refers to the corresponding Doob transform, or equivalently to the limit of finite-horizon conditioning, because the infinite-horizon event itself has probability zero. The discrete bijection developed in the next subtopic is the first step toward proving this statement.

One geometric interpretation is to place the increments of f_1 on top of the graph of f_2. The running gap is precisely the smallest upward translation required up to time y, as depicted in Figure 34. Equivalently, $(Wf)_1$ is the **Skorokhod reflection** of f_1 above the lower barrier f_2:

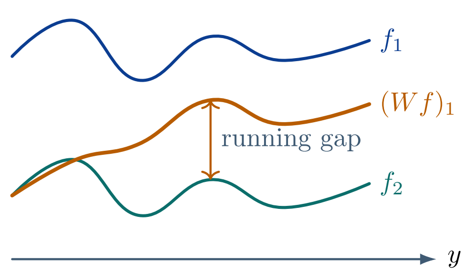

Figure 34

the path is pushed upward only when it would otherwise cross the barrier. This viewpoint is shown in Figure 35. The two-line Pitman transform has two fundamental structural properties.

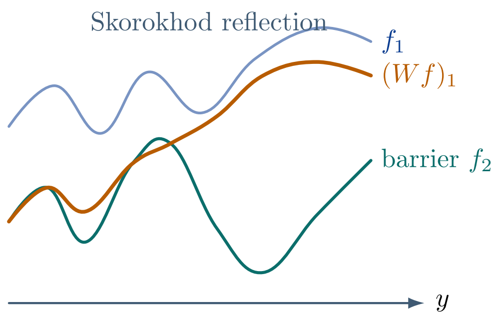

FIGURE 35

Proposition 6.1 The Pitman transform has the following properties:

1. It is idempotent: $W(Wf) = Wf$.
2. It preserves two-line last-passage distances. More precisely, whenever $0 \le x \le y$,

$$d_f\big((x,2),(y,1)\big) = d_{Wf}\big((x,2),(y,1)\big).$$

Proof. For (1), the definition gives $(Wf)_1 \ge f_1$ and $(Wf)_2 \le f_2$. In particular, $(Wf)_1 \ge (Wf)_2$, so the running gap of the already transformed pair is identically 0. Applying W again therefore makes no change.

For (2), write

$$H(t) := f_2(t) - f_1(t) \qquad \text{and} \qquad G(t) := \sup_{0\le s\le t} H(s).$$

Then $(Wf)_1 = f_1 + G$ and $(Wf)_2 = f_2 - G$. A direct use of the last-passage formula reduces the desired equality to the running maximum identity

$$G(y) + G(x) + \sup_{x\le z\le y} \big(H(z) - 2G(z)\big) = \sup_{x\le z\le y} H(z).$$

If H does not exceed $G(x)$ on $[x,y]$, then G stays constant there and the identity is immediate. Otherwise, consider the first time after x at which H reaches a new running maximum. Before that time the largest value of $H - 2G$ approaches $-G(x)$, and after that time it cannot exceed this value. The displayed identity follows in the second case as well. $\square$

Idempotence means that W maps onto the ordered pairs, but it is not injective without one additional endpoint coordinate. To make this precise, let

$$\mathcal{F}_t := \{f\colon [0,t] \to \mathbb{R} : f \text{ is continuous and } f(0) = 0\}$$

and

$$\mathcal{F}_t^{\downarrow 2} := \{(g_1, g_2) \in \mathcal{F}_t^2 : g_1 \geq g_2\}.$$

Define

$$\widetilde{W}_t(f_1, f_2) := \big(W(f_1, f_2), f_2(t)\big)$$

and set

$$\mathcal{F}_t^\star := \Big\{\big((g_1, g_2), z\big) \in \mathcal{F}_t^{\downarrow 2} \times \mathbb{R} : g_2(t) \leq z \leq g_1(t)\Big\}.$$

Proposition 6.2 The map

$$\widetilde{W}_t\colon \mathcal{F}_t^2 \to \mathcal{F}_t^\star$$

is a bijection.

Proof. The inverse problem is summarized in Figure 36: the ordered output pair and one original endpoint must determine the two input paths. The transform commutes with adding the same

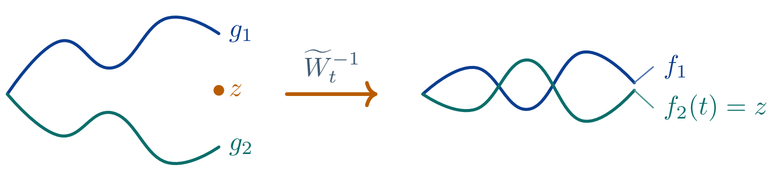

Figure 36

function to both coordinates, and it preserves their sum. Given $\big((g_1, g_2), z\big) \in \mathcal{F}_t^\star$, define

$$a := \frac{g_1 + g_2}{2}, \qquad H := \frac{g_1 - g_2}{2} \qquad \text{and} \qquad z_0 := z - a(t).$$

Then $H \geq 0$ and

$$-H(t) \leq z_0 \leq H(t).$$

It is therefore enough to invert the centred transform, in which the input has the form $(-X, X)$ and the output has the form $(H, -H)$. In this case, if

$$M(y) := \max_{0 \leq x \leq y} X(x),$$

then the Pitman formula becomes

$$H(y) = 2M(y) - X(y).$$

The two contributions on the right are isolated schematically in Figure 37.

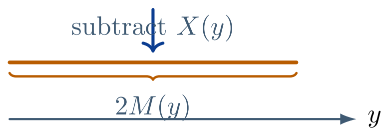

Figure 37

Theorem 6.3 (Pitman's $2M - X$ theorem) If X is a **simple symmetric random walk** on $\mathbb{Z}$ and $M_n = \max_{0 \le k \le n} X_k$, then the transformed process $(2M_n - X_n)_{n \ge 0}$ is Markov.

The reflection underlying Theorem 6.3 is illustrated in Figure 38. To construct the inverse, first

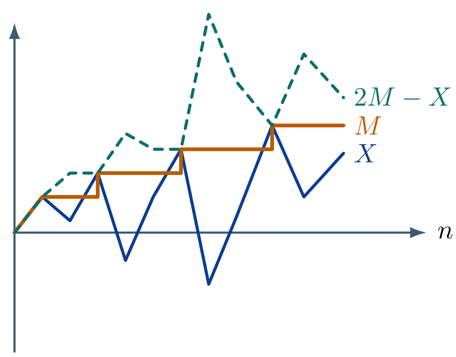

Figure 38

recover the terminal maximum from H and $z_0 = X(t)$:

$$M(t) = \frac{H(t) + z_0}{2}.$$

For $0 \le x \le t$, define

$$M(x) := \min\left\{M(t),\ \min_{x \le y \le t} H(y)\right\} \qquad \text{and} \qquad X(x) := 2M(x) - H(x).$$

The backward running minimum construction of M is depicted in Figure 39. One checks directly that $M(x) = \max_{0 \le s \le x} X(s)$, $X(t) = z_0$, and $H = 2M - X$. Hence the unique centred preimage

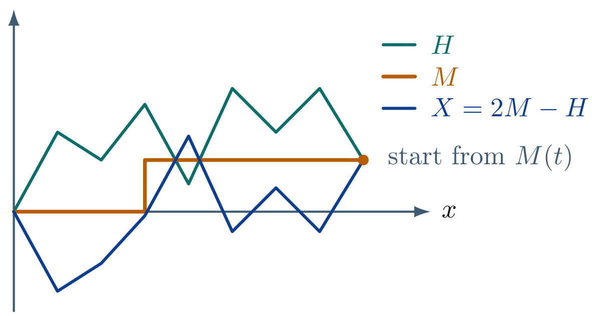

FIGURE 39

is $(-X, X)$. Adding the common function a back to both coordinates gives

$$(f_1, f_2) = (a - X, a + X).$$

It satisfies $f_2(t) = z$ and $W(f_1, f_2) = (g_1, g_2)$, completing the construction and uniqueness proof. $\square$

The continuous bijection restricts to Bernoulli paths. For an integer $t \geq 0$, let

$$\mathcal{G}_t \coloneqq \{f \colon \{0, 1, \ldots, t\} \to \mathbb{Z} : f(0) = 0,\ f(j+1) - f(j) \in \{0, 1\}\}.$$

We identify each element with its piecewise linear interpolation on $[0, t]$, and define $\mathcal{G}_t^{\downarrow 2}$ and $\mathcal{G}_t^{\star}$ by imposing the same ordering and endpoint conditions as in the continuous setting. The endpoint coordinate in $\mathcal{G}_t^{\star}$ is required to be an integer. By Proposition 6.2,

$$\widetilde{W}_t \colon \mathcal{G}_t^2 \to \mathcal{G}_t^{\star}$$

is a bijection of finite sets. For $t = 2$, there are four possible single-line paths and hence $4^2 = 16$ input pairs; the output set also has 16 elements. The finite correspondence is represented schematically in Figure 40. The finite bijection leads naturally to a Gibbs characterization of

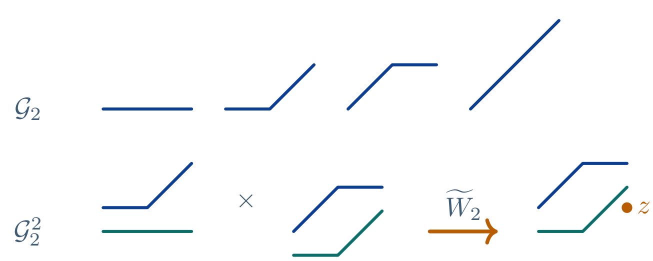

FIGURE 40

nonintersecting random walks. Consider an ordered **Bernoulli line ensemble** of k paths

$$X_1 \geq X_2 \geq \cdots \geq X_k.$$

Fix integer times $0 \leq t_1 < t_2$ and line indices $1 \leq i_1 \leq i_2 \leq k$. The corresponding space-time block is shown in Figure 41.

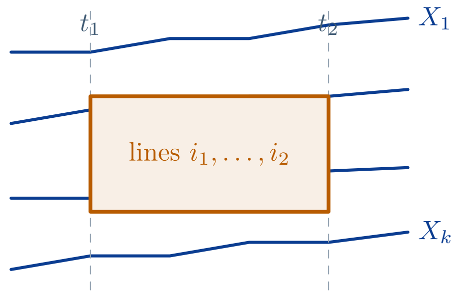

Figure 41

Definition 6.4 An ordered Bernoulli line ensemble has the **Gibbs property** if, conditional on all paths outside the block

$$\{(i, s) : i_1 \leq i \leq i_2,\ t_1 \leq s \leq t_2\},$$

the paths inside the block are uniformly distributed over all Bernoulli paths that match the revealed entrance and exit values and remain ordered between the neighbouring exterior paths.

The conditioning fixes the number of upward steps made by each bridge. Consequently, every admissible bridge configuration has the same Bernoulli weight, which explains why the conditional distribution in Definition 6.4 is uniform even when the lines have different step parameters.

Exercise 6.5 Verify the Gibbs property directly when $p_1 > p_2 > \cdots > p_k$, using the fact that conditioning on both endpoints fixes the number of upward steps on each line.

Solution. Fix a spacetime block with time interval $[t_1, t_2]$, and condition on the paths outside the block. This conditioning fixes the entrance and exit heights of every resampled line. If line i makes m_i upward steps inside the block, then it makes $t_2 - t_1 - m_i$ horizontal steps there, and its Bernoulli weight is

$$p_i^{m_i}(1 - p_i)^{t_2 - t_1 - m_i}.$$

The endpoints determine m_i, so this weight is the same for every bridge of line i with the prescribed endpoints. Independence of the lines shows that every admissible configuration in the block has the same product weight

$$\prod_{i=i_1}^{i_2} p_i^{m_i}(1-p_i)^{t_2-t_1-m_i}.$$

Conditioning further on the requirement that the resampled bridges remain ordered and lie between their revealed neighbouring paths merely restricts the set of admissible configurations; it does not change their common weight. The conditional distribution is therefore uniform on that set, which is precisely the Gibbs property in Definition 6.4. □

Theorem 6.6 Given parameters

$$p_1 \geq p_2 \geq \cdots \geq p_k,$$

there is a unique law on ordered k-tuples of Bernoulli paths $X_1 \geq \cdots \geq X_k$, starting at the origin, that has the Gibbs property and satisfies

$$\lim_{t\to\infty} \frac{X_i(t)}{t} = p_i$$

almost surely for each $i \in \{1, \ldots, k\}$.

This law is called the law of **weakly nonintersecting Bernoulli walks with drift $\boldsymbol{p}$**.

Proof outline for two lines. Take independent Bernoulli paths f_1, f_2 with parameters $p_1 \geq p_2$, and apply the finite bijection $\widetilde{W}_t$. The product probability of an input pair depends only on its two endpoints. Therefore, after conditioning on the transformed boundary data, all admissible ordered pairs have equal mass. This is exactly the finite-volume Gibbs property.

The transform preserves the sum of the two paths. If $p_1 > p_2$, the running gap G_f is almost surely eventually constant; if $p_1 = p_2$, the strong law gives $G_f(t)/t \to 0$. In either case, the transformed paths have slopes p_1 and p_2. Passing to a consistent infinite-volume limit proves existence. Uniqueness follows by sandwiching finite-volume Gibbs measures between measures with ordered boundary data; the required monotonicity is established in Proposition 6.7. □

Recall that a law μ stochastically dominates a law ν if they admit a coupling (U, V) with $U \geq V$ almost surely. For example, Bernoulli(p_1) dominates Bernoulli(p_2) whenever $p_1 \geq p_2$: take a single uniform random variable U on $[0, 1]$ and use the indicators of $\{U \leq p_1\}$ and $\{U \leq p_2\}$.

The bridge version of this statement is the key monotonicity input. It is illustrated in Figure 42.

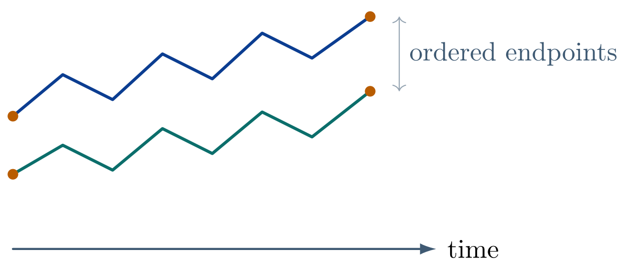

FIGURE 42

Proposition 6.7 The uniform law of a Bernoulli bridge is stochastically nondecreasing in each endpoint. The same conclusion holds for a collection of ordered Bernoulli bridges when either exterior boundary path is moved upward.

Proof. On a finite time interval, consider the local resampling dynamics that chooses an interior time and resamples the path height there uniformly among the values compatible with the two neighbouring times and the prescribed boundaries. The uniform bridge law is stationary for this dynamics.

Couple two such dynamics by choosing the same update locations and the same uniform random variables. If the endpoints and exterior boundaries of the first system lie above those of the second, every local admissible interval in the first system has endpoints at least as high as the corresponding interval in the second. The update using a common uniform variable therefore preserves the pointwise order. Starting from the maximal and minimal admissible configurations and letting the dynamics converge yields an ordered coupling of the two uniform bridge measures. This is precisely stochastic monotonicity. $\square$

Lecture 7 — Friday, February 05, 2021

Dyson Brownian Motion

Let $B_1, \dots, B_n$ be independent standard Brownian motions. A vector of independent Brownian particles may collide, but Dyson Brownian motion adds a singular drift that repels every pair of coordinates.

Definition 7.1 Let

$$\mathsf{W}_n^\circ := \{\boldsymbol{x} \in \mathbb{R}^n : x_1 > x_2 > \cdots > x_n\}$$

be the interior of the **type-A Weyl chamber**. Starting from $\boldsymbol{\lambda}(0) \in \mathsf{W}_n^\circ$, **Dyson Brownian motion** is the solution of

$$\lambda_i(t) = \lambda_i(0) + B_i(t) + \int_0^t \sum_{\substack{j=1 \\ j \neq i}}^{n} \frac{1}{\lambda_i(s) - \lambda_j(s)} \, \mathrm{d}s, \qquad i \in \{1, \dots, n\}.$$

The Brownian term supplies independent fluctuations. As recalled in Figure 43, an increment over an interval of length $T - s$ is centered Gaussian with variance $T - s$. For an upper particle λ_i,

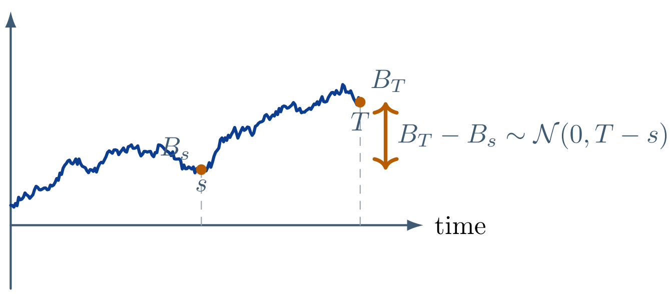

FIGURE 43

the terms arising from lower particles are positive; for a lower particle, the corresponding terms are negative. Thus the singular drift pushes close particles apart. The effect for two particles is illustrated in Figure 44. Two descriptions make Definition 7.1 especially important:

1. It is the Doob transform of n independent Brownian motions killed when two coordinates collide.
2. It is the ordered eigenvalue process of Hermitian matrix Brownian motion. In particular, at

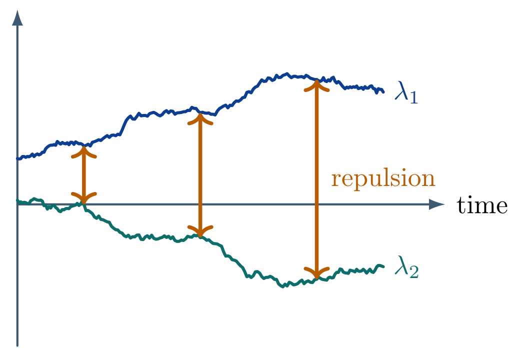

FIGURE 44

time 1 the eigenvalues have the Gaussian unitary ensemble law.

We first establish (1). This requires a short review of stochastic calculus and Doob transforms.

Let B be a one-dimensional Brownian motion, and let σ be an adapted process satisfying the usual square-integrability condition. Recall that the **Itô integral** is the L^2-limit of left-endpoint sums

$$\int_0^t \sigma(s)\,\mathrm{d}B_s := \lim_{|\Pi|\to 0} \sum_{\ell=0}^{N-1} \sigma(t_\ell)\big(B_{t_{\ell+1}} - B_{t_\ell}\big).$$

Proposition 7.2 The process

$$M_t := \int_0^t \sigma(s)\,\mathrm{d}B_s$$

is a martingale.

Now suppose that $\boldsymbol{X}$ is an $\mathbb{R}^n$-valued diffusion satisfying

$$\mathrm{d}\boldsymbol{X}_t = \boldsymbol{\mu}(\boldsymbol{X}_t)\,\mathrm{d}t + \mathrm{d}\boldsymbol{B}_t,$$

where $\boldsymbol{B}$ is standard n-dimensional Brownian motion and $\boldsymbol{\mu}$ is the drift vector field. For $f \in C^2(\mathbb{R}^n)$, Itô's formula gives

$$\mathrm{d}f(\boldsymbol{X}_t) = \left(\langle \boldsymbol{\mu}(\boldsymbol{X}_t), \nabla f(\boldsymbol{X}_t)\rangle + \frac{1}{2}\Delta f(\boldsymbol{X}_t)\right)\mathrm{d}t + \langle \nabla f(\boldsymbol{X}_t), \mathrm{d}\boldsymbol{B}_t\rangle.$$

The stochastic integral in this decomposition is a martingale by Proposition 7.2.

The next ingredient is the **Doob h-transform**. Let $\boldsymbol{X}$ be a time-homogeneous Markov process, possibly killed at the boundary of its state space, with transition kernel

$$\mathbb{P}_t(\boldsymbol{x}, \mathrm{d}\boldsymbol{y}) := \mathbb{P}_{\boldsymbol{x}}(\boldsymbol{X}_t \in \mathrm{d}\boldsymbol{y}).$$

The kernel may therefore be sub-Markov. Suppose that $h > 0$ is harmonic for the killed process, meaning

$$\int h(\boldsymbol{y})\, \mathbb{P}_t(\boldsymbol{x}, \mathrm{d}\boldsymbol{y}) = h(\boldsymbol{x}).$$

It has transition kernel

$$\widetilde{\mathbb{P}}_t(\boldsymbol{x}, \mathrm{d}\boldsymbol{y}) := \frac{h(\boldsymbol{y})}{h(\boldsymbol{x})}\mathbb{P}_t(\boldsymbol{x}, \mathrm{d}\boldsymbol{y}).$$

Harmonicity ensures that this is a probability kernel. When $h(x)$ is the probability of a future event starting from x, the formula is simply Bayes' rule and describes conditioning on that event.

The generator of the diffusion $\boldsymbol{X}$ is

$$Lf(\boldsymbol{x}) := \lim_{t\downarrow 0} \frac{\mathbb{E}_{\boldsymbol{x}}[f(\boldsymbol{X}_t)] - f(\boldsymbol{x})}{t} = \langle \boldsymbol{\mu}(\boldsymbol{x}), \nabla f(\boldsymbol{x})\rangle + \frac{1}{2}\Delta f(\boldsymbol{x}).$$

For the transformed process,

$$\widetilde{L}f(\boldsymbol{x}) = \frac{1}{h(\boldsymbol{x})}L(hf)(\boldsymbol{x}) = \langle \boldsymbol{\mu}(\boldsymbol{x}) + \nabla \log h(\boldsymbol{x}), \nabla f(\boldsymbol{x})\rangle + \frac{1}{2}\Delta f(\boldsymbol{x}),$$

where the second equality uses $Lh = 0$. Thus the Doob transform adds the drift $\nabla \log h$:

$$\mathrm{d}\widetilde{\boldsymbol{X}}_t = \big(\boldsymbol{\mu}(\widetilde{\boldsymbol{X}}_t) + \nabla \log h(\widetilde{\boldsymbol{X}}_t)\big)\,\mathrm{d}t + \mathrm{d}\boldsymbol{B}_t.$$

For Brownian motion in the Weyl chamber, the relevant harmonic function is the Vandermonde determinant.

Definition 7.3 For $\boldsymbol{x} \in \mathbb{R}^n$, define the **Vandermonde determinant** by

$$\Delta_n(\boldsymbol{x}) := \prod_{1\le i<j\le n} (x_i - x_j).$$

On W_n°, this function is strictly positive, and it vanishes precisely on the collision boundary.

Proposition 7.4 The Vandermonde determinant satisfies the following identities.

1. One has $\Delta_n(\boldsymbol{x}) = 0$ exactly when two coordinates of $\boldsymbol{x}$ agree.
2. For $\boldsymbol{x} \in \mathsf{W}_n^\circ$,
$$\partial_{x_i} \log \Delta_n(\boldsymbol{x}) = \sum_{\substack{j=1 \\ j\neq i}}^{n} \frac{1}{x_i - x_j}.$$
3. The Vandermonde determinant is harmonic:
$$\Delta\Delta_n(\boldsymbol{x}) = \sum_{i=1}^{n} \partial_{x_i}^2 \Delta_n(\boldsymbol{x}) = 0.$$

Proof. Property (1) follows immediately from the product in Definition 7.3. Logarithmic differentiation of that product proves (2).

For (3), use the determinant representation
$$\Delta_n(\boldsymbol{x}) = \det\left[x_i^{n-j}\right]_{i,j=1}^{n}.$$
This is an alternating polynomial of the smallest possible total degree, namely $n(n-1)/2$. Its Laplacian is again alternating but has total degree two less. Every alternating polynomial is divisible by Δ_n, so no nonzero alternating polynomial can have this smaller degree. Hence the Laplacian vanishes. □

Theorem 7.5 Dyson Brownian motion is the Doob Δ_n-transform of n independent Brownian motions killed upon leaving W_n°. Equivalently, it is Brownian motion conditioned, in the sense of the Doob Δ_n-transform, never to experience a collision.

Proof. For $n = 1$, the Vandermonde determinant is identically one and the claim reduces to ordinary Brownian motion. Assume henceforth that $n \geq 2$. Let $\boldsymbol{B}$ be standard Brownian motion in $\mathbb{R}^n$, started from $\boldsymbol{x} \in \mathsf{W}_n^\circ$, and let τ_0 be the first collision time. By Itô's formula and (3),
$$\mathrm{d}\Delta_n(\boldsymbol{B}_t) = \langle \nabla\Delta_n(\boldsymbol{B}_t), \mathrm{d}\boldsymbol{B}_t\rangle$$
up to time τ_0. Thus the stopped Vandermonde determinant is a local martingale.

For $c > \Delta_n(\boldsymbol{x})$, let
$$\tau_c := \inf\{t \geq 0 : \Delta_n(\boldsymbol{B}_t) = c\}.$$

Optional stopping for the bounded process $\Delta_n(\boldsymbol{B}_{t\wedge\tau_0\wedge\tau_c})$ gives the gambler's-ruin identity

$$\mathbb{P}_{\boldsymbol{x}}(\tau_c < \tau_0) = \frac{\Delta_n(\boldsymbol{x})}{c}.$$

Here $\tau_0 < \infty$ almost surely because the difference of any two coordinate Brownian motions is a one-dimensional Brownian motion and hence hits zero almost surely; this justifies passage to the unbounded stopping time after first stopping at a deterministic time. Consequently, conditioning on $\tau_c < \tau_0$ is the Doob transform with $h_c = \Delta_n/c$. Its additional drift is

$$\partial_{x_i} \log h_c(\boldsymbol{x}) = \partial_{x_i} \log \Delta_n(\boldsymbol{x}) = \sum_{\substack{j=1\\ j\neq i}}^{n} \frac{1}{x_i - x_j},$$

where (2) was used. This drift is independent of c. Letting $c \to \infty$ produces the standard Q-process conditioning the killed Brownian motion to survive forever. Its stochastic differential equation is exactly the one in Definition 7.1. □

Gaussian Unitary Ensembles

Definition 7.6 A random Hermitian matrix $\mathbf{M} = [\xi_{i,j}]_{i,j=1}^n$ has the **Gaussian unitary ensemble** law, denoted $\mathbf{M} \sim \mathrm{GUE}(n)$, if

1. The diagonal entries are independent real random variables with $\xi_{i,i} \sim \mathcal{N}(0,1)$.
2. For $i < j$, the entries $\xi_{i,j}$ are independent complex Gaussian random variables of the form
$$\xi_{i,j} = \frac{X_{i,j} + \mathrm{i}Y_{i,j}}{\sqrt{2}} \qquad \text{and} \qquad X_{i,j}, Y_{i,j} \sim \mathcal{N}(0,1),$$
and $\xi_{j,i} = \overline{\xi_{i,j}}$.
3. All real random variables appearing in (1) and (2) are mutually independent.

Proposition 7.7 Let $\mathcal{H}_n$ denote the real vector space of $n \times n$ Hermitian matrices. Then the following statements hold.

1. The map
$$\mathbf{A} \longmapsto \big((a_{i,i})_{i=1}^n, (\sqrt{2}\,\mathrm{Re}\,a_{i,j}, \sqrt{2}\,\mathrm{Im}\,a_{i,j})_{i<j}\big)$$
identifies $\mathcal{H}_n$ isometrically with $\mathbb{R}^{n^2}$, where
$$\langle \mathbf{A}, \mathbf{B}\rangle_{\mathcal{H}_n} = \mathrm{Tr}(\mathbf{AB}) \qquad \text{and} \qquad \|\mathbf{A}\|_{\mathcal{H}_n}^2 = \mathrm{Tr}(\mathbf{A}^2) = \sum_{i,j=1}^n |a_{i,j}|^2.$$
2. Under the identification in (1), a matrix $\mathbf{M} \sim \mathrm{GUE}(n)$ is a standard Gaussian vector in $\mathbb{R}^{n^2}$. Its density with respect to Lebesgue measure on $\mathcal{H}_n$ is
$$c_n \exp\left(-\frac{1}{2}\,\mathrm{Tr}(\mathbf{M}^2)\right).$$
3. For every deterministic unitary matrix $\mathbf{U}$,
$$\mathbf{UMU}^* \overset{\mathrm{d}}{=} \mathbf{M}.$$
4. If $\mathbf{A}(t)$ is standard Brownian motion in the Euclidean space $\mathcal{H}_n$, then
$$\mathbf{A}(1) - \mathbf{A}(0) \sim \mathrm{GUE}(n).$$

Proof. The coordinate formula in (1) follows by expanding $\mathrm{Tr}(\mathbf{AB})$ and using Hermitian symmetry. The entry normalizations in Definition 7.6 then make all n^2 real coordinates independent standard Gaussians, proving (2).

Conjugation $\mathbf{A} \mapsto \mathbf{UAU}^*$ preserves $\mathrm{Tr}(\mathbf{A}^2)$, so it is an orthogonal transformation of $\mathcal{H}_n$. The standard Gaussian law is orthogonally invariant, which proves (3). Finally, (4) is the defining time-one increment law of standard Brownian motion in $\mathbb{R}^{n^2}$. $\square$

Theorem 7.8 Let $\mathbf{A}(t)$ be standard Brownian motion in $\mathcal{H}_n$, started from a matrix with simple ordered eigenvalues
$$\lambda_1(0) > \lambda_2(0) > \cdots > \lambda_n(0).$$

Then its ordered eigenvalues satisfy

$$\mathrm{d}\lambda_i(t) = \mathrm{d}B_i(t) + \sum_{\substack{j=1\\ j\neq i}}^{n} \frac{\mathrm{d}t}{\lambda_i(t)-\lambda_j(t)},$$

where $B_1,\dots,B_n$ are independent standard Brownian motions. Thus the eigenvalue process is Dyson Brownian motion.

Proof. Before the first eigenvalue collision, the map $\mathbf{A}\mapsto\lambda_i(\mathbf{A})$ is smooth. Let $\boldsymbol{u}_1,\dots,\boldsymbol{u}_n$ be an orthonormal eigenbasis of $\mathbf{A}$. The Hadamard variation formulas state that, for $\mathbf{G}\in\mathcal{H}_n$,

$$D\lambda_i(\mathbf{A})[\mathbf{G}] = \boldsymbol{u}_i^*\mathbf{G}\boldsymbol{u}_i,$$

$$D^2\lambda_i(\mathbf{A})[\mathbf{G},\mathbf{G}] = 2\sum_{\substack{j=1\\ j\neq i}}^{n} \frac{\left|\boldsymbol{u}_j^*\mathbf{G}\boldsymbol{u}_i\right|^2}{\lambda_i-\lambda_j}.$$

Apply Itô's formula to $\lambda_i(\mathbf{A}(t))$. Its martingale part is

$$\mathrm{d}M_i(t) = \boldsymbol{u}_i^*\,\mathrm{d}\mathbf{A}(t)\,\boldsymbol{u}_i.$$

The gradient of λ_i is the rank-one projector $\boldsymbol{u}_i\boldsymbol{u}_i^*$. Therefore,

$$\mathrm{d}\langle M_i, M_j\rangle_t = \mathrm{Tr}(\boldsymbol{u}_i\boldsymbol{u}_i^*\boldsymbol{u}_j\boldsymbol{u}_j^*)\,\mathrm{d}t = \delta_{i,j}\,\mathrm{d}t.$$

Lévy's characterization shows that $M_1,\dots,M_n$ are independent standard Brownian motions.

For the drift, let $\mathcal{E}$ be any orthonormal basis of $\mathcal{H}_n$. The second variation formula gives

$$\frac{1}{2}\Delta\lambda_i(\mathbf{A}) = \frac{1}{2}\sum_{\mathbf{G}\in\mathcal{E}} D^2\lambda_i(\mathbf{A})[\mathbf{G},\mathbf{G}] = \sum_{\substack{j=1\\ j\neq i}}^{n} \frac{1}{\lambda_i-\lambda_j}\sum_{\mathbf{G}\in\mathcal{E}}\left|\boldsymbol{u}_j^*\mathbf{G}\boldsymbol{u}_i\right|^2.$$

By unitary invariance, choose the basis after rotating $\boldsymbol{u}_i$ and $\boldsymbol{u}_j$ to coordinate vectors. The diagonal directions and the normalized real and imaginary off-diagonal directions then give

$$\sum_{\mathbf{G}\in\mathcal{E}}\left|\boldsymbol{u}_j^*\mathbf{G}\boldsymbol{u}_i\right|^2 = 1.$$

Consequently,

$$\frac{1}{2}\Delta\lambda_i(\mathbf{A}) = \sum_{\substack{j=1\\ j\neq i}}^{n} \frac{1}{\lambda_i - \lambda_j}.$$

Combining the martingale and drift terms proves the stated stochastic differential equation. By Theorem 7.5, the eigenvalues do not collide, so the calculation extends for all positive times. □

Lecture 8 — Wednesday, February 10, 2021

Pitman Transforms and Brownian Scaling

The two-line Pitman transform reflects a pair of paths into the ordered chamber. For an input pair $f = (f_1, f_2)$, it preserves the sum of the coordinates while transferring the running gap from the lower path to the upper path. Thus

$$(Wf)_1(t) = f_1(t) + \sup_{0\le s\le t} \big(f_2(s) - f_1(s)\big),$$
$$(Wf)_2(t) = f_2(t) - \sup_{0\le s\le t} \big(f_2(s) - f_1(s)\big).$$

The reflection is represented schematically in Figure 45. The upper coordinate also has a last-

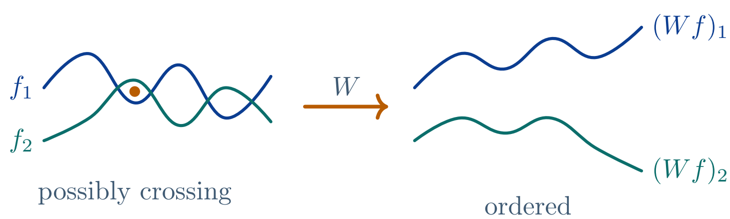

Figure 45

passage interpretation:

$$(Wf)_1(t) = \sup_{0\le s\le t} \big(f_2(s) + f_1(t) - f_1(s)\big).$$

This formula makes the transform compatible with diffusive scaling. Let X_1 and X_2 be independent Bernoulli random walks with parameter $1/2$, extended from the integers by linear interpolation, and define, for $i \in \{1,2\}$,

$$\widehat{X}_i^{(a)}(s) := \frac{X_i(a^2 s) - a^2 s/2}{a/2}, \qquad 0 \le s \le 1.$$

By Donsker's theorem,

$$\widehat{\boldsymbol{X}}^{(a)} \xrightarrow[a\to\infty]{\text{law}} \boldsymbol{B} \quad \text{in } C([0,1])^2,$$

where $\boldsymbol{B} = (B_1, B_2)$ consists of independent standard Brownian motions. The map W is continuous in the uniform topology because the running maximum map is Lipschitz. Consequently, the continuous mapping theorem gives

$$W\widehat{\boldsymbol{X}}^{(a)} \xrightarrow[a\to\infty]{\text{law}} W\boldsymbol{B} \quad \text{in } C([0,1])^2$$

for the pair of walks and the corresponding pair of Brownian motions.

The discrete Pitman representation identifies the transformed Bernoulli walks with the weakly nonintersecting Gibbs ensemble of Theorem 6.6. Passing to the diffusive limit therefore identifies WB with a pair of Brownian motions conditioned not to intersect, understood through the Doob transform and its entrance law at the origin. In particular, the last-passage value across two independent Brownian motions has the same law as the upper line of this conditioned pair.

The structural facts used here were proved in Propositions 6.1 and 6.2: the transform is idempotent, and its restriction to $[0,t]$ becomes a bijection after retaining the endpoint $f_2(t)$. We now give the detailed coupling and finite-volume arguments behind the Bernoulli Gibbs characterization.

Bernoulli Bridges and Monotone Couplings

Definition 8.1 Fix an integer $t \ge 0$ and integers x, y satisfying $x \le y \le x+t$. A **Bernoulli bridge** from $(0,x)$ to (t,y) is a uniformly chosen path

$$\gamma\colon \{0,1,\dots,t\} \longrightarrow \mathbb{Z}$$

such that $\gamma(0) = x$, $\gamma(t) = y$, and every increment belongs to $\{0,1\}$.

A typical bridge is shown in Figure 46. Its number of upward steps is fixed by its endpoints, namely $y - x$. By Proposition 6.7, the uniform bridge law is stochastically nondecreasing in its endpoints. The corresponding statement for an ordered line ensemble allows the entrance data, exit data, or exterior boundary paths to be raised, as suggested by Figure 47. The following corner-flip construction, due to Corwin and Hammond, provides a concrete proof of this monotonicity.

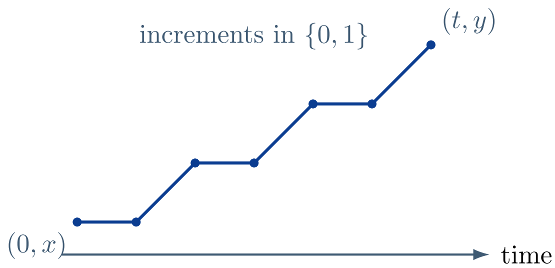

Figure 46

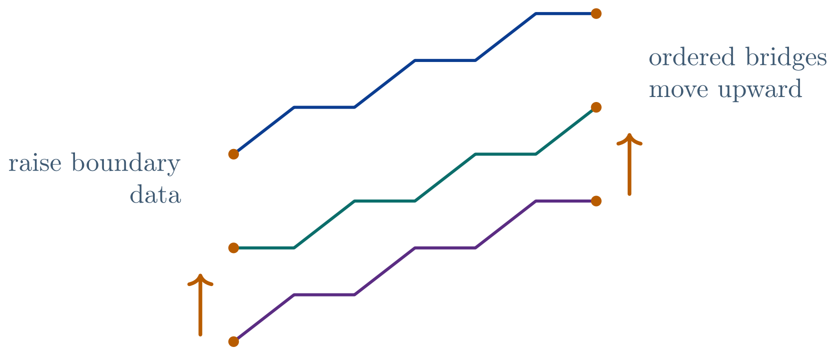

Figure 47

On a finite interval, attach an independent rate-one clock to every interior time. When the clock at time j rings, resample $\gamma(j)$ from the values compatible with $\gamma(j-1)$ and $\gamma(j+1)$. The endpoints are held fixed. A bridge and one of its update locations are depicted in Figure 48. If the two neighbouring increments are different, then precisely two middle heights are possible.

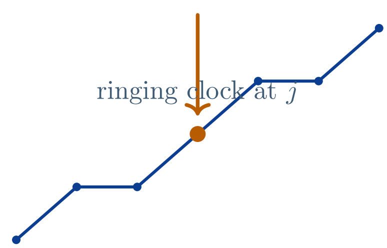

Figure 48

Resampling interchanges an upward–horizontal corner and a horizontal–upward corner, as shown in Figure 49. If the increments agree, only one middle height is possible and the update leaves the path unchanged. The resulting continuous-time Markov chain is irreducible on the finite set of bridges with the prescribed endpoints. Its transition rates are symmetric, so the uniform bridge measure is reversible and hence stationary. The unique stationary distribution is therefore uniform, and the chain converges to it from every initial bridge.

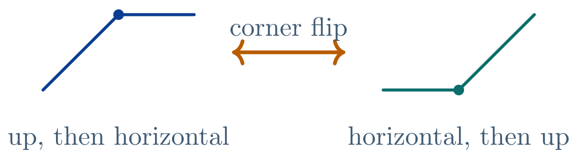

FIGURE 49

For two bridge systems with ordered boundary data, use the same clocks and the same uniform random variable at each update. Equivalently, update the two middle heights together and reject any proposed value that is not admissible for the corresponding bridge. This coupling is illustrated in Figure 50. The common update preserves pointwise order. Letting both chains converge to

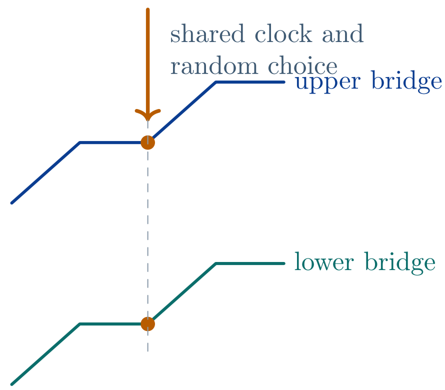

FIGURE 50

stationarity gives an ordered coupling of their uniform bridge laws. The same argument applies to an ordered collection of bridges: update one space-time site at a time and regard the adjacent lines as additional boundaries. This proves the ensemble assertion in Proposition 6.7.

The Two-Line Bernoulli Gibbs Law

We finish the finite-volume calculation underlying Theorem 6.6. Let Y_1, Y_2 be independent Bernoulli random walks with parameters $p_1 \geq p_2$, both started from zero, and set $X = WY$. Then X is the unique two-line Gibbs ensemble with asymptotic slopes p_1, p_2.

For a fixed integer t, the augmented transform

$$\widetilde{W}_t(Y_1, Y_2) = \big((WY)_1, (WY)_2, Y_2(t)\big)$$

is a bijection by Proposition 6.2. The extra endpoint coordinate is essential: the ordered output alone does not determine the input. The finite correspondence is depicted in Figure 51. Assume first that $p_1, p_2 \in (0, 1)$, and write $q_i = 1 - p_i$. A particular Bernoulli path γ of length t, started

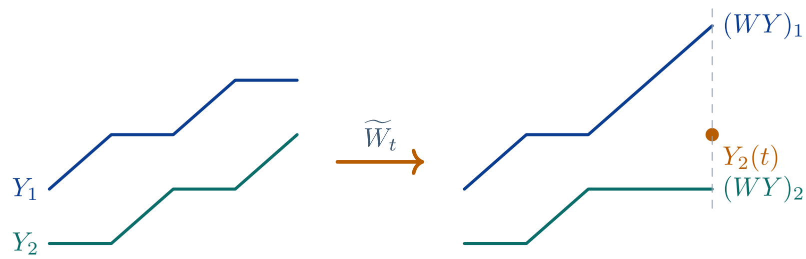

FIGURE 51

from zero, has probability

$$\mathbb{P}(Y_i = \gamma) = p_i^{\gamma(t)} q_i^{t-\gamma(t)} = q_i^t \left(\frac{p_i}{q_i}\right)^{\gamma(t)}.$$

Thus its probability depends on the path only through its endpoint, as illustrated in Figure 52. For an input pair, the corresponding product weight is

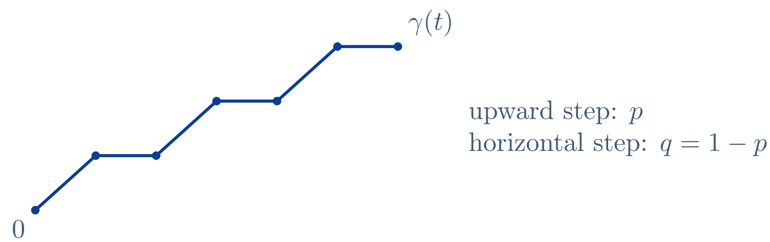

FIGURE 52

$$q_1^t q_2^t \left(\frac{p_1}{q_1}\right)^{Y_1(t)} \left(\frac{p_2}{q_2}\right)^{Y_2(t)}.$$

If the augmented output is $((g_1, g_2), z)$, then preservation of the sum gives

$$Y_2(t) = z \qquad \text{and} \qquad Y_1(t) = g_1(t) + g_2(t) - z.$$

Consequently, after conditioning on $g_1(t), g_2(t)$, and z, every admissible ordered output pair has the same weight. Summing over z changes the common weight by a factor that still depends only on the transformed endpoints. The conditional law of the paths between those endpoints is therefore uniform, which is precisely the Gibbs property in Definition 6.4.

The cases in which one of the parameters equals 0 or 1 follow by the same counting argument after the corresponding Bernoulli path is treated as deterministic.

Finally, the transform preserves $Y_1 + Y_2$. When $p_1 > p_2$, the process $Y_2 - Y_1$ has negative drift, so its running maximum is almost surely eventually constant. When $p_1 = p_2$, its running maximum is sublinear by the strong law of large numbers. Hence

$$\lim_{t\to\infty} \frac{X_1(t)}{t} = p_1 \qquad \text{and} \qquad \lim_{t\to\infty} \frac{X_2(t)}{t} = p_2$$

almost surely. Existence follows from this construction, and uniqueness follows from the monotone bridge coupling in Proposition 6.7. This completes the two-line proof of Theorem 6.6.

Lecture 9 — Friday, February 12, 2021

Tracy–Widom Distributions

For every $\beta > 0$, let $\mathcal{T}_\beta$ denote a random variable with the Tracy–Widom distribution of parameter β.

Let $q \in (0,1)$, and place independent geometric weights $w(i,j)$ on $\mathbb{Z}^2_{\geq 1}$, with

$$\mathbb{P}\big(w(i,j) = m\big) = (1-q)q^m, \qquad m \in \mathbb{Z}_{\geq 0}.$$

The point-to-point last-passage time across an N-by-N square is

$$G_N := \max_{\pi:(1,1)\nearrow(N,N)} \sum_{(i,j)\in\pi} w(i,j),$$

where π ranges over up-right lattice paths. The model and a maximizing path are illustrated in Figure 53.

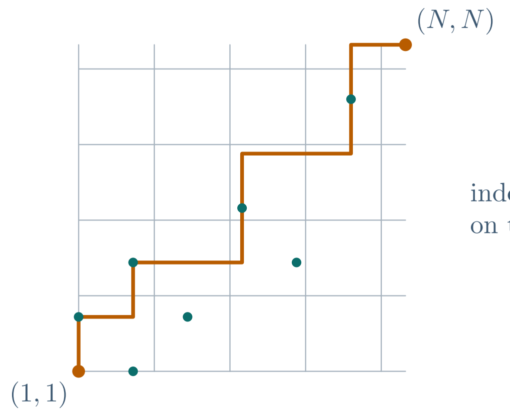

Figure 53

Theorem 9.1 There are positive constants $\gamma(q)$ and $\rho(q)$ such that

$$\frac{G_N}{N} \longrightarrow 2\gamma(q)$$

almost surely, and

$$\frac{G_N - 2\gamma(q)N}{\rho(q)N^{1/3}} \xrightarrow[N\to\infty]{\text{law}} \mathcal{T}_2.$$

The first limit is a law of large numbers. The second is the Kardar–Parisi–Zhang analogue of the central limit theorem, but with $N^{1/3}$ fluctuations and a non-Gaussian limit. For comparison, if $S_N = X_1 + \cdots + X_N$ is a sum of independent and identically distributed variables with mean μ and variance σ^2, then

$$\frac{S_N}{N} \longrightarrow \mu \qquad \text{and} \qquad \frac{S_N - N\mu}{\sqrt{N}} \xrightarrow[N\to\infty]{\text{law}} \mathcal{N}(0, \sigma^2).$$

Changing the symmetry of the last-passage model changes the limiting Tracy–Widom parameter. The three fundamental geometries are shown in Figure 54: the ordinary point-to-point model has parameter 2, the point-to-line model has parameter 1, and reflection symmetry across a diagonal produces parameter 4, subject to a conventional deterministic rescaling. For example, in a model symmetric across an antidiagonal $\mathcal{D}_N$, every path is obtained by reflecting its first half. Its passage time can therefore be written as

$$G_N^{\text{sym}} = 2 \max_{x\in\mathcal{D}_N} \max_{\pi:(1,1)\nearrow x} \sum_{(i,j)\in\pi} w(i,j).$$

Coupling the ordinary and symmetric models with the same weights yields a finite-N comparison,

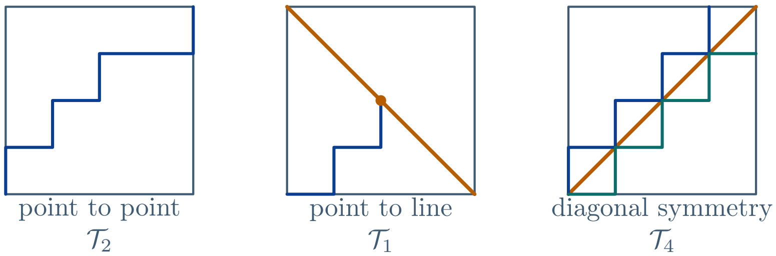

FIGURE 54

as represented in Figure 55. After centring, rescaling, and passing to the limit, one obtains the classical comparison

$$\mathcal{T}_2 \succeq_{\mathrm{st}} 2^{2/3}\mathcal{T}_4.$$

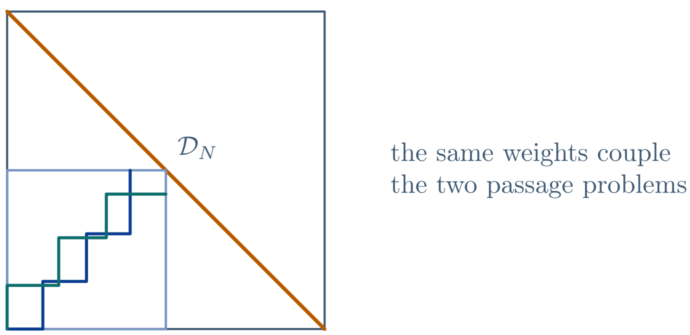

FIGURE 55

Random Matrices and the Stochastic Airy Operator

The same distributions arise at the spectral edge of Gaussian random matrices. Let $\mathbf{G}$ be an n-by-n real matrix with independent $\mathcal{N}(0,1)$ entries, and set

$$\mathbf{M}_n = \frac{\mathbf{G} + \mathbf{G}^{\mathsf{T}}}{\sqrt{2}}.$$

Then $\mathbf{M}_n$ belongs to the **Gaussian orthogonal ensemble** in the normalization for which its largest eigenvalue is asymptotic to $2\sqrt{n}$. Tracy and Widom proved that

$$n^{1/6}\big(\lambda_{\max}(\mathbf{M}_n) - 2\sqrt{n}\big) \xrightarrow[n\to\infty]{\text{law}} \mathcal{T}_1.$$

The analogous limits for the Gaussian unitary and **Gaussian symplectic ensemble** are $\mathcal{T}_2$ and $\mathcal{T}_4$, respectively.

The tridiagonal matrix model extends these three classical cases to every $\beta > 0$. Let

$$\mathbf{J}_{n,\beta} \coloneqq \frac{1}{\sqrt{\beta}} \begin{bmatrix} a_1 & b_1 & & & 0 \\ b_1 & a_2 & b_2 & & \\ & b_2 & \ddots & \ddots & \\ & & \ddots & a_{n-1} & b_{n-1} \\ 0 & & & b_{n-1} & a_n \end{bmatrix},$$

where the variables are independent and

$$a_i \sim \mathcal{N}(0,2) \qquad \text{and} \qquad b_j \sim \chi_{\beta(n-j)}.$$

For $\beta \in \{1,2,4\}$, this tridiagonal matrix has the same joint eigenvalue law as the corresponding Gaussian orthogonal, unitary, or symplectic ensemble. For arbitrary $\beta > 0$, it defines the **Gaussian $\boldsymbol{\beta}$-ensemble**.

Ramírez, Rider, and Virág showed that the edge-scaled matrices

$$\mathbf{H}_{n,\beta} \coloneqq n^{1/6}\big(2\sqrt{n}\,\mathbf{I} - \mathbf{J}_{n,\beta}\big)$$

converge, in an appropriate operator sense, to the **stochastic Airy operator**

$$\mathcal{H}_\beta \coloneqq -\partial_x^2 + x + \frac{2}{\sqrt{\beta}} B'_x$$

on $L^2(\mathbb{R}_{\geq 0})$, with a Dirichlet boundary condition at the origin. Here B' denotes white noise. For a sufficiently regular test function f satisfying $f(0) = 0$, its quadratic form is

$$\langle \mathcal{H}_\beta f, f\rangle = \int_0^\infty \big((f'(x))^2 + x f(x)^2\big)\,\mathrm{d}x + \frac{2}{\sqrt{\beta}} \int_0^\infty f(x)^2\,\mathrm{d}B(x).$$

Let

$$\Lambda_0^\beta < \Lambda_1^\beta < \Lambda_2^\beta < \cdots$$

be the eigenvalues of $\mathcal{H}_\beta$, listed in increasing order, and let $\lambda_{\beta,1}^{(n)} \geq \lambda_{\beta,2}^{(n)} \geq \cdots$ be the eigenvalues of $\mathbf{J}_{n,\beta}$. For each fixed k,

$$\Big(n^{1/6}\big(2\sqrt{n} - \lambda_{\beta,\ell}^{(n)}\big)\Big)_{\ell=1}^{k} \xrightarrow[n\to\infty]{\text{law}} \Big(\Lambda_{\ell-1}^\beta\Big)_{\ell=1}^{k}.$$

This gives the general definition

$$\mathcal{T}_\beta \overset{\text{law}}{=} -\Lambda_0^\beta, \qquad \beta > 0.$$

Stochastic Comparison Across β

For random variables U and V, write $U \succeq_{\mathrm{st}} V$ when

$$\mathbb{P}(U > a) \geq \mathbb{P}(V > a)$$

for every $a \in \mathbb{R}$. The comparison suggested by the symmetric last-passage models extends to arbitrary positive parameters.

Theorem 9.2 Let $\beta' > \beta > 0$ and $\alpha > 0$. Then

$$\mathcal{T}_\beta \succeq_{\mathrm{st}} \alpha \mathcal{T}_{\beta'}$$

if and only if

$$\left(\frac{\beta'}{\beta}\right)^{1/3} \leq \alpha \leq \left(\frac{\beta'}{\beta}\right)^{2/3}.$$

Proof. *Forward implication.* The two tails of the Tracy–Widom law satisfy, as $a \to \infty$,

$$\mathbb{P}\big(\mathcal{T}_\beta > a\big) = \exp\left(-\frac{2}{3}\beta a^{3/2}\big(1 + o(1)\big)\right),$$
$$\mathbb{P}\big(\mathcal{T}_\beta < -a\big) = \exp\left(-\frac{1}{24}\beta a^{3}\big(1 + o(1)\big)\right).$$

Suppose that $\mathcal{T}_\beta \succeq_{\mathrm{st}} \alpha \mathcal{T}_{\beta'}$. Comparison of the right tails gives

$$\exp\left(-\frac{2}{3}\frac{\beta'}{\alpha^{3/2}} a^{3/2}\big(1 + o(1)\big)\right) \leq \exp\left(-\frac{2}{3}\beta a^{3/2}\big(1 + o(1)\big)\right),$$

and hence

$$\alpha \leq \left(\frac{\beta'}{\beta}\right)^{2/3}.$$

Comparison of the left tails gives

$$\mathbb{P}\big(\mathcal{T}_\beta < -a\big) \leq \mathbb{P}\big(\mathcal{T}_{\beta'} < -a/\alpha\big),$$

which similarly implies

$$\alpha \geq \left(\frac{\beta'}{\beta}\right)^{1/3}.$$

Reverse implication. We couple all stochastic Airy operators using the same Brownian motion and use the Loewner order: $A \geq B$ when $\langle (A-B)f, f\rangle \geq 0$ for every f in their common form domain. This order implies the corresponding ordering of all eigenvalues by the min–max principle.

Fix $s > 0$ and make the change of variables $y = sx$. Brownian scaling gives the identity in distribution

$$s^{-1}\mathcal{H}_\beta \overset{\text{law}}{=} -s\partial_y^2 + \frac{1}{s^2}y + \frac{2}{\sqrt{s\beta}}B'_y =: \mathcal{H}^s_\beta.$$

Moreover, the eigenvalues of $\mathcal{H}^s_\beta$ are $s^{-1}\Lambda^\beta_k$. Set

$$\gamma := \sqrt{\frac{\beta'}{s\beta}}.$$

The white noise terms cancel in the difference

$$\gamma\mathcal{H}_{\beta'} - \mathcal{H}^s_\beta = -\left(\sqrt{\frac{\beta'}{s\beta}} - s\right)\partial_y^2 + \left(\sqrt{\frac{\beta'}{s\beta}} - \frac{1}{s^2}\right)y.$$

Both coefficients are nonnegative precisely when

$$\left(\frac{\beta}{\beta'}\right)^{1/3} \leq s \leq \left(\frac{\beta'}{\beta}\right)^{1/3}.$$

For such s, the difference is nonnegative in the Loewner order. The min–max principle therefore gives

$$\gamma\Lambda^{\beta'}_k \geq \frac{1}{s}\Lambda^\beta_k$$

for every $k \geq 0$. With

$$\alpha = s\gamma = \sqrt{s}\sqrt{\frac{\beta'}{\beta}},$$

the permitted range of s is equivalent to

$$\left(\frac{\beta'}{\beta}\right)^{1/3} \leq \alpha \leq \left(\frac{\beta'}{\beta}\right)^{2/3}.$$

Taking $k = 0$, multiplying by -1, and using $\mathcal{T}_\beta \overset{\text{law}}{=} -\Lambda^\beta_0$ proves the required stochastic domination. $\square$

Lecture 10 — Wednesday, February 24, 2021

Multiline Last-Passage Geometry

Let $f = (f_1, \ldots, f_n)$ be continuous functions on $\mathbb{R}_{\geq 0}$, all vanishing at the origin. A directed path from (x, n) to $(y, 1)$, where $0 \leq x \leq y$, is specified by jump times

$$x = t_n \leq t_{n-1} \leq \cdots \leq t_1 \leq t_0 = y.$$

The path follows line i from t_i to t_{i-1}, as shown in Figure 56. Its weight is the sum of the corresponding increments of the functions. The resulting last-passage value is

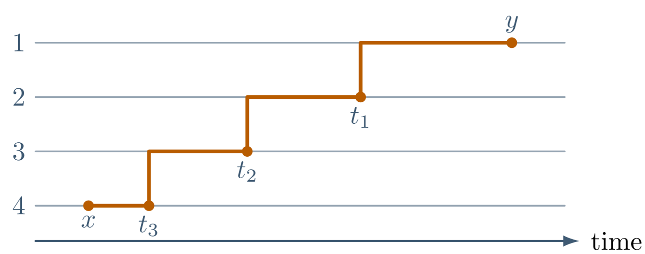

Figure 56

$$d_f\big((x, n), (y, 1)\big) := \sup_{x = t_n \leq \cdots \leq t_0 = y} \sum_{i=1}^{n} \big(f_i(t_{i-1}) - f_i(t_i)\big).$$

For $n = 2$, the two-line Pitman transform is characterized by

$$(Wf)_1(y) = d_f\big((0, 2), (y, 1)\big) \qquad \text{and} \qquad (Wf)_1(y) + (Wf)_2(y) = f_1(y) + f_2(y).$$

For $k \geq 1$, we may similarly maximize the total weight of a **disjoint k-tuple** of directed paths. The paths may share their entrance and exit points, but are otherwise disjoint. We write $d_f^{(k)}(\vec{p}, \vec{q})$ for the supremum of the total weight over all such k-tuples with entrance tuple $\vec{p}$ and exit tuple $\vec{q}$. A two-path configuration with entrance positions $x_1 \leq x_2$ and exit positions $y_1 \leq y_2$ is depicted in Figure 57. The two-line transform is idempotent and preserves last-passage values by Proposition 6.1. The next construction applies this local operation successively to every adjacent pair of lines.

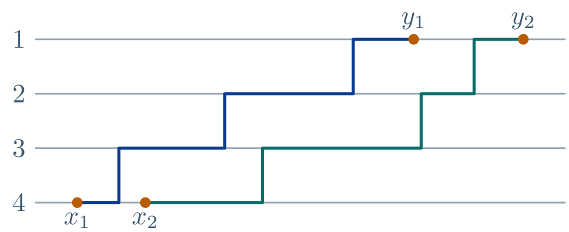

FIGURE 57

Functional Robinson–Schensted–Knuth Correspondence

For linear functions, the Pitman transform acts exactly like an adjacent comparison in bubble sort. If $f_i(x) = s_i x$ and $f_{i+1}(x) = s_{i+1}x$, then it places the larger slope above the smaller one. The two possible cases are shown in Figure 58. For $i \in \{1, \dots, n-1\}$, define the **adjacent transform**

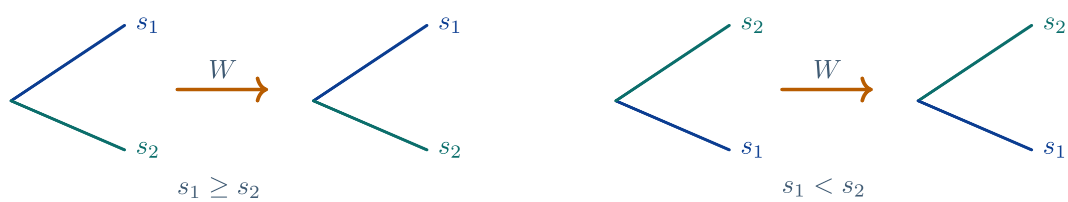

FIGURE 58

$$\sigma_i f := \big(f_1, \dots, f_{i-1}, (W(f_i, f_{i+1}))_1, (W(f_i, f_{i+1}))_2, f_{i+2}, \dots, f_n\big).$$

Set τ_n equal to the identity and, for $m < n$, let

$$\tau_m := \sigma_m \sigma_{m+1} \cdots \sigma_{n-1}.$$

Composition is read from right to left. The **functional Robinson–Schensted–Knuth transform** is

$$Wf := \tau_{n-1}\tau_{n-2} \cdots \tau_1 f.$$

This is a composition of

$$1 + 2 + \cdots + (n-1) = \binom{n}{2}$$

two-line Pitman transforms. It is also called the **melon** or the **functional Robinson–Schensted–Knuth image** of f. The reduced word applies adjacent comparisons in the same pattern as bubble sort, as illustrated in Figure 59.

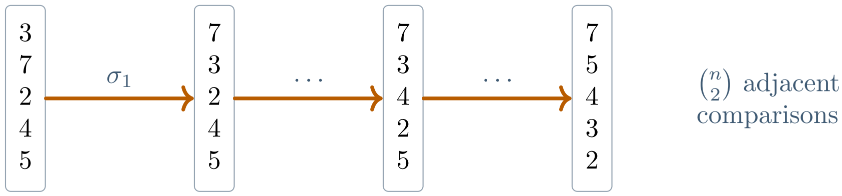

Figure 59

Theorem 10.1 The functional Robinson–Schensted–Knuth transform has the following properties:

1. It preserves outer-line last-passage values:
$$d_f\big((x,n),(y,1)\big) = d_{Wf}\big((x,n),(y,1)\big).$$
2. Its output is ordered:
$$(Wf)_1 \geq (Wf)_2 \geq \cdots \geq (Wf)_n.$$
3. For $k \in \{1,\ldots,n\}$, the sum of the top k output lines equals the k-path last-passage value:
$$\sum_{i=1}^{k} (Wf)_i(y) = d_f\big((0,n)^k,(y,1)^k\big).$$

The identity in (3) is Greene's theorem. Its proof requires path-swapping arguments. We first prove (1) and (2).

Proof of (1). It is enough to show that every adjacent transform σ_i preserves the last-passage value. A path from (x,n) to $(y,1)$ enters the two-line block $\{i,i+1\}$ at some time w and leaves it at some time z, as shown in Figure 60. With the empty boundary pieces omitted when necessary,

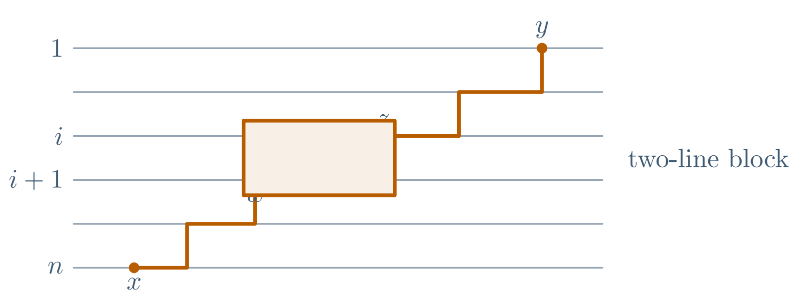

Figure 60

decomposition at the entry and exit times gives

$$d_f\big((x,n),(y,1)\big) = \sup_{x\le w\le z\le y}\Big[d_f\big((x,n),(w,i+2)\big) + d_f\big((w,i+1),(z,i)\big) + d_f\big((z,i-1),(y,1)\big)\Big].$$

Only the middle term uses lines i and $i+1$. By the two-line isometry in (2), it is unchanged when f is replaced by $\sigma_i f$; the other two terms do not use the transformed lines. Therefore

$$d_f\big((x,n),(y,1)\big) = d_{\sigma_i f}\big((x,n),(y,1)\big).$$

Applying this identity through the defining composition of W proves (1). □

To prove the ordering property, it is helpful to reinterpret the partial transforms τ_m as a traffic process. The lowest path moves freely. Each successive path is reflected upward by the paths below it, like a car that must queue behind the traffic ahead. This interpretation is depicted in Figure 61.

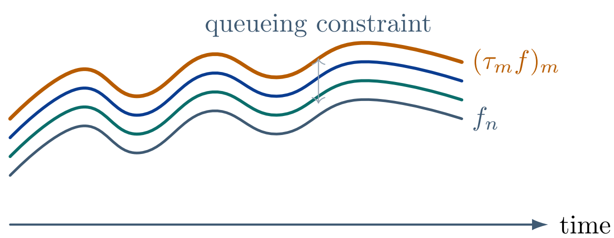

Figure 61

Lemma 10.2 For $m \in \{1,\dots,n\}$ and $y \ge 0$,

$$d_f\big((0,n),(y,m)\big) = (\tau_m f)_m(y).$$

Proof. We argue by induction on $n-m$. When $m=n$, the path remains on the bottom line, and both sides equal $f_n(y)$. Suppose now that $m<n$. Decomposing according to the time z at which the path enters line m and applying the induction hypothesis gives

$$\begin{aligned} d_f\big((0,n),(y,m)\big) &= \sup_{0\le z\le y}\big[d_f\big((0,n),(z,m+1)\big) + f_m(y) - f_m(z)\big] \\ &= \sup_{0\le z\le y}\,[(\tau_{m+1}f)_{m+1}(z) + f_m(y) - f_m(z)] \end{aligned}$$

$$
\begin{aligned}
&= (\sigma_m \tau_{m+1} f)_m(y) \\
&= (\tau_m f)_m(y).
\end{aligned}
$$

□

Lemma 10.3 If $m < n$, then
$$\tau_m^2 = \tau_{m+1}\tau_m.$$

Proof. Fix f, and set $g = \tau_{m+1}\tau_m f$. Since $\tau_m = \sigma_m \tau_{m+1}$, it suffices to prove $\sigma_m g = g$. By Lemma 10.2 and the local isometry in (2),

$$(\tau_m^2 f)_m(y) = d_{\tau_m f}\big((0,n),(y,m)\big) = d_f\big((0,n),(y,m)\big) = (\tau_m f)_m(y).$$

The operator τ_{m+1} does not change coordinate m, so

$$g_m = (\tau_m f)_m.$$

On the other hand, $\tau_m^2 f = \sigma_m \tau_{m+1}\tau_m f = \sigma_m g$. Hence $(\sigma_m g)_m = g_m$. The two-line transform preserves the sum of coordinates m and $m+1$, so the other coordinate is unchanged as well. Thus $\sigma_m g = g$, which proves the identity. □

Proof of (2). For any line ensemble g,
$$\sigma_i g = g$$
if and only if $g_i \geq g_{i+1}$. It is therefore enough to prove $\sigma_m W = W$ for every $m \in \{1, \dots, n-1\}$.

When $m = n-1$, the leftmost factor in W is $\tau_{n-1} = \sigma_{n-1}$. Idempotence of the two-line transform gives $\sigma_{n-1} W = W$.

Now let $m < n-1$. Adjacent transforms σ_i and σ_j commute when $|i-j| \geq 2$. We may therefore move σ_m to the right until it reaches $\tau_{m+1}\tau_m$. Since $\tau_m = \sigma_m \tau_{m+1}$, Lemma 10.3 gives

$$\sigma_m \tau_{m+1}\tau_m = \tau_m^2 = \tau_{m+1}\tau_m.$$

Consequently,

$$\sigma_m W = \tau_{n-1}\cdots\tau_{m+2}\sigma_m\tau_{m+1}\tau_m\cdots\tau_1 = \tau_{n-1}\cdots\tau_{m+2}\tau_{m+1}\tau_m\cdots\tau_1 = W.$$

□

The remaining structural questions are the proof of Greene's identity (3) and the invertibility of the multiline transform. As in the two-line case, invertibility requires retaining appropriate endpoint data.

Lecture 11 — Friday, February 26, 2021

Determinantal Point Processes and Nonintersecting Paths

Determinantal point processes encode many systems of nonintersecting paths and random matrix eigenvalues through their correlation functions. The development is based in part on Hough, Krishnapur, Peres, and Virág's *Zeros of Gaussian Analytic Functions and Determinantal Point Processes* and Tao's exposition of determinantal processes.

Let Λ be a locally compact Polish space. Thus Λ is homeomorphic to a complete separable metric space, and every point has a compact neighbourhood. A **Radon measure** on Λ is a Borel measure that is finite on compact sets and satisfies the usual inner and outer regularity properties.

Definition 11.1 A **point process** on Λ is a random nonnegative, integer-valued Radon measure $\mathcal{X}$. It is **simple** if

$$\mathbb{P}\big(\mathcal{X}(\{x\}) \leq 1 \text{ for every } x \in \Lambda\big) = 1.$$

Every realization can be written as

$$\mathcal{X} = \sum_{i \in I} \delta_{x_i},$$

where I is at most countable and the collection $\{x_i : i \in I\}$ is locally finite, with repetitions allowed when the process is not simple. For a Borel set $D \subseteq \Lambda$,

$$\mathcal{X}(D) = \#\{i \in I : x_i \in D\}$$

counts the points in D, including multiplicities.

On a countable state space, independently retaining the site x with probability p_x produces a simple **Bernoulli point process**. By contrast, placing an independent Poisson number of particles at each site generally produces multiplicities. On a continuous space with an atomless intensity measure, a **Poisson point process** is simple almost surely.

Fix a Radon reference measure μ on Λ. Typical choices are counting measure on a countable space and Lebesgue measure on an open subset of $\mathbb{R}^d$.

Definition 11.2 The **joint intensities**, or **correlation functions**, of a simple point process $\mathcal{X}$ are measurable functions $\rho_k \colon \Lambda^k \to [0,\infty)$ such that, for pairwise disjoint Borel sets $D_1, \dots, D_k$,

$$\mathbb{E}\left[\prod_{i=1}^{k} \mathcal{X}(D_i)\right] = \int_{D_1\times\cdots\times D_k} \rho_k(x_1,\dots,x_k) \prod_{i=1}^{k} \mu(\mathrm{d}x_i). \tag{11.1}$$

We set $\rho_k(x_1,\dots,x_k) = 0$ whenever two coordinates coincide.

When Λ is countable and μ is counting measure, (11.1) gives

$$\rho_k(x_1,\dots,x_k) = \mathbb{P}(x_1,\dots,x_k \in \mathcal{X})$$

for distinct sites $x_1,\dots,x_k$. Thus ρ_k is the continuous analogue of a joint inclusion probability. For a Poisson point process on $\mathbb{R}^d$ with constant intensity λ,

$$\rho_k(x_1,\dots,x_k) = \lambda^k.$$

Subject to standard local growth conditions, the full collection of joint intensities determines the law of the process.

The determinantal form of these intensities first appears naturally in a fermionic model. Suppose that a single electron has a normalized **wavefunction** ψ, so that $|\psi(x)|^2$ is its spatial density and

$$\int_\Lambda |\psi(x)|^2 \,\mu(\mathrm{d}x) = 1.$$

For n distinguishable particles in one-particle states $\psi_1,\dots,\psi_n$, independence would suggest the product $\prod_{i=1}^n \psi_i(x_i)$. This model misses two physical effects. First, Coulomb's law gives a repulsive force of magnitude proportional to q_1q_2/r^2 between equal charges, so electron positions are not genuinely independent. Second, electrons are indistinguishable **fermions**, so exchanging two particle labels must reverse the sign of the **many-particle wavefunction**. The following **antisymmetrization** addresses the second issue and creates an additional exclusion effect even before Coulomb interactions are included. For two particles it is

$$\Psi(x_1,x_2) = \frac{\psi_1(x_1)\psi_2(x_2) - \psi_2(x_1)\psi_1(x_2)}{\sqrt{2}}.$$

For n particles it becomes the **Slater determinant**

$$\Psi(\boldsymbol{x}) = \frac{1}{\sqrt{n!}} \sum_{\pi \in S_n} \operatorname{sgn}(\pi) \prod_{i=1}^{n} \psi_{\pi(i)}(x_i) = \frac{1}{\sqrt{n!}} \det\left[\psi_j(x_i)\right]_{i,j=1}^{n}. \tag{11.2}$$

The factor $1/\sqrt{n!}$ gives a normalized many-particle wavefunction when the one-particle states are orthonormal; without orthonormality, the displayed formula should be read only as their antisymmetrization. The sign changes because composing a permutation with a transposition reverses its sign. In particular, Ψ vanishes whenever two particle locations coincide. This exclusion mechanism is illustrated in Figure 62.

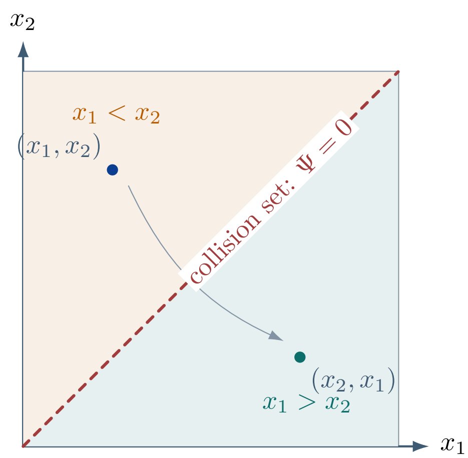

Figure 62

Assume now that $\psi_1, \dots, \psi_n$ are orthonormal in $L^2(\Lambda, \mu)$, and define

$$K(x, y) \coloneqq \sum_{j=1}^{n} \psi_j(x)\overline{\psi_j(y)}. \tag{11.3}$$

If $\mathbf{V} = [\psi_j(x_i)]_{i,j=1}^{n}$, then

$$|\Psi(\boldsymbol{x})|^2 = \frac{1}{n!} \det(\mathbf{V}) \det(\mathbf{V}^*) = \frac{1}{n!} \det(\mathbf{V}\mathbf{V}^*) = \frac{1}{n!} \det\left[K(x_i, x_j)\right]_{i,j=1}^{n}.$$

The determinant is therefore the joint density created by fermionic antisymmetry.

Definition 11.3 Let (Λ, μ) be a locally compact Polish space equipped with a Radon measure. A simple point process $\mathcal{X}$ is a **determinantal point process** with **correlation**

kernel $K\colon \Lambda^2 \to \mathbb{C}$, relative to μ, if its joint intensities satisfy

$$\rho_k(x_1, \dots, x_k) = \det\left[K(x_i, x_j)\right]_{i,j=1}^k \tag{11.4}$$

for every $k \geq 1$, almost everywhere with respect to $\mu^{\otimes k}$.

Several qualifications in Definition 11.3 are important.

1. A kernel of an integral operator is initially defined only up to a $\mu^{\otimes 2}$-null set. On an atomless space, this equivalence class does not determine its values on the diagonal. We must therefore choose a suitable measurable version; the correlation functions themselves are likewise defined only almost everywhere.
2. Local finiteness of the process requires local integrability. The locally trace-class condition introduced below gives the standard operator-theoretic formulation of this requirement.
3. Equation (11.4) requires every principal determinant to be nonnegative. Hermitian positive semidefiniteness of the kernel is the natural sufficient structure, but pointwise nonnegativity of $K(x, y)$ is neither required nor generally true.
4. The pair (K, μ) is not unique. If $h\colon \Lambda \to \mathbb{C} \setminus \{0\}$ is measurable, then

   $$K_h(x, y) = h(x)\overline{h(y)}K(x, y) \qquad \text{and} \qquad \mu_h(\mathrm{d}x) = |h(x)|^{-2}\mu(\mathrm{d}x)$$

   gives the same joint intensities as measures, because the determinant gains the factor $\prod_i |h(x_i)|^2$, which is cancelled by the change in reference measure.

Definition 11.4 A measurable kernel K defines the **integral operator** K on $L^2(\Lambda, \mu)$ by

$$(\mathsf{K}f)(x) = \int_\Lambda K(x, y) f(y)\, \mu(\mathrm{d}y),$$

whenever the integral is defined. The operator is **locally trace class** if $\mathbb{1}_D \mathsf{K} \mathbb{1}_D$ is trace class for every compact set $D \subseteq \Lambda$. A compact operator is trace class when the sum of its singular values is finite.

The following criterion answers both the existence and uniqueness questions for Hermitian kernels.

Theorem 11.5 Let K be a locally square-integrable Hermitian kernel whose associated operator K on $L^2(\Lambda, \mu)$ is self-adjoint and locally trace class. There exists a unique determinantal point process with kernel K if and only if

$$0 \leq \mathsf{K} \leq \mathsf{I}.$$

Equivalently, the spectrum of K is contained in $[0, 1]$.

The finite-rank kernel in (11.3) is the **projection kernel** of the orthogonal projection onto $\operatorname{span}\{\psi_1, \ldots, \psi_n\}$. Its spectrum consists only of zeros and ones, so Theorem 11.5 applies. The resulting process has exactly n points almost surely.

Determinantal processes also arise from nonintersection. We first record the ordered one-dimensional form of the Karlin–McGregor formula; a more general signed version and its proof appear in Theorem 12.1.

Theorem 11.6 (Karlin–McGregor theorem) Let $X_1, \ldots, X_n$ be independent copies of a possibly time-inhomogeneous nearest-neighbour Markov chain on $\mathbb{Z}$, with transition probabilities

$$\mathbb{P}_{a,b}(s, t) = \mathbb{P}(X(t) = b \mid X(s) = a).$$

Start the chains from the ordered configuration $\boldsymbol{i} = (i_1, \ldots, i_n)$, where $i_1 < \cdots < i_n$, and fix an ordered terminal configuration $\boldsymbol{j} = (j_1, \ldots, j_n)$, where $j_1 < \cdots < j_n$. After linearly interpolating the walks between successive times,

$$\mathbb{P}\left(\begin{array}{c} X_p(t) = j_p \text{ for } 1 \leq p \leq n, \\ X_1(s) < \cdots < X_n(s) \text{ for } 0 < s < t \end{array} \right) = \det\left[\mathbb{P}_{i_p, j_q}(0, t)\right]_{p,q=1}^n.$$

Consider now n independent, time-homogeneous, reversible nearest-neighbour chains that start from the ordered configuration $x_1 < \cdots < x_n$, return to the same configuration at time $2t$, and are conditioned not to intersect. The geometry is shown in Figure 63; the random ordered configuration at time t is denoted by $\boldsymbol{U} = (U_1, \ldots, U_n)$.

Corollary 11.7 Let $\mathbb{P}_t(a, b)$ be the transition probabilities of a time-homogeneous, reversible, nearest-neighbour Markov chain on $\mathbb{Z}$, and let π be a strictly positive reversible measure. Define the matrix

$$\mathbf{A} = \left[\frac{\mathbb{P}_{2t}(x_p, x_q)}{\pi(x_q)}\right]_{p,q=1}^n \tag{11.5}$$

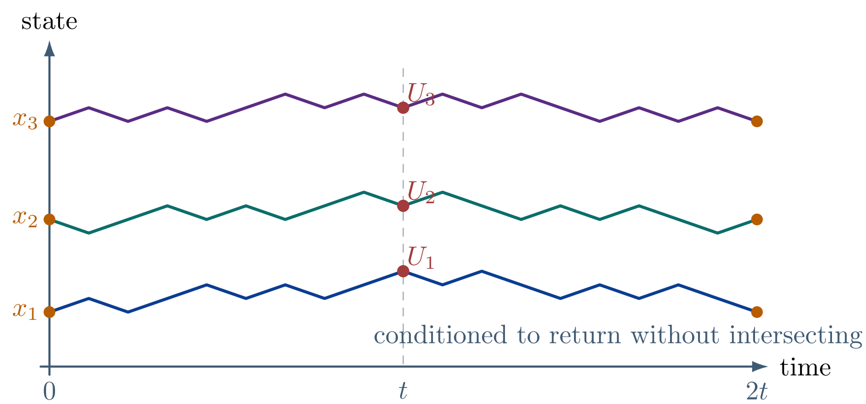

Figure 63

or, explicitly,

$$\mathbf{A} = \begin{bmatrix} \dfrac{\mathbb{P}_{2t}(x_1, x_1)}{\pi(x_1)} & \cdots & \dfrac{\mathbb{P}_{2t}(x_1, x_n)}{\pi(x_n)} \\ \vdots & \ddots & \vdots \\ \dfrac{\mathbb{P}_{2t}(x_n, x_1)}{\pi(x_1)} & \cdots & \dfrac{\mathbb{P}_{2t}(x_n, x_n)}{\pi(x_n)} \end{bmatrix},$$

and the functions

$$\psi_j(u) = \sum_{k=1}^{n} \left[\mathbf{A}^{-1/2}\right]_{j,k} \frac{\mathbb{P}_t(x_k, u)}{\pi(u)}, \qquad 1 \le j \le n. \tag{11.6}$$

If the nonintersection conditioning event has positive probability, then the midpoint configuration in Figure 63 is a determinantal point process relative to π, with correlation kernel

$$K(u, v) = \sum_{j=1}^{n} \psi_j(u)\psi_j(v). \tag{11.7}$$

Proof. Let $\mathsf{NI}(s_1, s_2)$ denote the event that the paths do not intersect between times s_1 and s_2. For an ordered midpoint configuration $\boldsymbol{u} = (u_1, \ldots, u_n)$, Bayes' rule expresses the desired conditional probability as the probability of the paths passing through $\boldsymbol{u}$, divided by the probability of the conditioned endpoint event. The Markov property splits the numerator at time t, time homogeneity translates the second half back to an interval of length t, and Theorem 11.6 gives

$$\mathbb{P}\big(\boldsymbol{U} = \boldsymbol{u} \mid \boldsymbol{X}(0) = \boldsymbol{x},\ \boldsymbol{X}(2t) = \boldsymbol{x},\ \mathsf{NI}(0, 2t)\big) = \frac{\det[\mathbb{P}_t(x_p, u_q)]_{p,q=1}^{n} \det[\mathbb{P}_t(u_p, x_q)]_{p,q=1}^{n}}{\det[\mathbb{P}_{2t}(x_p, x_q)]_{p,q=1}^{n}}.$$

Introduce

$$\mathbf{B}(\boldsymbol{u}) = \left[\frac{\mathbb{P}_t(x_p, u_q)}{\pi(u_q)}\right]_{p,q=1}^n .$$

Reversibility implies that $\mathbf{A}$ is symmetric and that

$$\det[\mathbb{P}_t(x_p, u_q)] = \det(\mathbf{B}(\boldsymbol{u})) \prod_{q=1}^n \pi(u_q),$$

$$\det[\mathbb{P}_t(u_p, x_q)] = \det\Big(\mathbf{B}(\boldsymbol{u})^\mathsf{T}\Big) \prod_{q=1}^n \pi(x_q),$$

$$\det[\mathbb{P}_{2t}(x_p, x_q)] = \det(\mathbf{A}) \prod_{q=1}^n \pi(x_q).$$

Consequently, the conditional probability is

$$\begin{aligned}\frac{\det(\mathbf{B})\det\big(\mathbf{B}^\mathsf{T}\big)}{\det(\mathbf{A})} \prod_{q=1}^n \pi(u_q) &= \det\Big[(\mathbf{A}^{-1/2}\mathbf{B})^\mathsf{T}(\mathbf{A}^{-1/2}\mathbf{B})\Big] \prod_{q=1}^n \pi(u_q) \\ &= \det[K(u_p, u_q)]_{p,q=1}^n \prod_{q=1}^n \pi(u_q).\end{aligned}$$

Here $\mathbf{A}$ is the Gram matrix in $L^2(\mathbb{Z}, \pi)$ of the functions $u \mapsto \mathbb{P}_t(x_p, u)/\pi(u)$. It is positive definite whenever the conditioning event has positive probability, and (11.6) orthonormalizes these functions. Hence (11.7) is a rank-n orthogonal projection kernel.

The displayed conditional probability identifies the n-point intensity. Integrating out $n-k$ variables and using orthonormality yields

$$\frac{1}{(n-k)!} \sum_{u_{k+1},\ldots,u_n \in \mathbb{Z}} \det[K(u_p, u_q)]_{p,q=1}^n \prod_{q=k+1}^n \pi(u_q) = \det[K(u_p, u_q)]_{p,q=1}^k .$$

Thus all joint intensities have the determinantal form in (11.4) for $k \le n$. For $k > n$, the joint intensity vanishes because the process has exactly n points, while the determinant in (11.4) vanishes because K has rank n. $\square$

We finish by identifying the eigenvalues of a Gaussian unitary ensemble as a determinantal point process. Equivalently to Definition 7.6, we may take a matrix $\mathbf{G}$ with independent standard complex Gaussian entries, where "standard complex Gaussian" means $(X+\mathrm{i}Y)/\sqrt{2}$ for independent

$X, Y \sim \mathcal{N}(0,1)$, and set

$$\mathbf{M} = \frac{\mathbf{G} + \mathbf{G}^*}{\sqrt{2}}.$$

Proposition 11.8 Let $\mathbf{M} \sim \mathrm{GUE}(n)$, as in Definition 7.6, and list its eigenvalues in a uniformly chosen random order. Their joint density on $\mathbb{R}^n$ is

$$\frac{1}{(2\pi)^{n/2}\prod_{k=1}^{n} k!} \exp\left(-\frac{1}{2}\|\boldsymbol{\lambda}\|_2^2\right) \prod_{1\le i<j\le n} (\lambda_i - \lambda_j)^2. \tag{11.8}$$

Proof sketch. By Theorem 7.8, the ordered eigenvalues of Hermitian Brownian motion evolve as Dyson Brownian motion. For $\boldsymbol{\lambda}$ and $\boldsymbol{\nu}$ in the decreasing Weyl chamber, the **Johansson transition density** is

$$p_t(\boldsymbol{\lambda} \mid \boldsymbol{\nu}) = \frac{1}{(2\pi t)^{n/2}} \frac{\Delta_n(\boldsymbol{\lambda})}{\Delta_n(\boldsymbol{\nu})} \det\left[\exp\left(-\frac{(\lambda_i - \nu_j)^2}{2t}\right)\right]_{i,j=1}^{n}. \tag{11.9}$$

Set $t = 1$ and let $\boldsymbol{\nu} \to \mathbf{0}$ within the Weyl chamber. The determinant factors as

$$\det\left[e^{-(\lambda_i-\nu_j)^2/2}\right]_{i,j=1}^{n} = e^{-\|\boldsymbol{\lambda}\|_2^2/2} e^{-\|\boldsymbol{\nu}\|_2^2/2} \det\left[e^{\lambda_i \nu_j}\right]_{i,j=1}^{n}.$$

Expanding the last determinant in power series gives

$$\det\left[e^{\lambda_i \nu_j}\right]_{i,j=1}^{n} = \frac{\Delta_n(\boldsymbol{\lambda})\Delta_n(\boldsymbol{\nu})}{\prod_{k=0}^{n-1} k!} + o\big(\Delta_n(\boldsymbol{\nu})\big).$$

Substitution into (11.9) gives the density on the ordered chamber. A uniformly random ordering divides that density by $n!$, and $n!\prod_{k=0}^{n-1} k! = \prod_{k=1}^{n} k!$, which proves (11.8). ☐

Let H_k denote the monic probabilists' Hermite polynomial of degree k. For the standard Gaussian measure

$$\gamma(\mathrm{d}x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}\,\mathrm{d}x,$$

these polynomials satisfy

$$\int_{\mathbb{R}} H_k(x) H_\ell(x)\,\gamma(\mathrm{d}x) = k!\,\mathbb{1}_{\{k=\ell\}}.$$

Thus $\varphi_k = H_k/\sqrt{k!}$ form an orthonormal family in $L^2(\mathbb{R}, \gamma)$. They may be defined explicitly by

$$H_k(x) = (-1)^k e^{x^2/2} \frac{\mathrm{d}^k}{\mathrm{d}x^k} e^{-x^2/2}.$$

Theorem 11.9 The eigenvalues of $\mathbf{M} \sim \mathrm{GUE}(n)$ form a determinantal point process relative to γ, with the **Hermite kernel**

$$K_n(x, y) = \sum_{k=0}^{n-1} \varphi_k(x)\varphi_k(y) = \sum_{k=0}^{n-1} \frac{H_k(x)H_k(y)}{k!}. \tag{11.10}$$

Proof. For distinct $x_1, \ldots, x_n$, let $\mathbf{V} = [\varphi_{j-1}(x_i)]_{i,j=1}^n$. Then

$$[K_n(x_i, x_j)]_{i,j=1}^n = \mathbf{V}\mathbf{V}^\mathsf{T},$$

so its determinant is nonnegative and vanishes whenever two coordinates coincide. As a polynomial in $x_1, \ldots, x_n$, it has total degree at most $n(n-1)$. It is divisible by $\Delta_n(\boldsymbol{x})^2$, which has exactly the same degree, and is therefore a constant multiple of that polynomial. Since H_k is monic,

$$\det[K_n(x_i, x_j)]_{i,j=1}^n = \frac{\Delta_n(\boldsymbol{x})^2}{\prod_{k=0}^{n-1} k!}.$$

It follows that the density of a randomly ordered eigenvalue vector, relative to $\gamma^{\otimes n}$, is

$$\frac{1}{n!} \det[K_n(x_i, x_j)]_{i,j=1}^n,$$

which agrees with (11.8). Since K_n is the rank-n orthogonal projection onto $\operatorname{span}\{\varphi_0, \ldots, \varphi_{n-1}\}$, the same integration identity used in the proof of Corollary 11.7 gives every lower-order joint intensity. Hence the eigenvalue process is determinantal. $\square$

There is also a direct path construction behind this result. Apply the continuous-state Karlin–McGregor formula to n Brownian bridges over the interval $[0, 2t]$, with Gaussian heat kernel

$$p_t(x, y) = \frac{1}{\sqrt{2\pi t}} \exp\left(-\frac{(x-y)^2}{2t}\right).$$

Condition the bridges not to intersect and condition their initial and terminal configurations to equal $x_1 < \cdots < x_n$, as illustrated in Figure 64. Their midpoint density is a determinant of Gaussian heat kernels. As the endpoints coalesce, $\boldsymbol{x} \to \mathbf{0}$ within the Weyl chamber, the Vandermonde limit used in the proof of Proposition 11.8 shows that the midpoint configuration converges to the Gaussian unitary ensemble eigenvalue process, after the corresponding Brownian scaling.

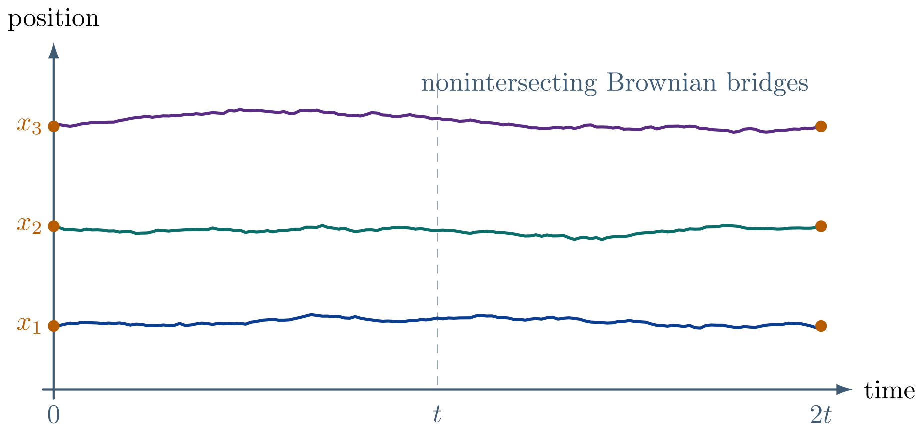

FIGURE 64

Lecture 12 — Wednesday, March 03, 2021

The Karlin–McGregor Formula

Let $\mathbb{P}_t(x, y)$ denote the transition probability of a Markov chain from x to y in time t. The determinant of the one-particle transition matrix has a path interpretation in terms of independent particles killed at their first collision.

Theorem 12.1 (Karlin–McGregor formula) Let $X_1, \ldots, X_n$ be independent copies of a Markov chain on a discrete state space, started from $\boldsymbol{x} = (x_1, \ldots, x_n)$, and let

$$T \coloneqq \inf\{s \geq 0 : X_i(s) = X_j(s) \text{ for some } i \neq j\}$$

be the first collision time. For $\boldsymbol{y} = (y_1, \ldots, y_n)$,

$$\det\left[\mathbb{P}_t(x_i, y_j)\right]_{i,j=1}^n = \sum_{\sigma \in S_n} \operatorname{sgn}(\sigma)\, \mathbb{P}_{\boldsymbol{x}}\big(\boldsymbol{X}(t) = \boldsymbol{y}_\sigma,\ T > t\big),$$

where $\boldsymbol{y}_\sigma = (y_{\sigma(1)}, \ldots, y_{\sigma(n)})$.

Proof. Consider the random matrix

$$\mathbf{A}_{ij} \coloneqq \mathbb{I}_{\{X_i(t) = y_j\}}.$$

The rows of $\mathbf{A}$ are independent. Since the determinant is multilinear in its rows,

$$\mathbb{E}_{\boldsymbol{x}}[\det \mathbf{A}] = \det\left[\mathbb{E}_{\boldsymbol{x}}[\mathbf{A}_{ij}]\right]_{i,j=1}^n = \det\left[\mathbb{P}_t(x_i, y_j)\right]_{i,j=1}^n.$$

We next show that the contribution from $\{T \leq t\}$ vanishes. On this event, choose the first collision time and, by a fixed tie-breaking rule, a pair $i < j$ that collides then. Swap the future trajectories of particles i and j after that time. The strong Markov property and the fact that the particles have identical transition laws make this a measure-preserving involution. It exchanges two rows of the terminal indicator matrix, so it changes the sign of its determinant. Pairing each configuration with its image gives

$$\mathbb{E}_{\boldsymbol{x}}\big[\det \mathbf{A}\, \mathbb{I}_{\{T \leq t\}}\big] = 0.$$

Consequently,

$$\begin{aligned}
\det\left[\mathbb{P}_t(x_i, y_j)\right]_{i,j=1}^n &= \mathbb{E}_{\boldsymbol{x}}\big[\det \mathbf{A}\, \mathbb{I}_{\{T > t\}}\big] \\
&= \mathbb{E}_{\boldsymbol{x}}\left[\sum_{\sigma \in S_n} \operatorname{sgn}(\sigma) \prod_{i=1}^n \mathbb{I}_{\{X_i(t) = y_{\sigma(i)}\}} \mathbb{I}_{\{T > t\}}\right] \\
&= \sum_{\sigma \in S_n} \operatorname{sgn}(\sigma)\, \mathbb{P}_{\boldsymbol{x}}\big(\boldsymbol{X}(t) = \boldsymbol{y}_\sigma,\ T > t\big).
\end{aligned}$$

$\square$

In many ordered state spaces, paths cannot exchange their order without colliding. If both $\boldsymbol{x}$ and

$\boldsymbol{y}$ are strictly ordered in the same direction, then only the identity permutation contributes, and

$$\det\left[\mathbb{P}_t(x_i, y_j)\right]_{i,j=1}^n = \mathbb{P}_{\boldsymbol{x}}\big(\boldsymbol{X}(t) = \boldsymbol{y},\ T > t\big).$$

Exercise 12.2 Identify other state spaces and transition rules for which noncollision prevents the particles from changing order.

Solution. The basic setting is any linearly ordered state space in which a sample path cannot pass from one side of another path to the other side without the two paths first taking the same value. Several important examples satisfy this condition.

Continuous real-valued paths have this property by the intermediate value theorem. Consequently, it applies to independent one-dimensional diffusions, including Brownian motions and Ornstein–Uhlenbeck processes. It also applies to continuous-time nearest-neighbour walks on $\mathbb{Z}$: jumps occur one at a time almost surely, and a particle must visit the intervening lattice site before it can pass another particle.

In discrete time, Bernoulli walks with increments in $\{0, 1\}$ are **skip-free**. If $X_i(t) > X_j(t)$, their gap can decrease by at most one during the next step, so it must become zero before it can become negative. The same conclusion holds for any one-dimensional transition rule whose coupled particle gaps cannot jump across zero. By contrast, simultaneously updated nearest-neighbour walks with increments in $\{-1, 1\}$ can exchange adjacent positions in one step; unless crossing edges are counted as collisions, that transition rule does not have the required property. $\square$

The Multiline Pitman Bijection

Fix $t > 0$, and regard the functional Robinson–Schensted–Knuth transform as a map on $\mathcal{F}_t^n$. In the two-line case, Proposition 6.2 shows that the ordered output pair together with the endpoint $f_2(t)$ uniquely determines the input.

For general n, write the reduced word defining the multiline transform as

$$W = \sigma_{i_N}\sigma_{i_{N-1}}\cdots\sigma_{i_1} \qquad \text{and} \qquad N = \binom{n}{2},$$

where σ_{i_1} acts first. Starting with $f^{(0)} = f$, define

$$f^{(r)} := \sigma_{i_r} f^{(r-1)} \qquad \text{and} \qquad z_r := f^{(r-1)}_{i_r+1}(t), \qquad r \in \{1, \dots, N\}.$$

Thus $f^{(N)} = W_t f$, and the vector

$$\overline{W}_t f := \boldsymbol{z} = (z_1, \dots, z_N)$$

records the lower input endpoint at every two-line step.

To describe the image, start with an ordered output $g^{(N)}$ and read the steps backward. At stage r, the two-line inverse exists precisely when

$$g^{(r)}_{i_r+1}(t) \le z_r \le g^{(r)}_{i_r}(t).$$

When this inequality holds, apply the inverse from Proposition 6.2 to coordinates i_r, i_r+1, using z_r as the retained endpoint, and call the result $g^{(r-1)}$. Let $\mathcal{A}_t$ be the set of pairs $(g^{(N)}, \boldsymbol{z})$ for which this backward recursion is admissible at every stage. These recursive inequalities are the multiline analogue of

$$(W_t f)_2(t) \le f_2(t) \le (W_t f)_1(t)$$

from the two-line case.

Proposition 12.3 The map

$$\widetilde{W}_t \colon \mathcal{F}^n_t \longrightarrow \mathcal{A}_t \qquad \text{and} \qquad \widetilde{W}_t f := (W_t f, \overline{W}_t f),$$

is a bijection.

Proof. At every forward step, the corresponding augmented two-line transform is bijective by Proposition 6.2. Composing these finite-step bijections proves injectivity. Conversely, if $(g^{(N)}, \boldsymbol{z}) \in \mathcal{A}_t$, the defining backward recursion applies the unique two-line inverse at each stage and produces a unique $g^{(0)} \in \mathcal{F}^n_t$. This path tuple maps to $(g^{(N)}, \boldsymbol{z})$, proving surjectivity. □

Greene's Theorem

We now prove the k-path identity (3). A typical disjoint path ensemble is shown in Figure 65.

Proof of (3). We first establish a local multipath isometry. Within a two-line strip, a disjoint

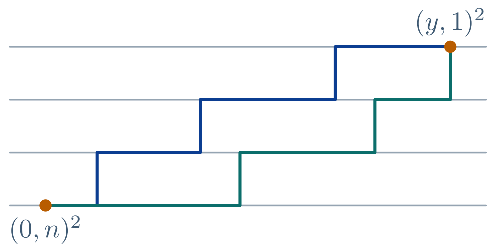

FIGURE 65

path ensemble decomposes into the three types shown in Figure 66: a single path, a pair that occupies both lines, or an inadmissible attempt to place three paths through the strip. For a single

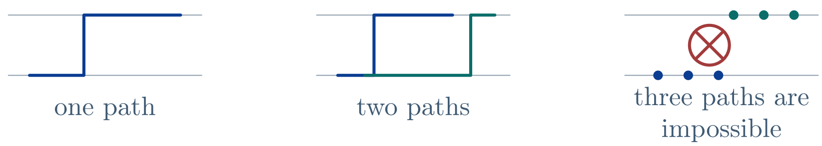

FIGURE 66

path, invariance is exactly the two-line isometry (2). When two paths occupy the strip, their total weight is unchanged because the Pitman transform preserves the sum of the two functions. There is no third case. It follows that, for any admissible entrance and exit tuples $\vec{p}, \vec{q}$,

$$d_f^{(k)}(\vec{p},\vec{q}) = d_{\sigma_i f}^{(k)}(\vec{p},\vec{q}).$$

Metric composition now allows the local identity to be applied successively through the reduced word for W. Hence

$$d_f^{(k)}\big((0,n)^k,(y,1)^k\big) = d_{Wf}^{(k)}\big((0,n)^k,(y,1)^k\big).$$

Finally, Wf is ordered by (2). In an ordered environment, the optimal k-tuple follows the top k lines, as illustrated in Figure 67. Its total weight is therefore

$$d_{Wf}^{(k)}\big((0,n)^k,(y,1)^k\big) = \sum_{i=1}^{k}(Wf)_i(y),$$

which proves Greene's identity. $\square$

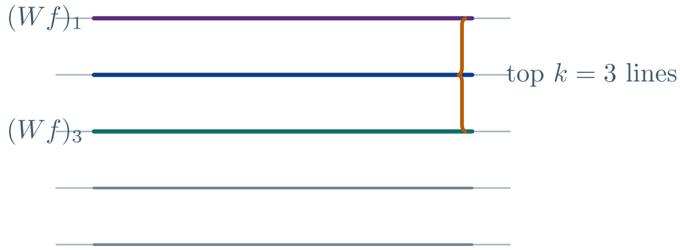

Figure 67

Lecture 13 — Friday, March 05, 2021

The Eden Model and Exponential First-Passage Percolation

The Eden model was introduced in 1961 as a model of cell growth. Begin with one occupied vertex at the origin. At each discrete step, choose uniformly from the edges having one occupied endpoint and one unoccupied endpoint, and adjoin the unoccupied endpoint to the cluster. Write $\mathcal{A}_k$ for the cluster after k additions, with $\mathcal{A}_0 = \{0\}$.

This discrete growth process is the **embedded jump chain** of first-passage percolation with exponential edge weights. Assign independent Exp(1) variables τ_e to the nearest-neighbour edges of $\mathbb{Z}^d$, and define

$$T(\boldsymbol{x}, \boldsymbol{y}) \coloneqq \inf_{\pi : \boldsymbol{x} \to \boldsymbol{y}} \sum_{e \in \pi} \tau_e.$$

Let

$$B(t) \coloneqq \{\boldsymbol{x} \in \mathbb{Z}^d : T(\mathbf{0}, \boldsymbol{x}) \le t\}$$

and define the time at which the growing ball first contains $k+1$ vertices by

$$\sigma_k \coloneqq \inf\{t \ge 0 : |B(t)| = k+1\}.$$

Proposition 13.1 The two cluster-valued processes

$$(\mathcal{A}_k)_{k \ge 0} \qquad \text{and} \qquad \big(B(\sigma_k)\big)_{k \ge 0}$$

have the same law.

Proof. Suppose the current first-passage cluster has K active boundary edges. By the memoryless property of the exponential distribution, the residual passage times along these edges are independent $\mathrm{Exp}(1)$ variables. Their minimum is $\mathrm{Exp}(K)$, and each boundary edge is equally likely to attain that minimum. Conditional on the present cluster, the next vertex is therefore obtained by choosing a boundary edge uniformly. This is exactly the Eden update rule. □

For every $\boldsymbol{x} \in \mathbb{Q}^d$, the subadditive argument leading to Proposition 3.4 gives a deterministic time constant $\mu(\boldsymbol{x})$ such that

$$\frac{T(\mathbf{0}, n\boldsymbol{x})}{n} \longrightarrow \mu(\boldsymbol{x})$$

almost surely and in L^1, along the integers n for which $n\boldsymbol{x} \in \mathbb{Z}^d$. The function μ extends to a norm on $\mathbb{R}^d$, and its unit ball is the deterministic limit shape. Even for independent exponential weights, the exact shape is unknown.

High-Dimensional Time Constants

Let

$$\mu_d \coloneqq \lim_{n\to\infty} \frac{T(\mathbf{0}, n\boldsymbol{e}_1)}{n}$$

be the coordinate direction time constant in dimension d.

Theorem 13.2 As $d \to \infty$,

$$\frac{d\,\mu_d}{\log d} \longrightarrow \frac{1}{2}.$$

Numerical bounds can reveal anisotropy of the limit shape. Introduce the coordinate hyperplane

$$L_n \coloneqq \{\boldsymbol{x} \in \mathbb{Z}^d : x_1 = n\}$$

and the diagonal hyperplane

$$J_n \coloneqq \left\{\boldsymbol{x} \in \mathbb{Z}^d : x_1 + \cdots + x_d = n\sqrt{d}\right\},$$

with the harmless convention that $n\sqrt{d}$ is replaced by a nearby integer when necessary. Then

$$\mu_d = \lim_{n\to\infty} \frac{T(\mathbf{0}, L_n)}{n} \qquad \text{and} \qquad \mu_d^\star \coloneqq \lim_{n\to\infty} \frac{T(\mathbf{0}, J_n)}{n}.$$

These two directions have the same Euclidean normalization. Thus $\mu_d \neq \mu_d^\star$ implies that the limit shape is not a Euclidean ball. The geometry of the two hyperplanes is shown in Figure 68.

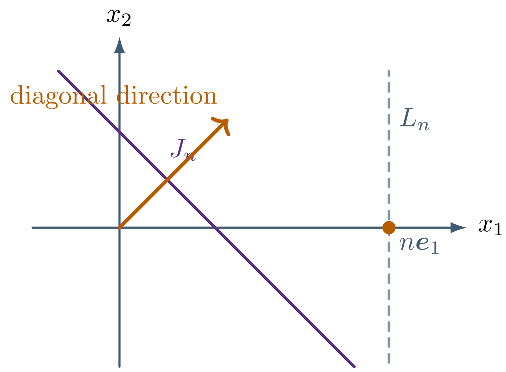

FIGURE 68

Theorem 13.3 Let $\alpha_\star$ be the unique positive solution of

$$\coth(\alpha) = \alpha.$$

Then

$$\mu_d^\star \geq \frac{\sqrt{\alpha_\star^2 - 1}}{2\sqrt{d}} \approx \frac{0.3313}{\sqrt{d}}.$$

For $d = 35$, Dhar's numerical estimate gives

$$\mu_{35} \leq \frac{0.93 \log(2d)}{2d} < \frac{0.3313}{\sqrt{d}} \leq \mu_{35}^\star.$$

Hence the exponential first-passage limit shape in dimension 35 is not a Euclidean ball.

Competition Models for First-Passage Percolation

Fix distinct seeds $\boldsymbol{x}_1, \boldsymbol{x}_2 \in \mathbb{Z}^d$. Two infections spread through the lattice, and a vertex retains permanently the type that reaches it first. Let

$$C_i \coloneqq \{\boldsymbol{y} \in \mathbb{Z}^d : \boldsymbol{y} \text{ is eventually occupied by infection } i\}.$$

The principal questions are the following:

1. Can both C_1 and C_2 be infinite?
2. If the infections have different rates, is coexistence possible?
3. What is the geometry and fluctuation behaviour of the interface separating the two infec-

tions?

Suppose type i crosses each active edge after an independent $\mathrm{Exp}(\lambda_i)$ waiting time. If type i currently has K_i active boundary edges, then its next infection time is

$$T_i \overset{\text{law}}{=} \min_{1\le j\le K_i} \tau_{i,j} \qquad \text{and} \qquad \tau_{i,j} \sim \mathrm{Exp}(\lambda_i),$$

and therefore

$$\mathbb{P}(T_i > t) = e^{-K_i\lambda_i t}.$$

The exponential race identity gives

$$\mathbb{P}(T_1 < T_2) = \frac{K_1\lambda_1}{K_1\lambda_1 + K_2\lambda_2} =: p.$$

Thus we may generate the next discrete update by first sampling a Bernoulli(p) variable to choose the infection type, and then choosing uniformly from that type's active boundary edges.

For equal rates on $\mathbb{Z}^2$, coexistence has positive probability. Let A denote the event that neither infection eventually surrounds the other; in this setting, A is the coexistence event. Write $\mathbb{P}_{\mathbf{0},\boldsymbol{x}}$ for the law with seeds at $\mathbf{0}$ and $\boldsymbol{x}$.

Theorem 13.4 The equal-rate competition model on $\mathbb{Z}^2$ has the following properties:

1. For nonzero $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{Z}^2$,
$$\mathbb{P}_{\mathbf{0},\boldsymbol{x}}(A) > 0 \quad \text{if and only if} \quad \mathbb{P}_{\mathbf{0},\boldsymbol{y}}(A) > 0.$$

2. If μ is the coordinate time constant in dimension 2, then
$$\mathbb{P}_{\mathbf{0},\boldsymbol{e}_1}(A) \geq \frac{4\mu - 1}{3} > 0.$$

By (2), coexistence occurs with positive probability for adjacent seeds; by (1), the same qualitative conclusion holds for every pair of distinct seeds.

When the rates differ, coexistence is exceptional.

Theorem 13.5 Fix distinct seeds $\boldsymbol{x}_1, \boldsymbol{x}_2$ and a rate $\lambda_1 > 0$. For all but countably many rates $\lambda_2 \neq \lambda_1$,
$$\mathbb{P}\big(|C_1| = |C_2| = \infty\big) = 0.$$

Comparable results are substantially more limited for nonexponential passage time distributions because the memoryless competition dynamics is no longer available.

Lecture 14 — Wednesday, March 10, 2021

Weakly Nonintersecting Bernoulli Walks

Fix a drift vector $\boldsymbol{d} = (d_1, \ldots, d_n) \in [0,1]^n$. Let $\boldsymbol{Y}^{\boldsymbol{d}} = (Y_1^{\boldsymbol{d}}, \ldots, Y_n^{\boldsymbol{d}})$ be independent Bernoulli walks starting at the origin, so that
$$Y_i^{\boldsymbol{d}}(t+1) - Y_i^{\boldsymbol{d}}(t) \sim \text{Bernoulli}(d_i).$$

The walks are interpolated linearly between consecutive integer times. Write
$$\boldsymbol{d}^\circ = (d_{(1)}, \ldots, d_{(n)}) \qquad \text{and} \qquad d_{(1)} \geq \cdots \geq d_{(n)},$$

for the decreasing rearrangement of the coordinates of $\boldsymbol{d}$. A typical Bernoulli path is shown in Figure 69.

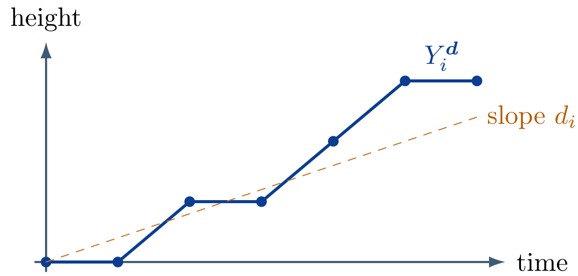

FIGURE 69

Definition 14.1 An ordered Bernoulli line ensemble $\boldsymbol{X} = (X_1, \ldots, X_n)$, where $X_1 \geq \cdots \geq X_n$ and every line starts at the origin, is an ensemble of **weakly nonintersecting Bernoulli walks with drift $\boldsymbol{d}^\circ$** if it has the following two properties:

1. It has the Gibbs property of Definition 6.4. Equivalently, conditional on the configuration from time t onward, its restriction to $[0, t]$ is uniform among all ordered Bernoulli path ensembles from the origin to $\boldsymbol{X}(t)$.
2. For every $i \in \{1, \ldots, n\}$,
$$\frac{X_i(t)}{t} \longrightarrow d_{(i)}$$
almost surely as $t \to \infty$.

The definition prescribes both local resampling and the asymptotic boundary condition. Neither existence nor uniqueness follows immediately from the definition; both are consequences of the functional Robinson–Schensted–Knuth correspondence.

Theorem 14.2 Let $\boldsymbol{d} \in [0,1]^n$.

1. There exists a unique law of weakly nonintersecting Bernoulli walks with drift $\boldsymbol{d}^\circ$.
2. If $\boldsymbol{Y^d}$ consists of independent Bernoulli walks with drift vector $\boldsymbol{d}$, then $W\boldsymbol{Y^d}$ has this unique law. In particular, the law of $W\boldsymbol{Y^d}$ depends on $\boldsymbol{d}$ only through the multiset of its coordinates.

Before proving the theorem, we record its metric consequence. The geometry behind this consequence is depicted in Figure 70.

Corollary 14.3 (Metric Burke property) Suppose that $\boldsymbol{d}, \boldsymbol{d'} \in [0,1]^n$ have the same decreasing rearrangement. Then the entire last-passage fields generated by $\boldsymbol{Y^d}$ and $\boldsymbol{Y^{d'}}$ have the same law:

$$\big(d_{\boldsymbol{Y^d}}((x,n),(y,1)) : 0 \le x \le y\big) \overset{\mathcal{L}}{=} \big(d_{\boldsymbol{Y^{d'}}}((x,n),(y,1)) : 0 \le x \le y\big).$$

Proof. The functional Robinson–Schensted–Knuth isometry (1) and (2) give

$$d_{\boldsymbol{Y^d}} = d_{W\boldsymbol{Y^d}} \overset{\mathcal{L}}{=} d_{W\boldsymbol{Y^{d'}}} = d_{\boldsymbol{Y^{d'}}}.$$

These identities hold simultaneously for every pair of admissible endpoints, which proves the equality of the full fields. ☐

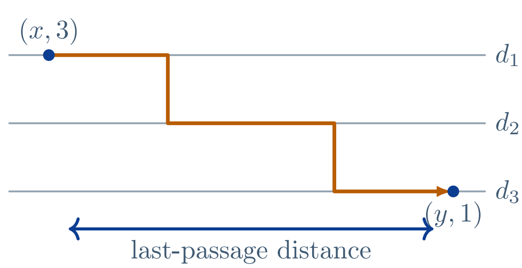

Figure 70

Proof of Theorem 14.2. We first prove the Gibbs property for two walks. Fix a terminal time t. The augmented two-line transform

$$(Y_1, Y_2) \longmapsto \big(W_t Y, Y_2(t)\big)$$

is a bijection by Proposition 6.2; see Figure 71. When $d_1 = d_2 = 1/2$, all input path pairs on $[0,t]$ have the same probability. After conditioning on the transformed endpoint $W_tY(t)$ and the retained value $Y_2(t)$, the transformed path is therefore uniform over the admissible ordered paths with that endpoint. For arbitrary $d_1, d_2 \in [0,1]$, an input pair $y = (y_1, y_2)$ has probability mass

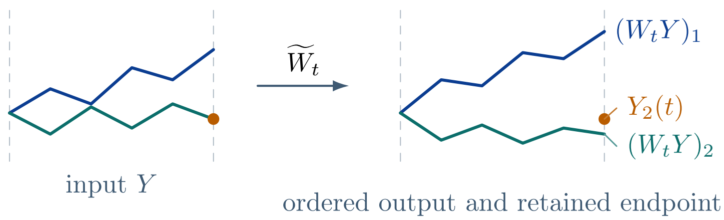

FIGURE 71

$$d_1^{y_1(t)}(1-d_1)^{t-y_1(t)} d_2^{y_2(t)}(1-d_2)^{t-y_2(t)}.$$

Here and below, powers with exponent zero are interpreted as 1. When a parameter equals 0 or 1, the corresponding path is deterministic, and the same argument applies after that path is fixed.

The transform preserves the sum of the two paths, so

$$y_1(t) = (W_ty)_1(t) + (W_ty)_2(t) - y_2(t).$$

Consequently, the input mass depends only on the ordered output endpoint and the retained endpoint. Conditional on those data, every admissible ordered output path again has equal mass. Summing over the retained endpoint shows that, conditional on $W_tY(t)$, the law of W_tY on $[0,t]$ is uniform over admissible ordered bridges. Thus WY has the Gibbs property.

For $n > 2$, apply the multiline bijection of Proposition 12.3. Its $\binom{n}{2}$ retained endpoint values provide exactly the auxiliary data needed to invert the successive two-line transforms. The probability of an input path tuple is

$$\prod_{i=1}^{n} d_i^{y_i(t)}(1-d_i)^{t-y_i(t)},$$

and each $y_i(t)$ is determined by the ordered output endpoint together with the retained data. The same conditioning argument proves the Gibbs property for $W\boldsymbol{Y}^{\boldsymbol{d}}$.

It remains to identify the asymptotic slopes. For $r > 0$, define the rescaled path tuple

$$\boldsymbol{Y}^{(r)}(s) := \frac{\boldsymbol{Y}^{\boldsymbol{d}}(rs)}{r}, \qquad s \geq 0.$$

The strong law of large numbers, applied coordinatewise, implies that $\boldsymbol{Y}^{(r)}$ converges almost surely and uniformly on compact time intervals to the linear path $s \mapsto s\boldsymbol{d}$. The multiline Pitman transform is a finite composition of running maximum operations and is therefore continuous in the topology of uniform convergence on compact intervals. It also commutes with this scaling. Hence

$$\frac{(W\boldsymbol{Y}^{\boldsymbol{d}})(r)}{r} = (W\boldsymbol{Y}^{(r)})(1) \longrightarrow W(s\boldsymbol{d})(1) = \boldsymbol{d}^{\circ}.$$

The last equality expresses the elementary fact, illustrated in Figure 72, that W sorts the slopes of a tuple of linear functions. This proves existence and the identification in (2). We now prove

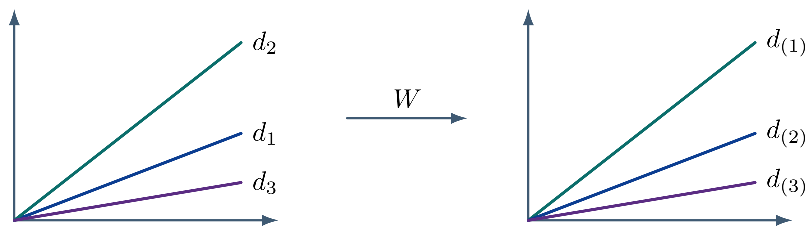

Figure 72

uniqueness. Two facts are needed.

1. For each finite set of times, the law of $W\boldsymbol{Y}^{\boldsymbol{d}}$ at those times is a continuous function of $\boldsymbol{d}$. Indeed, on a fixed finite time interval, every input path probability is a polynomial in the coordinates of $\boldsymbol{d}$ and $\boldsymbol{1} - \boldsymbol{d}$, while W is a deterministic map of the finite path space.
2. If $\boldsymbol{a}, \boldsymbol{b} \in [0,1]^n$ are decreasing and $a_i < b_i$ for every i, then any Gibbs ensemble with slope $\boldsymbol{b}$ stochastically dominates any Gibbs ensemble with slope $\boldsymbol{a}$.

To verify (2), let $\boldsymbol{X}^{\boldsymbol{a}}$ and $\boldsymbol{X}^{\boldsymbol{b}}$ be such ensembles. Their slope limits imply

$$\mathbb{P}\big(\boldsymbol{X}^{\boldsymbol{a}}(t) \leq \boldsymbol{X}^{\boldsymbol{b}}(t)\big) \longrightarrow 1$$

as $t \to \infty$, where the inequality is coordinatewise. Conditional on ordered endpoints at time t, both restrictions to $[0,t]$ are uniform ordered Bernoulli bridges. By Proposition 6.7, whenever $\boldsymbol{X}^{\boldsymbol{a}}(t) \leq \boldsymbol{X}^{\boldsymbol{b}}(t)$, these restrictions admit a coupling in which $\boldsymbol{X}^{\boldsymbol{a}} \leq \boldsymbol{X}^{\boldsymbol{b}}$. For any fixed initial time interval, the probability that this coupling fails tends to zero with t. Passing to the limit gives the claimed stochastic domination. The coupling is represented schematically in Figure 73. Finally, let $\boldsymbol{X}$ be any Gibbs ensemble with slope vector $\boldsymbol{d}^{\circ}$. Choose decreasing drift vectors $\boldsymbol{d}_r^-$ and $\boldsymbol{d}_r^+$ such that

$$\boldsymbol{d}_r^- < \boldsymbol{d}^{\circ} < \boldsymbol{d}_r^+, \qquad \boldsymbol{d}_r^- \longrightarrow \boldsymbol{d}^{\circ}, \qquad \text{and} \qquad \boldsymbol{d}_r^+ \longrightarrow \boldsymbol{d}^{\circ}.$$

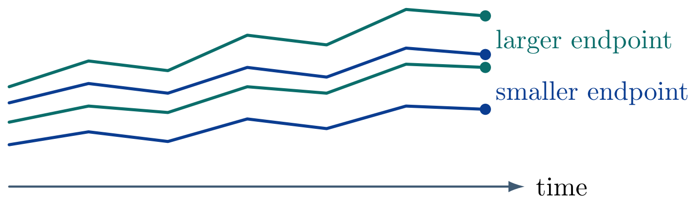

FIGURE 73

By (2),

$$W\boldsymbol{Y}^{\boldsymbol{d}_r^-} \preceq_{\mathrm{st}} \boldsymbol{X} \preceq_{\mathrm{st}} W\boldsymbol{Y}^{\boldsymbol{d}_r^+}.$$

Both bounding laws converge in every finite-dimensional distribution to the law of $W\boldsymbol{Y}^{\boldsymbol{d}^\circ}$ by (1). The squeeze therefore forces $\boldsymbol{X}$ to have that same law. If a coordinate of $\boldsymbol{d}^\circ$ equals 0 or 1, the conclusion follows by a one-sided approximation; the corresponding Bernoulli increments are deterministic. This proves uniqueness and completes the theorem. $\square$

Spacetime Determinantal Processes

Let $\boldsymbol{X}^{(m)}$ denote m weakly nonintersecting Bernoulli walks with drift 1/2. A central asymptotic problem is to understand the top line $X_1^{(m)}(n)$ when $m, n \to \infty$ at comparable rates. The two parameters are displayed in Figure 74. Direct formulas obtained from the Karlin–McGregor theorem become unwieldy as the number of walks grows. The **spacetime point process** provides a more effective description. Associate with a line ensemble $\boldsymbol{X}$ the random subset

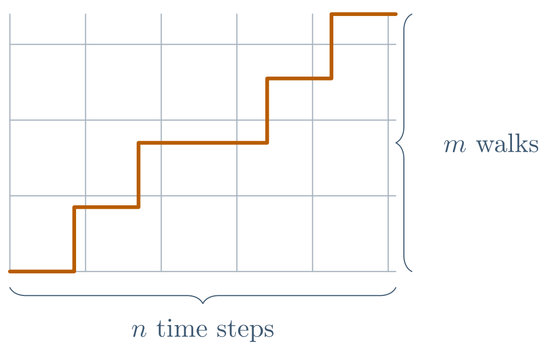

FIGURE 74

$$\Delta_{\boldsymbol{X}} := \{(t, X_i(t) - i) : t \in \mathbb{Z}_{\geq 0},\ 1 \leq i \leq m\} \subseteq \mathbb{Z}_{\geq 0} \times \mathbb{Z}.$$

The shift by the line index converts weak ordering into strict ordering, so every walk contributes a distinct point at each integer time, as in Figure 75. Instead of first computing the complete distribution of $\boldsymbol{X}(n)$, one studies the **inclusion probabilities**

$$\mathbb{P}(x \in \Delta_{\boldsymbol{X}}) \qquad \text{and} \qquad \mathbb{P}(x_1, \ldots, x_k \in \Delta_{\boldsymbol{X}}).$$

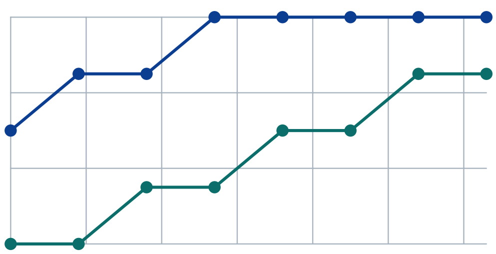

FIGURE 75

Theorem 14.4 The spacetime point process of weakly nonintersecting Bernoulli walks is a determinantal point process. Its **correlation kernel** can be obtained from the Eynard–Mehta theorem.

The proof and the explicit kernel are developed after the necessary determinantal identities have been established.

Definition 14.5 Let S be a countable discrete space. A random subset $\Delta \subseteq S$ is a **determinantal point process** if there is a kernel $K\colon S \times S \to \mathbb{C}$ such that, for every collection of distinct points $x_1, \ldots, x_k \in S$,

$$\mathbb{P}(x_1, \ldots, x_k \in \Delta) = \det\big[K(x_i, x_j)\big]_{i,j=1}^{k}.$$

Taking $k = 1$ in Definition 14.5 gives

$$K(x, x) = \mathbb{P}(x \in \Delta) \in [0, 1].$$

Thus the diagonal of the kernel records the one-point intensities, while its principal minors encode all higher inclusion probabilities.

Exercise 14.6 Let $\mathbf{L}$ be a matrix indexed by a finite set S. Prove that, when $\mathbf{L}$ is Hermitian, all of its principal minors are nonnegative if and only if $\mathbf{L}$ is positive semidefinite. Determine why the Hermitian hypothesis is essential by constructing a non-Hermitian counterexample.

Solution. *Forward implication.* Suppose that $\mathbf{L}$ is positive semidefinite. Every principal submatrix $\mathbf{L}_A$ is then positive semidefinite, so all of its eigenvalues are nonnegative. Hence $\det \mathbf{L}_A \geq 0$ for every $A \subseteq S$.

Reverse implication. Suppose that $\mathbf{L}$ is Hermitian and that all of its principal minors are nonnegative. For $\varepsilon > 0$ and $A \subseteq S$, the principal minor expansion gives

$$\det\bigl(\mathbf{L}_A + \varepsilon \mathbf{I}_A\bigr) = \sum_{B \subseteq A} \varepsilon^{|A|-|B|} \det \mathbf{L}_B > 0.$$

Thus every leading principal minor of $\mathbf{L} + \varepsilon \mathbf{I}$ is positive. Since this matrix is Hermitian, Sylvester's criterion shows that it is positive definite. Therefore, for every vector $\boldsymbol{z}$,

$$\boldsymbol{z}^* \mathbf{L} \boldsymbol{z} = \lim_{\varepsilon \downarrow 0} \boldsymbol{z}^* (\mathbf{L} + \varepsilon \mathbf{I}) \boldsymbol{z} \geq 0,$$

and $\mathbf{L}$ is positive semidefinite.

The Hermitian hypothesis cannot be omitted. The non-Hermitian matrix

$$\mathbf{N} = \begin{pmatrix} 1 & 4 \\ 0 & 1 \end{pmatrix}$$

has principal minors 1, 1, and 1, all of which are positive. Nevertheless, for $\boldsymbol{z} = (1, -1)^{\mathsf{T}}$,

$$\boldsymbol{z}^{\mathsf{T}} \mathbf{N} \boldsymbol{z} = -2 < 0.$$

Hence nonnegativity of all principal minors does not characterize positive semidefiniteness without Hermitian symmetry. $\square$

Example 14.7 (L-ensembles) Let S be finite and let $\mathbf{L}$ be a Hermitian positive semidefinite matrix indexed by S. For $A \subseteq S$, write $\mathbf{L}_A$ for the principal submatrix indexed by A, with $\det \mathbf{L}_\varnothing = 1$. The associated **L-ensemble** is defined by

$$\mathbb{P}(\Delta = A) := \frac{\det \mathbf{L}_A}{\det(\mathbf{I} + \mathbf{L})}$$

and is a point process on S. Indeed, the principal minor expansion gives

$$\det(\mathbf{I} + \mathbf{L}) = \sum_{A \subseteq S} \det \mathbf{L}_A,$$

and every summand is nonnegative. This process is determinantal with kernel

$$\mathbf{K} = \mathbf{L}(\mathbf{I} + \mathbf{L})^{-1}.$$

To verify that this process is determinantal, let $\mathbf{Z} = \operatorname{diag}(z_x : x \in S)$. The generating function of the inclusion probabilities is

$$\mathbb{E}\left[\prod_{x \in \Delta} (1 + z_x)\right] = \frac{\det(\mathbf{I} + \mathbf{L}(\mathbf{I} + \mathbf{Z}))}{\det(\mathbf{I} + \mathbf{L})} = \det(\mathbf{I} + \mathbf{K}\mathbf{Z}).$$

Comparing the coefficient of $\prod_{x \in A} z_x$ on both sides yields $\mathbb{P}(A \subseteq \Delta) = \det \mathbf{K}_A$, which is precisely the determinantal property.

Lecture 15 — Friday, March 12, 2021

3. Airy Line Ensembles and Determinantal Methods

The Airy Line Ensemble

We use the stationary convention for the Airy line ensemble. This convention separates the determinantal description from the parabolic shift that reveals the Brownian Gibbs property.

Definition 15.1 The **Airy line ensemble** is a random sequence $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \dots)$ of continuous functions on $\mathbb{R}$ such that

$$\mathcal{A}_1(t) > \mathcal{A}_2(t) > \cdots$$

almost surely for every $t \in \mathbb{R}$. Its spacetime point process is determinantal with the **extended Airy kernel**

$$K_{\mathrm{Ai}}\big((x,s),(y,t)\big) = \begin{cases} \displaystyle\int_0^\infty e^{-\lambda(s-t)}\,\mathrm{Ai}(x+\lambda)\,\mathrm{Ai}(y+\lambda)\,\mathrm{d}\lambda, & s \ge t, \\ \displaystyle -\int_{-\infty}^0 e^{-\lambda(s-t)}\,\mathrm{Ai}(x+\lambda)\,\mathrm{Ai}(y+\lambda)\,\mathrm{d}\lambda, & s < t. \end{cases}$$

The law of $\mathcal{A}(\,\cdot + s)$ does not depend on s. In particular, $\mathcal{A}_1$ is the stationary Airy process. It is also convenient to introduce the parabolically shifted and Brownian-normalized ensembles

$$\mathcal{P}_i(t) \coloneqq \mathcal{A}_i(t) - t^2 \qquad \text{and} \qquad \mathcal{L}_i(t) \coloneqq 2^{-1/2}\mathcal{P}_i(t).$$

The spacetime point process of $\mathcal{P}$ has the equivalent shifted kernel

$$K_{\mathcal{P}}\big((x,s),(y,t)\big) = \begin{cases} \displaystyle\int_0^\infty e^{-\lambda(s-t)}\,\mathrm{Ai}(x+s^2+\lambda)\,\mathrm{Ai}(y+t^2+\lambda)\,\mathrm{d}\lambda, & s \ge t, \\ \displaystyle -\int_{-\infty}^0 e^{-\lambda(s-t)}\,\mathrm{Ai}(x+s^2+\lambda)\,\mathrm{Ai}(y+t^2+\lambda)\,\mathrm{d}\lambda, & s < t. \end{cases}$$

This is the same kernel written in the moving coordinates $(t,x) \mapsto (t, x+t^2)$.

Brownian Bridges at the Spectral Edge

For each N, let $B^N = (B_1^N, \dots, B_N^N)$ be N standard Brownian bridges on $[-N, N]$, all with value 0 at the two endpoints, conditioned to satisfy

$$B_1^N(t) > \cdots > B_N^N(t), \qquad -N < t < N.$$

Although the curves meet at the terminal points, they are strictly ordered in the interior. At any finite set of times, their positions form a determinantal point process. The geometry is illustrated in Figure 76. Define the **edge-scaled Dyson line ensemble** by

$$\begin{aligned} \boldsymbol{D}^N &= (D_1^N, \dots, D_N^N), \\ D_i^N(t) &\coloneqq N^{-1/3}\left(B_i^N(N^{2/3}t) - \sqrt{2}\,N\right), \end{aligned}$$

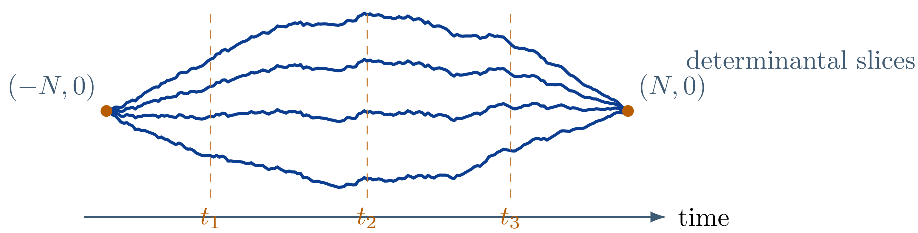

FIGURE 76

$$\widetilde{D}_i^N(t) := \sqrt{2}\, D_i^N(t) + t^2.$$

The horizontal scale $N^{2/3}$ and vertical scale $N^{1/3}$ magnify the spectral edge. Subtracting $\sqrt{2}N$ centers the top curves. The transformation is depicted in Figure 77. For every fixed collection of

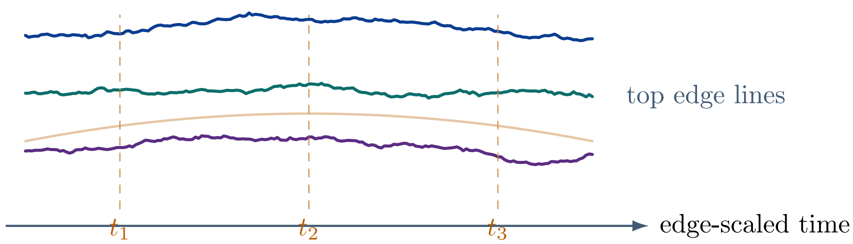

FIGURE 77

times, the point process associated with $\widetilde{\boldsymbol{D}}^N$ has a correlation kernel K_N. After conjugating by an irrelevant **determinantal gauge factor**, K_N converges locally in **trace norm** to K_{Ai}. Consequently, the finite-dimensional distributions of $\widetilde{\boldsymbol{D}}^N$ converge to those prescribed in Definition 15.1.

Definition 15.2 An ordered continuous line ensemble $L = (L_1, L_2, \dots)$ has the **Brownian Gibbs property** if the following resampling rule holds. Fix integers $1 \le k_1 \le k_2$ and an interval $[a, b]$. Conditional on all curves outside

$$\{k_1, \dots, k_2\} \times (a, b),$$

the curves $L_{k_1}, \dots, L_{k_2}$ on $[a, b]$ have the law of independent standard Brownian bridges between their revealed endpoints, conditioned to remain mutually ordered and between the neighbouring curves L_{k_1-1} and L_{k_2+1}. When $k_1 = 1$, the upper boundary is understood to be $+\infty$.

Theorem 15.3 There is a unique probability measure on continuous, strictly ordered line ensembles whose finite-dimensional distributions are those in Definition 15.1. Moreover, the parabolically shifted and normalized ensemble

$$\mathcal{L}_i(t) = 2^{-1/2}\big(\mathcal{A}_i(t) - t^2\big)$$

has the Brownian Gibbs property.

Tightness and Passage of the Gibbs Property

Finite-dimensional convergence alone does not give convergence as random continuous functions. We must also control the curves uniformly on compact intervals. For fixed k and $T > 0$, define the smallest gap among the first k curves by

$$M_{k,T}^N := \min_{1 \le i \le k-1} \min_{s \in [-T,T]} \big|D_i^N(s) - D_{i+1}^N(s)\big|.$$

The following two estimates isolate the required compactness input.

Proposition 15.4 For every $k \ge 2$, $T > 0$, and $\varepsilon > 0$, there are $\delta > 0$ and N_0 such that

$$\mathbb{P}(M_{k,T}^N < \delta) < \varepsilon$$

for every $N \ge N_0$.

Proposition 15.5 For every $k \ge 1$, $T > 0$, and $\varepsilon > 0$, there are $R > 0$ and N_0 such that

$$\mathbb{P}\left(\max_{1 \le i \le k} \sup_{t \in [-T,T]} \big|D_i^N(t)\big| > R\right) < \varepsilon$$

for every $N \ge N_0$.

For $k \le N$, write

$$\boldsymbol{D}^{N,k}(t) := (D_1^N(t), \dots, D_k^N(t)).$$

The one-time determinantal convergence gives tightness of the vectors $\boldsymbol{D}^{N,k}(0)$. On the event supplied by Propositions 15.4 and 15.5, the Brownian Gibbs property compares the modulus of continuity of the curves with that of Brownian bridges in a bounded corridor. Standard bridge estimates then show that, for every $\eta > 0$,

$$\lim_{\delta \downarrow 0} \sup_{N \geq N_0} \mathbb{P}\left(\max_{1 \leq i \leq k} \sup_{\substack{s,t \in [-T,T] \\ |s-t| \leq \delta}} \left| D_i^N(s) - D_i^N(t) \right| > \eta \right) = 0.$$

The Arzelà–Ascoli tightness criterion therefore makes $\boldsymbol{D}^{N,k}$ tight in $C([-T,T])^k$. A diagonal argument over k and T produces a subsequential limit $\boldsymbol{L}^\infty = (L_1^\infty, L_2^\infty, \dots)$. Its finite-dimensional distributions identify it with $\mathcal{L}$, and Proposition 15.4 ensures strict ordering.

It remains to pass the Brownian Gibbs property to the limit. The acceptance-rejection construction makes this step explicit. By the Skorokhod representation theorem, a convergent subsequence and its limit may be realized on one probability space so that convergence is almost surely uniform on every compact set.

Fix a resampling block $\{k_1, \dots, k_2\} \times [a,b]$. For each positive integer r, let $(\beta_{k_1}^r, \dots, \beta_{k_2}^r)$ be an independent tuple of standard Brownian bridges on $[a,b]$ with zero endpoints. Define the r-th candidate bridge for the i-th curve by

$$D_i^{N,r}(t) := \beta_i^r(t) + \frac{b-t}{b-a} D_i^N(a) + \frac{t-a}{b-a} D_i^N(b).$$

Let ℓ_N be the first r for which the candidate curves are mutually ordered and lie strictly between the revealed upper and lower boundary curves throughout $[a,b]$. Then the accepted tuple $(D_i^{N,\ell_N} : k_1 \leq i \leq k_2)$ has exactly the conditional bridge law in Definition 15.2. Replacing the original curves by this tuple therefore leaves the law of $\boldsymbol{D}^N$ unchanged.

Uniform convergence of the endpoints and boundary curves implies

$$\ell_N \longrightarrow \ell_\infty$$

almost surely, where ℓ_∞ is the first accepted candidate for the limiting ensemble, defined by the same rule. Indeed, each fixed candidate is either separated from the limiting boundary or violates an ordering inequality by a definite amount; a Brownian bridge has probability zero of merely touching a continuous boundary without crossing it at the first critical contact. Hence the accepted resamplings converge uniformly. Since every prelimit resampling preserves the law, the limiting resampling preserves the law as well. This proves the Brownian Gibbs assertion in Theorem 15.3.

Local Absolute Continuity and the Unique Maximizer

Proposition 15.6 For $s \in \mathbb{R}$, $t > 0$, and $k \geq 1$, the increment process

$$\big(\mathcal{A}_k(s+u) - \mathcal{A}_k(s) : 0 \leq u \leq t\big)$$

is absolutely continuous with respect to Brownian motion of variance parameter 2, after adding the deterministic parabolic drift. More precisely, the centred process

$$\mathcal{A}_k(s+u) - \left(\frac{t-u}{t}\mathcal{A}_k(s) + \frac{u}{t}\mathcal{A}_k(s+t)\right), \qquad 0 \leq u \leq t,$$

is absolutely continuous with respect to a variance-2 Brownian bridge from 0 to 0.

Proof. Condition on the exterior curves. Then the Brownian Gibbs law is an ordinary bridge law restricted to a positive probability nonintersection event. □

Theorem 15.7 Let

$$H(t) \coloneqq \mathcal{A}_1(t) - t^2.$$

For every $T > 0$, the function H attains its maximum on $[-T, T]$ at a unique point almost surely. It also attains its maximum on $\mathbb{R}$ at a unique point almost surely.

Proof. On a compact interval, the absolute continuity in Proposition 15.6 reduces uniqueness to the corresponding fact for Brownian motion with a continuous deterministic drift, whose maximizer is almost surely unique. For the whole line, stationarity and the upper-tail bounds for $\mathcal{A}_1(t)$, combined with the Brownian modulus estimate and the Borel–Cantelli lemma, give

$$\mathcal{A}_1(t) - t^2 \longrightarrow -\infty$$

almost surely as $|t| \to \infty$. The global maximum is therefore attained on some random compact interval, where it is unique by the first part. □

Lecture 16 — Wednesday, March 17, 2021

The Eynard–Mehta Theorem

The Eynard–Mehta theorem converts a product of determinants into the correlation kernel of a determinantal point process. This is precisely the form produced by the Karlin–McGregor formula for nonintersecting Markov paths. It remains useful when both the number of paths and the number of time steps diverge, as suggested by Figure 78.

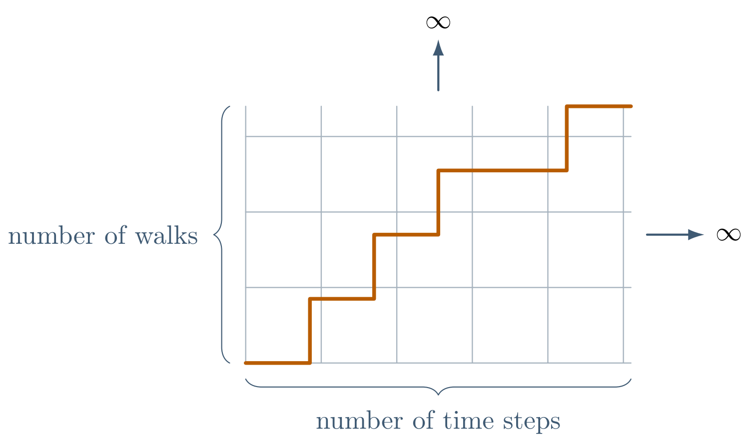

FIGURE 78

Nonintersecting Markov Bridges

Let $p_{s,t}(x,y)$ be the transition probability of a time-homogeneous Markov chain on a discrete state space. Thus

$$p_{s,t}(x,y) = p_{0,t-s}(x,y).$$

For example, the transition probability of a Bernoulli walk with parameter p is

$$p_{0,t}(x,y) = \binom{t}{y-x} p^{y-x}(1-p)^{t-y+x}, \qquad 0 \le y-x \le t,$$

and is zero otherwise. If $\boldsymbol{x} = (x_1 < \cdots < x_n)$ and $\boldsymbol{y} = (y_1 < \cdots < y_n)$, define

$$q_{s,t}(\boldsymbol{x},\boldsymbol{y}) \coloneqq \det\big[p_{s,t}(x_i,y_j)\big]_{i,j=1}^{n}.$$

By the Karlin–McGregor theorem, this determinant is the transition probability for n walks to move from $\boldsymbol{x}$ to $\boldsymbol{y}$ without intersecting. The bridge geometry is shown in Figure 79. Fix times

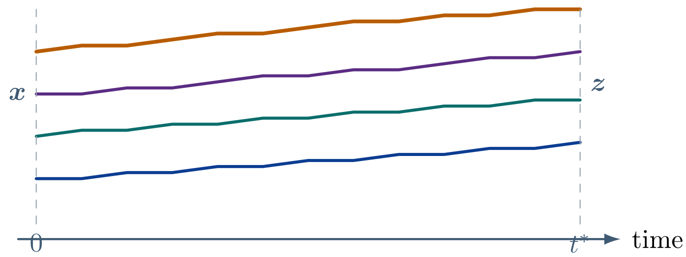

FIGURE 79

$$0 = t_0 < t_1 < \cdots < t_m < t_{m+1} = t^*$$

and endpoint configurations $\boldsymbol{y}^{(0)} = \boldsymbol{x}$ and $\boldsymbol{y}^{(m+1)} = \boldsymbol{z}$. The joint law of a nonintersecting bridge at the intermediate times is

$$\mathbb{P}\big(\boldsymbol{X}(t_1) = \boldsymbol{y}^{(1)}, \dots, \boldsymbol{X}(t_m) = \boldsymbol{y}^{(m)} \mid \boldsymbol{X}(0) = \boldsymbol{x},\ \boldsymbol{X}(t^*) = \boldsymbol{z}\big) = \frac{\prod_{r=0}^{m} q_{t_r, t_{r+1}}\big(\boldsymbol{y}^{(r)}, \boldsymbol{y}^{(r+1)}\big)}{q_{0,t^*}(\boldsymbol{x}, \boldsymbol{z})}. \quad (16.1)$$

Every factor in the numerator of (16.1) is a determinant. The Eynard–Mehta theorem applies to exactly this product.

A General Product-of-Determinants Measure

Let S be a finite set and let n be the number of particles. Fix matrices

$$\boldsymbol{\Phi} \colon \{1, \dots, n\} \times S \longrightarrow \mathbb{C} \qquad \text{and} \qquad \boldsymbol{\Psi} \colon S \times \{1, \dots, n\} \longrightarrow \mathbb{C},$$

together with matrices $\mathbf{W}_1, \dots, \mathbf{W}_{m-1}$ indexed by $S \times S$. These matrices connect the successive levels in Figure 80. For each level r, let $Y_r = \{y_1^{(r)} < \cdots < y_n^{(r)}\}$ be an n-element subset of S.

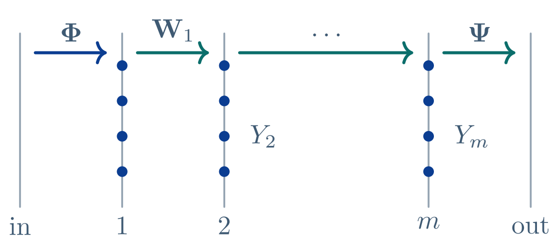

FIGURE 80

Assign the configuration $\boldsymbol{Y} = (Y_1, \ldots, Y_m)$ the weight

$$w(\boldsymbol{Y}) \coloneqq \det\big[\boldsymbol{\Phi}(i, y_j^{(1)})\big]_{i,j=1}^n \prod_{r=1}^{m-1} \det\big[\mathbf{W}_r(y_i^{(r)}, y_j^{(r+1)})\big]_{i,j=1}^n \det\big[\boldsymbol{\Psi}(y_i^{(m)}, j)\big]_{i,j=1}^n. \tag{16.2}$$

Assume that all these weights are nonnegative and that their sum is positive. Define

$$\mathbf{M} \coloneqq \boldsymbol{\Phi}\mathbf{W}_1 \cdots \mathbf{W}_{m-1}\boldsymbol{\Psi}.$$

For $1 \leq r, s \leq m$, set

$$\boldsymbol{\Phi}_s \coloneqq \boldsymbol{\Phi}\mathbf{W}_1 \cdots \mathbf{W}_{s-1} \qquad \text{and} \qquad \boldsymbol{\Psi}_r \coloneqq \mathbf{W}_r \cdots \mathbf{W}_{m-1}\boldsymbol{\Psi},$$

where an empty product is the appropriate identity matrix. Also set

$$\mathbf{W}_{r,s} \coloneqq \begin{cases} \mathbf{W}_r\mathbf{W}_{r+1} \cdots \mathbf{W}_{s-1}, & r < s, \\ 0, & r \geq s. \end{cases}$$

Theorem 16.1 (Eynard–Mehta theorem) Suppose that $\mathbf{M}$ is invertible. The normalization of (16.2) is

$$Z = \sum_{\boldsymbol{Y}} w(\boldsymbol{Y}) = \det \mathbf{M}. \tag{16.3}$$

Moreover, the random subset

$$\mathcal{X} \coloneqq \bigcup_{r=1}^{m} \big(\{r\} \times Y_r\big) \subseteq \{1, \ldots, m\} \times S$$

is a determinantal point process with correlation kernel

$$K\big((r,x),(s,y)\big) = \big[\boldsymbol{\Psi}_r\mathbf{M}^{-1}\boldsymbol{\Phi}_s\big]_{x,y} - \big[\mathbf{W}_{r,s}\big]_{x,y}. \tag{16.4}$$

Thus, for distinct points $u_1, \ldots, u_\ell$ in $\{1, \ldots, m\} \times S$,

$$\mathbb{P}(u_1, \ldots, u_\ell \in \mathcal{X}) = \det\big[K(u_i, u_j)\big]_{i,j=1}^{\ell}.$$

The first term in (16.4) propagates from level r back to the terminal data and from the initial data forward to level s. The second term removes direct forward propagation from level r to level s. In applications, the principal calculation is the inversion of the finite $n \times n$ matrix $\mathbf{M}$.

Proof of (16.3)**.** Apply the Cauchy–Binet identity successively to the sum over $Y_1, \dots, Y_m$. Each summation contracts two adjacent determinants into the determinant of the corresponding matrix product. After all m sums have been performed, the result is

$$\det\big(\mathbf{\Phi}\mathbf{W}_1 \cdots \mathbf{W}_{m-1}\mathbf{\Psi}\big) = \det \mathbf{M}.$$

$\square$

Proposition 16.2 (Cauchy–Binet identity) Let $\mathbf{A}$ be an $r \times q$ matrix and let $\mathbf{B}$ be a $q \times r$ matrix, where $q \geq r$. Then

$$\det(\mathbf{AB}) = \sum_{\substack{I \subseteq \{1,\dots,q\} \\ |I|=r}} \det \mathbf{A}_{[:,I]} \det \mathbf{B}_{[I,:]}.$$

Proof. Expand the determinant of $\mathbf{AB}$ multilinearly in its columns. Terms with a repeated intermediate index vanish by antisymmetry. Grouping the remaining terms by their r-element set of intermediate indices gives the stated sum. $\square$

L-Ensembles and Conditional L-Ensembles

The proof of the kernel formula (16.4) uses the $\mathbf{L}$-ensemble from Example 14.7. For a finite state space S and a matrix $\mathbf{L}$ whose principal minors are nonnegative, the law is

$$\mathbb{P}(X = A) = \frac{\det \mathbf{L}_A}{\det(\mathbf{I} + \mathbf{L})}, \qquad A \subseteq S.$$

If $\mathbf{L} = \mathbf{0}$, then $X = \varnothing$ almost surely. If $\mathbf{L} = \mathbf{I}$, every subset has probability $2^{-|S|}$, so X is independent Bernoulli percolation with parameter $1/2$. No finite $\mathbf{L}$-ensemble on a nonempty state space is concentrated on the full set, because $\mathbb{P}(X = \varnothing) = 1/\det(\mathbf{I} + \mathbf{L}) > 0$.

Remark 16.3 For a real symmetric matrix $\mathbf{L}$, nonnegativity of every principal minor is equivalent to positive semidefiniteness. For a general real matrix, write

$$\mathbf{L} = \mathbf{A} + \mathbf{B}, \qquad \mathbf{A} = \frac{\mathbf{L} + \mathbf{L}^\mathsf{T}}{2} \qquad \text{and} \qquad \mathbf{B} = \frac{\mathbf{L} - \mathbf{L}^\mathsf{T}}{2}.$$

The matrix $\mathbf{B}$ is skew-symmetric and $\boldsymbol{v}^\mathsf{T}\mathbf{B}\boldsymbol{v} = 0$, so $\boldsymbol{v}^\mathsf{T}\mathbf{L}\boldsymbol{v} \geq 0$ for every $\boldsymbol{v}$ precisely when $\mathbf{A}$ is positive semidefinite. This condition implies that all principal minors of $\mathbf{L}$ are nonnegative. Indeed, after a positive definite approximation,

$$\det(\mathbf{A}+\mathbf{B}) = \det(\mathbf{A})\det\bigl(\mathbf{I}+\mathbf{A}^{-1/2}\mathbf{B}\mathbf{A}^{-1/2}\bigr) \geq 0,$$

because the eigenvalues of the skew-symmetric factor occur in pairs $\pm i\lambda$. The converse fails: the matrix

$$\mathbf{N} = \begin{pmatrix} 1 & 3 \\ 0 & 1 \end{pmatrix}$$

has nonnegative principal minors, but its symmetric part is not positive semidefinite.

Conditional L-ensembles allow part of the configuration to be forced. For $A \subseteq S$, let $\mathbf{I}_A$ denote the diagonal matrix with ones on A and zeros on A^c. Condition the $\mathbf{L}$-ensemble on the event that every point of A^c is present, and regard only the random points in A.

Proposition 16.4 Assume that $\mathbf{I}_A + \mathbf{L}$ is invertible. The conditional $\mathbf{L}$-ensemble on A is determinantal with kernel

$$\mathbf{K}_A = \mathbf{I} - \bigl[(\mathbf{I}_A+\mathbf{L})^{-1}\bigr]_{A\times A}. \tag{16.5}$$

Proof. The normalization of the conditional law is the principal minor expansion

$$\det(\mathbf{I}_A+\mathbf{L}) = \sum_{Y\subseteq A} \det \mathbf{L}_{Y\cup A^c}.$$

Let $\mathbf{Q} = \mathbf{I}_A + \mathbf{L}$. The Jacobi complementary minor identity states that

$$\det\bigl[(\mathbf{Q}^{-1})_C\bigr] = \frac{\det \mathbf{Q}_{C^c}}{\det \mathbf{Q}}, \qquad C \subseteq A.$$

For $B \subseteq A$, the principal minor expansion and this identity give

$$\begin{aligned}
\det\bigl[(\mathbf{K}_A)_B\bigr] &= \sum_{C\subseteq B} (-1)^{|C|} \det\bigl[(\mathbf{Q}^{-1})_C\bigr] \\
&= \sum_{C\subseteq B} (-1)^{|C|} \mathbb{P}(C\cap X = \varnothing) \\
&= \mathbb{P}(B\subseteq X).
\end{aligned}$$

The last equality is inclusion-exclusion. These are exactly the determinantal inclusion probabilities.

□

Taking $A = S$ in (16.5) recovers

$$\mathbf{K} = \mathbf{I} - (\mathbf{I} + \mathbf{L})^{-1} = \mathbf{L}(\mathbf{I} + \mathbf{L})^{-1},$$

the ordinary L-ensemble kernel from Example 14.7.

Encoding the Layered Measure

Introduce an n-point virtual level V, and let

$$\mathfrak{S} \coloneqq V \sqcup \big(\{1, \dots, m\} \times S\big).$$

With the blocks ordered as $V, 1, \dots, m$, define

$$\mathbf{L} = \begin{pmatrix} 0 & -\mathbf{\Phi} & 0 & \cdots & 0 \\ 0 & 0 & -\mathbf{W}_1 & \cdots & 0 \\ 0 & 0 & 0 & \ddots & \vdots \\ \vdots & \vdots & \vdots & \ddots & -\mathbf{W}_{m-1} \\ \mathbf{\Psi} & 0 & 0 & \cdots & 0 \end{pmatrix}. \tag{16.6}$$

The physical state space $A = \{1, \dots, m\} \times S$ is random, whereas all points of $V = A^c$ are forced to be present. Expanding a principal minor of (16.6) around its cyclic block structure yields exactly the product of determinants in (16.2). Thus the layered process in Theorem 16.1 is a conditional $\mathbf{L}$-ensemble. It remains to invert $\mathbf{I}_A + \mathbf{L}$ and insert its physical block into (16.5).

Lecture 17 — Friday, March 19, 2021

The Conditional Ensemble Encoding

We first justify the conditional ensemble encoding. Consider three nonintersecting paths moving through a layered directed graph, as in Figure 81. The matrix $\mathbf{\Phi}$ connects the virtual source vertices to the first layer, the matrices $\mathbf{W}_r$ connect consecutive physical layers, and $\mathbf{\Psi}$ connects the last layer back to the virtual vertices. The determinant of a matrix is a signed sum over permutations. For a principal submatrix of $\mathbf{L}$ in (16.6), a nonzero permutation must follow the

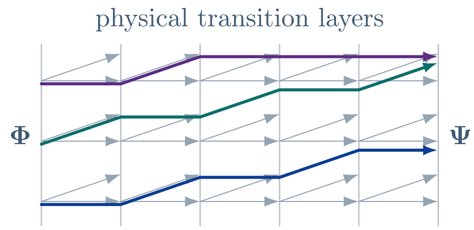

FIGURE 81

directed cycle of nonzero blocks

$$V \longrightarrow 1 \longrightarrow 2 \longrightarrow \cdots \longrightarrow m \longrightarrow V.$$

Because the forced virtual level has n points, every physical level must also contain exactly n points. Once $Y_r = \{y_1^{(r)}, \ldots, y_n^{(r)}\}$ is fixed at each level, the block determinant factors as

$$\begin{aligned}
\det \mathbf{L}_{V\cup\mathcal{X}} = \det\bigl[\mathbf{\Phi}(i, y_j^{(1)})\bigr]_{i,j=1}^{n} \\
\times \prod_{r=1}^{m-1} \det\bigl[\mathbf{W}_r(y_i^{(r)}, y_j^{(r+1)})\bigr]_{i,j=1}^{n} \\
\times \det\bigl[\mathbf{\Psi}(y_i^{(m)}, j)\bigr]_{i,j=1}^{n} \\
= w(\boldsymbol{Y}).
\end{aligned}$$

The minus signs in (16.6) cancel the sign of the cyclic block permutation. This proves that the measure of Theorem 16.1 is the conditional **L**-ensemble obtained by forcing all virtual points.

The same determinant expansion has a useful graph interpretation. Identify the outgoing virtual vertices with the incoming ones, wrapping the layered graph around a cylinder as in Figure 82. A nonzero determinant term is then a vertex-disjoint union of directed cycles. Cutting the cylinder along the virtual level turns those cycles into the nonintersecting paths of the layered ensemble.

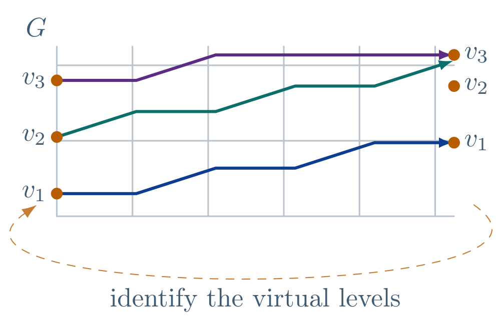

FIGURE 82

Block Inversion and the Correlation Kernel

Let $\mathbf{Q} = \mathbf{I}_A + \mathbf{L}$, where $\mathbf{I}_A$ is the identity on the physical levels and is zero on V. Separate the virtual block from the physical blocks:

$$\mathbf{Q} = \begin{pmatrix} 0 & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{pmatrix}, \tag{17.1}$$

where

$$\mathbf{B} = \begin{pmatrix} -\boldsymbol{\Phi} & 0 & \cdots & 0 \end{pmatrix} \qquad \text{and} \qquad \mathbf{C} = \begin{pmatrix} 0 \\ \vdots \\ 0 \\ \boldsymbol{\Psi} \end{pmatrix},$$

and

$$\mathbf{D} = \begin{pmatrix} \mathbf{I} & -\mathbf{W}_1 & 0 & \cdots & 0 \\ 0 & \mathbf{I} & -\mathbf{W}_2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & \mathbf{I} & -\mathbf{W}_{m-1} \\ 0 & \cdots & \cdots & 0 & \mathbf{I} \end{pmatrix}. \tag{17.2}$$

The sparsity patterns of $\mathbf{Q}$ and $\mathbf{D}^{-1}$ are shown schematically in Figure 83. Since $\mathbf{D}$ is block upper

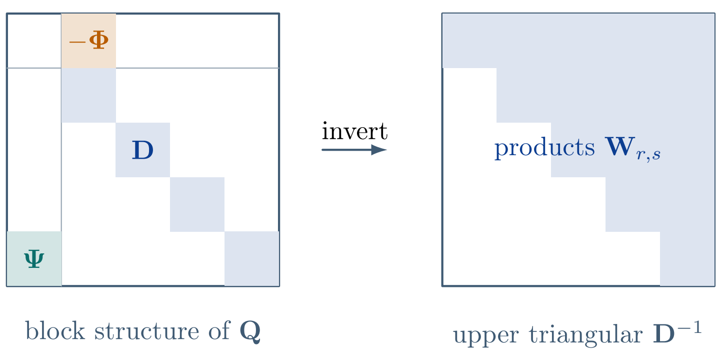

Figure 83

triangular with identity diagonal blocks, its inverse is

$$\mathbf{D}^{-1} = \begin{pmatrix} \mathbf{I} & \mathbf{W}_{1,2} & \mathbf{W}_{1,3} & \cdots & \mathbf{W}_{1,m} \\ 0 & \mathbf{I} & \mathbf{W}_{2,3} & \cdots & \mathbf{W}_{2,m} \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & \mathbf{I} & \mathbf{W}_{m-1,m} \\ 0 & \cdots & \cdots & 0 & \mathbf{I} \end{pmatrix}. \tag{17.3}$$

Indeed, multiplying (17.2) and (17.3) produces telescoping cancellations above the diagonal.

For a block matrix

$$\mathbf{R} = \begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{pmatrix}$$

with $\mathbf{D}$ invertible, let $\mathbf{S} = \mathbf{A} - \mathbf{B}\mathbf{D}^{-1}\mathbf{C}$ be the **Schur complement**. If $\mathbf{S}$ is also invertible, then

$$\mathbf{R}^{-1} = \begin{pmatrix} \mathbf{S}^{-1} & -\mathbf{S}^{-1}\mathbf{B}\mathbf{D}^{-1} \\ -\mathbf{D}^{-1}\mathbf{C}\mathbf{S}^{-1} & \mathbf{D}^{-1} + \mathbf{D}^{-1}\mathbf{C}\mathbf{S}^{-1}\mathbf{B}\mathbf{D}^{-1} \end{pmatrix}. \tag{17.4}$$

For $\mathbf{Q}$ in (17.1), the upper-left block is zero and

$$\mathbf{S} = -\mathbf{B}\mathbf{D}^{-1}\mathbf{C} = \mathbf{\Phi}\mathbf{W}_1 \cdots \mathbf{W}_{m-1}\mathbf{\Psi} = \mathbf{M}.$$

Furthermore, the r-th block of $\mathbf{D}^{-1}\mathbf{C}$ is $\mathbf{\Psi}_r$, while the s-th block of $\mathbf{B}\mathbf{D}^{-1}$ is $-\mathbf{\Phi}_s$. Therefore the physical block of $\mathbf{Q}^{-1}$ is

$$\big[\mathbf{Q}^{-1}\big]_{r,s} = \mathbb{1}_{\{r=s\}}\mathbf{I} + \mathbf{W}_{r,s} - \mathbf{\Psi}_r\mathbf{M}^{-1}\mathbf{\Phi}_s. \tag{17.5}$$

Insert (17.5) into the conditional $\mathbf{L}$-ensemble formula (16.5). The identity term cancels, leaving

$$K\big((r,x),(s,y)\big) = \big[\mathbf{\Psi}_r\mathbf{M}^{-1}\mathbf{\Phi}_s\big]_{x,y} - \big[\mathbf{W}_{r,s}\big]_{x,y}.$$

This is exactly (16.4) and completes the proof of Theorem 16.1.

Kernel for Nonintersecting Bridges

For the bridge law in (16.1), take

$$\mathbf{\Phi}(i,u) = p_{0,t_1}(x_i,u), \qquad \mathbf{W}_r(u,v) = p_{t_r,t_{r+1}}(u,v) \qquad \text{and} \qquad \mathbf{\Psi}(v,j) = p_{t_m,t^*}(v,z_j).$$

The Chapman–Kolmogorov identity gives

$$\mathbf{M}_{ij} = p_{0,t^*}(x_i,z_j),$$

so $\det \mathbf{M} = q_{0,t^*}(\boldsymbol{x},\boldsymbol{z})$, exactly the denominator in (16.1). The spacetime correlation kernel is therefore

$$\begin{aligned} K\big((r,u),(s,v)\big) = \sum_{a,b=1}^{n} & p_{t_r,t^*}(u,z_a)(\mathbf{M}^{-1})_{a,b}p_{0,t_s}(x_b,v) \\ & - \mathbb{1}_{\{r<s\}}p_{t_r,t_s}(u,v). \end{aligned}$$

This formula expresses every spacetime inclusion probability as a determinant of one-particle transition probabilities and the inverse of the endpoint matrix $\mathbf{M}$. It is the determinantal description needed for asymptotic analysis of nonintersecting walks.

Lecture 18 — Wednesday, March 24, 2021

Determinantal Martingales

Definition 18.1 Let $(\mathcal{F}_t)_{t\geq 0}$ be a filtration. An integrable adapted process $(M_t)_{t\geq 0}$ is a **martingale** if

$$\mathbb{E}[M_{t+s} \mid \mathcal{F}_t] = M_t$$

for all $s, t \geq 0$.

Linear combinations of martingales are martingales. In addition, if M and N are independent martingales with respect to their joint filtration, then MN is a martingale. The same argument applies coordinatewise to vector-valued martingales and to their tensor products.

Proposition 18.2 For $i \in \{1, \ldots, k\}$, let

$$\boldsymbol{M}_t^{(i)} = \left(M_{t,1}^{(i)}, \ldots, M_{t,k}^{(i)}\right)$$

be an $\mathbb{R}^k$-valued martingale. If the k vector-valued martingales are independent, then

$$D_t \coloneqq \det\left[M_{t,j}^{(i)}\right]_{i,j=1}^k$$

is a martingale.

Proof. The Leibniz formula for determinants gives

$$D_t = \sum_{\sigma \in \mathfrak{S}_k} \operatorname{sgn}(\sigma) \prod_{i=1}^k M_{t,\sigma(i)}^{(i)}.$$

For each permutation σ, the factors in the product come from independent vector-valued martingales, so their product is a martingale. The finite signed sum is therefore a martingale. $\square$

Remark 18.3 The same result can be expressed through exterior algebra. The antisymmetrized tensor

$$\boldsymbol{M}_t^{(1)} \wedge \cdots \wedge \boldsymbol{M}_t^{(k)} := \frac{1}{k!} \sum_{\sigma \in \mathfrak{S}_k} \operatorname{sgn}(\sigma) \boldsymbol{M}_t^{(\sigma(1))} \otimes \cdots \otimes \boldsymbol{M}_t^{(\sigma(k))}$$

is a vector-valued martingale. Its coefficient in the direction $\boldsymbol{e}_1 \wedge \cdots \wedge \boldsymbol{e}_k$ is a constant multiple of D_t.

Polynomial Martingales

Let

$$X(t) = \xi_1 + \cdots + \xi_t$$

be a random walk whose increments are independent and identically distributed, have mean zero, and have a moment generating function in a neighbourhood of the origin. Define the **cumulant generating function**

$$\kappa(\lambda) := \log \mathbb{E}[e^{\lambda \xi_1}].$$

For every admissible λ, the exponential process

$$\mathcal{E}_t(\lambda) := \exp\big(\lambda X(t) - t\kappa(\lambda)\big) \tag{18.1}$$

is a martingale. Indeed,

$$\mathcal{E}_t(\lambda) = \prod_{r=1}^{t} \exp\big(\lambda \xi_r - \kappa(\lambda)\big),$$

and the factors are independent with mean one.

Differentiating (18.1) with respect to λ produces polynomial martingales. More precisely, define

$$\mathcal{C}_r(x,t) := \frac{\partial^r}{\partial \lambda^r} \exp\big(\lambda x - t\kappa(\lambda)\big)\bigg|_{\lambda=0}. \tag{18.2}$$

Differentiation under conditional expectation is justified by the exponential moment bound. Hence $\mathcal{C}_r(X(t),t)$ is a martingale for every $r \geq 0$. The first three polynomials are

$$\mathcal{C}_0(x,t) = 1, \qquad \mathcal{C}_1(x,t) = x - t\kappa'(0) = x,$$
$$\text{and} \qquad \mathcal{C}_2(x,t) = \big(x - t\kappa'(0)\big)^2 - t\kappa''(0) = x^2 - t\kappa''(0).$$

For every r, the polynomial $\mathcal{C}_r(\,\cdot\,,t)$ is monic of degree r.

Let $X_1, \dots, X_k$ be independent copies of X, and write $\boldsymbol{X}(t) = (X_1(t), \dots, X_k(t))$. By Proposition 18.2,

$$\det\big[\mathcal{C}_{i-1}(X_j(t), t)\big]_{i,j=1}^{k}$$

is a martingale. Elementary row operations replace the monic polynomial $\mathcal{C}_{i-1}(\,\cdot, t)$ by the monomial of the same degree without changing the determinant. Consequently,

$$\Delta(\boldsymbol{X}(t)) \coloneqq \prod_{1\le i<j\le k} \big(X_j(t) - X_i(t)\big) \tag{18.3}$$

is a martingale. In particular, Δ is harmonic for the k-particle transition kernel:

$$\mathbb{E}_{\boldsymbol{x}}[\Delta(\boldsymbol{X}(1))] = \Delta(\boldsymbol{x}).$$

Endpoint Martingales and Karlin–McGregor

Fix a terminal time T and an ordered terminal vector $\boldsymbol{y} = (y_1 < \cdots < y_k)$. For one walk, the Doob martingale associated with the event $\{X(T) = y_j\}$ is

$$M_j(t) \coloneqq \mathbb{E}\big[\mathbb{1}_{\{X(T)=y_j\}} \mid \mathcal{F}_t\big] = p_{T-t}(X(t), y_j).$$

For independent walks $X_1, \dots, X_k$, set

$$\mathbf{M}_{ij}(t) \coloneqq p_{T-t}(X_i(t), y_j).$$

The determinant $D_t = \det[\mathbf{M}_{ij}(t)]_{i,j=1}^{k}$ is a martingale by Proposition 18.2. Let

$$\tau \coloneqq \inf\{t \ge 0 : X_i(t) = X_j(t) \text{ for some } i \ne j\}.$$

If $\tau \le T$, two rows of $\mathbf{M}(\tau)$ coincide, so $D_\tau = 0$. Optional stopping at $T \wedge \tau$ therefore gives

$$\begin{aligned}\det\big[p_T(x_i, y_j)\big]_{i,j=1}^{k} &= D_0 \\ &= \mathbb{E}_{\boldsymbol{x}}\left[\sum_{\sigma\in\mathfrak{S}_k} \operatorname{sgn}(\sigma) \prod_{i=1}^{k} \mathbb{1}_{\{X_i(T)=y_{\sigma(i)}\}};\ \tau > T\right].\end{aligned}$$

If the walks cannot change their order without meeting, as is the case for Bernoulli walks, ordered initial and terminal positions permit only the identity permutation on $\{\tau > T\}$. Thus

$$\det\big[p_T(x_i, y_j)\big]_{i,j=1}^{k} = \mathbb{P}_{\boldsymbol{x}}\big(\boldsymbol{X}(T) = \boldsymbol{y},\ \tau > T\big),$$

which recovers the Karlin–McGregor formula.

The Vandermonde Doob Transform

For the remainder of this discussion, suppose that the walks cannot change their order without meeting, as is the case for Bernoulli walks. This assumption identifies the Karlin–McGregor determinant with the transition kernel of the process killed at the boundary of the ordered chamber.

Let

$$\mathbb{W}_k := \{\boldsymbol{x} \in \mathbb{R}^k : x_1 < \cdots < x_k\}$$

be the type A Weyl chamber. Write

$$\mathbb{P}_t^*(\boldsymbol{x}, \boldsymbol{y}) := \det\big[p_t(x_i, y_j)\big]_{i,j=1}^k$$

for the transition kernel of the independent walks killed upon leaving $\mathbb{W}_k$. Since Δ is positive and harmonic on $\mathbb{W}_k$, its Doob transform is

$$q_t(\boldsymbol{x}, \boldsymbol{y}) := \frac{\Delta(\boldsymbol{y})}{\Delta(\boldsymbol{x})} \mathbb{P}_t^*(\boldsymbol{x}, \boldsymbol{y}). \tag{18.4}$$

The harmonicity identity shows that $\sum_{\boldsymbol{y}} q_t(\boldsymbol{x}, \boldsymbol{y}) = 1$. Hence (18.4) is the transition kernel of the nonintersecting walk ensemble obtained by conditioning the particles to remain ordered.

At observation times $0 < t_1 < \cdots < t_m$, the resulting line ensemble has the geometry in Figure 84. The product of transformed transition probabilities telescopes:

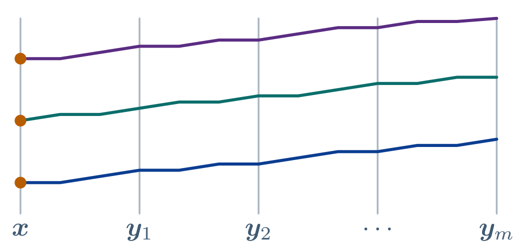

Figure 84

$$\begin{aligned}
\mathbb{P}_{\boldsymbol{x}}\big(\boldsymbol{X}(t_1) = \boldsymbol{y}_1, \ldots, \boldsymbol{X}(t_m) = \boldsymbol{y}_m\big) &= \prod_{r=1}^m q_{t_r - t_{r-1}}(\boldsymbol{y}_{r-1}, \boldsymbol{y}_r) \\
&= \frac{\Delta(\boldsymbol{y}_m)}{\Delta(\boldsymbol{x})} \prod_{r=1}^m \det\big[p_{t_r - t_{r-1}}(y_{r-1,i}, y_{r,j})\big]_{i,j=1}^k,
\end{aligned}$$

where $t_0 = 0$ and $\boldsymbol{y}_0 = \boldsymbol{x}$. This has the Eynard–Mehta form. Define the Vandermonde matrix

$$\mathbf{V}(\boldsymbol{x}) \coloneqq \big[\mathcal{C}_{j-1}(x_i, 0)\big]_{i,j=1}^{k} = \big[x_i^{j-1}\big]_{i,j=1}^{k}$$

and choose

$$\begin{aligned}\boldsymbol{\Phi}(i,u) &\coloneqq \sum_{a=1}^{k} \big[\mathbf{V}(\boldsymbol{x})^{-1}\big]_{i,a} p_{0,t_1}(x_a, u),\\ \mathbf{W}_r(u,v) &\coloneqq p_{t_r,t_{r+1}}(u,v),\\ \boldsymbol{\Psi}(u,j) &\coloneqq \mathcal{C}_{j-1}(u, t_m).\end{aligned}$$

The determinant of $\boldsymbol{\Phi}$ contributes $\Delta(\boldsymbol{x})^{-1}$, while the determinant of $\boldsymbol{\Psi}$ contributes $\Delta(\boldsymbol{y}_m)$. Moreover, the martingale property of $\mathcal{C}_{j-1}(X(t), t)$ implies

$$\sum_{v} p_{s,t}(u,v)\mathcal{C}_{j-1}(v,t) = \mathcal{C}_{j-1}(u,s).$$

Applying this identity through all transition matrices gives

$$\mathbf{M} = \boldsymbol{\Phi}\mathbf{W}_1 \cdots \mathbf{W}_{m-1}\boldsymbol{\Psi} = \mathbf{V}(\boldsymbol{x})^{-1}\mathbf{V}(\boldsymbol{x}) = \mathbf{I}.$$

Substituting $\mathbf{M}^{-1} = \mathbf{I}$ into (16.4) yields the explicit correlation kernel

$$\begin{aligned}K\big((r,u),(s,v)\big) = \sum_{a,b=1}^{k} &\mathcal{C}_{a-1}(u,t_r)\big[\mathbf{V}(\boldsymbol{x})^{-1}\big]_{a,b} p_{0,t_s}(x_b, v)\\ &- \mathbb{1}_{\{r<s\}} p_{t_r,t_s}(u,v).\end{aligned}$$

Thus polynomial martingales provide the basis that diagonalizes the finite matrix $\mathbf{M}$ in the Eynard–Mehta formula.

Lecture 19 — Friday, March 26, 2021

Random and Ballistic Deposition

In the simplest random deposition model, the columns are indexed by $\mathbb{Z}$, and unit blocks arrive independently at each column according to a Poisson process of rate λ. If $H_x(t)$ denotes the height of column x at time t, then

$$H_x(t) \sim \text{Poisson}(\lambda t).$$

The central limit theorem therefore gives

$$\frac{H_x(t) - \lambda t}{\sqrt{\lambda t}} \xrightarrow{\mathrm{d}} \mathcal{N}(0,1).$$

In ballistic deposition, a falling block sticks as soon as it touches the existing aggregate, rather than falling independently to the base of its chosen column. This interaction creates spatial correlations. For standard one-dimensional ballistic deposition, numerical evidence suggests Kardar–Parisi–Zhang fluctuations of the form

$$c_2 t^{-1/3}\big(H_0(t) - c_1 t\big) \xrightarrow{\mathrm{d}} \mathcal{T}_1$$

for suitable constants $c_1, c_2 > 0$, but this limiting statement is not known from the elementary deposition description alone. Here $\mathcal{T}_1$ is the Tracy–Widom law associated with the Gaussian orthogonal ensemble. The following exactly solvable anisotropic model instead produces $\mathcal{T}_2$, the law associated with the Gaussian unitary ensemble.

An Anisotropic Deposition Model

The following exactly solvable anisotropic model was introduced by Majumdar and Nechaev. Consider L columns with free boundary conditions. Time is discrete, and at each time exactly one column is chosen uniformly for deposition. Let $H_k(t)$ be the height of column k after t depositions. If column k is selected, the update is

$$H_k(t+1) = \max_{1 \le j \le k} H_j(t) + 1. \tag{19.1}$$

All other column heights remain unchanged. Thus a block deposited in column k screens every site at the same level in columns to its right: future blocks in those columns must land above that level. A possible configuration is shown in Figure 85. The first column evolves independently of

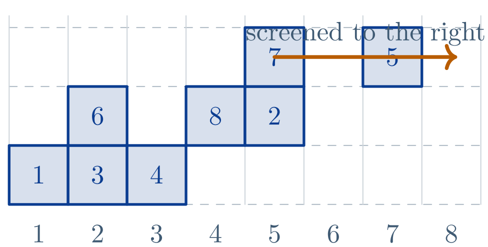

Figure 85

the others. Since it is selected with probability $1/L$ at each step,

$$\mathbb{P}(H_1(t) = r) = \binom{t}{r} \left(\frac{1}{L}\right)^r \left(1 - \frac{1}{L}\right)^{t-r}.$$

For $x \in \{0, \ldots, L-1\}$, measure height relative to the first column:

$$h_x(t) \coloneqq H_{x+1}(t) - H_1(t).$$

The goal is to determine the asymptotic law of $h_x(t)$. The key is an exact correspondence with longest increasing subsequences.

The Ulam Problem

For a permutation $\pi \in \mathfrak{S}_N$, let $\ell_N(\pi)$ be the length of its longest increasing subsequence. The entries of the subsequence need not be consecutive. For example,

$$\ell_4(4, 1, 2, 3) = 3.$$

Under the uniform distribution on $\mathfrak{S}_N$, write ℓ_N for this random length. Ulam conjectured a square-root growth law. Vershik and Kerov identified the constant:

$$\lim_{N\to\infty} \frac{\mathbb{E}[\ell_N]}{\sqrt{N}} = 2.$$

Baik, Deift, and Johansson determined the fluctuations.

Theorem 19.1 (Baik–Deift–Johansson theorem) For a uniformly random permutation of $\{1, \ldots, N\}$,

$$\frac{\ell_N - 2\sqrt{N}}{N^{1/6}} \xrightarrow{\mathrm{d}} \mathcal{T}_2.$$

The same limit holds in the Poissonized model: if N_λ is an independent Poisson random variable with mean λ, then

$$\frac{\ell_{N_\lambda} - 2\sqrt{\lambda}}{\lambda^{1/6}} \xrightarrow{\mathrm{d}} \mathcal{T}_2$$

as $\lambda \to \infty$.

Patience Sorting

The patience sorting algorithm extracts $\ell_N(\pi)$ without examining all subsequences. Process the entries of π from left to right. Place each entry on the leftmost pile whose top entry is larger; if no such pile exists, begin a new pile.

Example 19.2 For

$$\pi = (8, 3, 5, 1, 2, 6, 4, 7),$$

the successive piles are

$$\begin{array}{cccc} 1 & 2 & 4 & 7 \\ 3 & 5 & 6 & \\ 8 & & & \end{array}$$

when each pile is displayed from top to bottom. There are four piles, and the longest increasing subsequences have length four; examples include $(3, 5, 6, 7)$ and $(1, 2, 4, 7)$.

Proposition 19.3 The number of piles created by patience sorting is $\ell_N(\pi)$.

Proof. Let r be the number of piles. Every increasing subsequence uses at most one entry from each pile and moves strictly to the right through the piles. Hence $\ell_N(\pi) \le r$.

For the reverse inequality, whenever an entry is placed on pile $j > 1$, point from it to the top entry of pile $j-1$ that was present immediately before its placement. That earlier entry is smaller. Starting from the entry that first creates pile r and following pointers backward produces entries from piles $r, r-1, \ldots, 1$. Reading this chain in chronological order gives an increasing subsequence of length r. Thus $\ell_N(\pi) \ge r$, proving equality. $\square$

The Deposition–Subsequence Correspondence

Record the deposition time inside every occupied block. Project the level lines of the heap onto the column-time plane. Reading the event labels column by column from left to right and, within a column, from top to bottom produces a permutation of the event times. Under this encoding, the piles in patience sorting are in one-to-one correspondence with the projected level lines. The correspondence is summarized in Figure 86. For the rectangle

$$R_{x,t} = [0, x] \times [0, t],$$

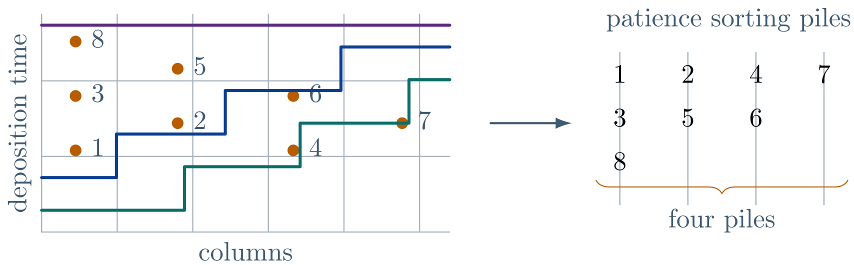

Figure 86

the **relative height** $h_x(t)$ is the number of level lines crossing $R_{x,t}$. By Proposition 19.3, this equals the length of the longest increasing subsequence formed by the deposition events inside the rectangle. If $N_{x,t}$ denotes the number of such events, then

$$h_x(t) = \ell_{N_{x,t}} \qquad \text{and} \qquad N_{x,t} \sim \mathrm{Bin}\left(t, \frac{x}{L}\right). \tag{19.2}$$

Indeed, each of the t independently selected columns lies among the first x columns with probability x/L.

Theorem 19.4 Consider the iterated regime in which $L, t \to \infty$ with

$$\frac{tx}{L} \longrightarrow \lambda$$

and then $\lambda \to \infty$. The relative height satisfies

$$\frac{h_x(t) - 2\sqrt{tx/L}}{(tx/L)^{1/6}} \xrightarrow{\mathrm{d}} \mathcal{T}_2.$$

Proof. In the first limit,

$$N_{x,t} \xrightarrow{\mathrm{d}} \mathrm{Poisson}(\lambda)$$

by the Poisson limit for binomial random variables. Conditional on $N_{x,t} = n$, the encoding in (19.2) identifies the height with the longest increasing subsequence length of the induced random permutation. The Poissonized statement in Theorem 19.1, together with (19.2), gives

$$h_x(t) = 2\sqrt{\lambda} + \lambda^{1/6}\chi + o_{\mathbb{P}}(\lambda^{1/6}) \qquad \text{and} \qquad \chi \sim \mathcal{T}_2.$$

Substituting $\lambda = tx/L$ proves the result. $\square$

Lecture 20 — Wednesday, March 31, 2021

4. Brownian Last-Passage Percolation and the Directed Landscape

Noncolliding Ornstein–Uhlenbeck Processes

Let B be a standard Brownian motion and define

$$U(t) \coloneqq 2^{-1/2}e^{-t/2}B(e^t), \qquad t \in \mathbb{R}.$$

This time change turns Brownian motion into an Ornstein–Uhlenbeck process. Indeed, U is a centered Gaussian process and, when $s \leq t$,

$$\mathbb{E}[U(s)U(t)] = \frac{1}{2}e^{-(s+t)/2}\mathbb{E}[B(e^s)B(e^t)] = \frac{1}{2}e^{-(t-s)/2}.$$

Thus $\mathbb{E}[U(s)U(t)] = \frac{1}{2}e^{-|t-s|/2}$, so U is stationary. Its covariance is unchanged when time is reversed, and hence the process is also reversible. Applying the same transformation to nonintersecting Brownian motions gives stationary, nonintersecting Ornstein–Uhlenbeck processes, as illustrated in Figure 87. For $k \geq 0$, let

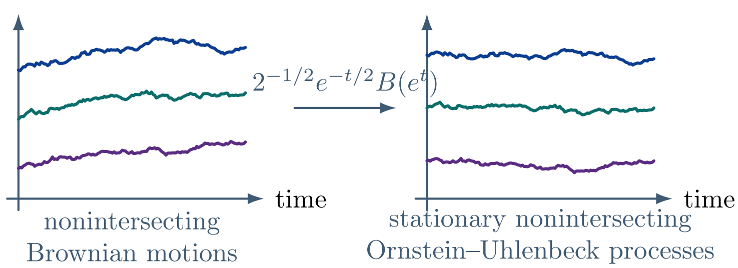

Figure 87

$$\varphi_k(x) \coloneqq \frac{1}{\sqrt{2^k k!\sqrt{\pi}}}H_k(x)e^{-x^2/2},$$

where H_k is the kth Hermite polynomial. These Hermite functions form an orthonormal basis of $L^2(\mathbb{R})$, so

$$\int_{\mathbb{R}} \varphi_k(x)\varphi_\ell(x)\,\mathrm{d}x = \delta_{k\ell}.$$

Theorem 20.1 (Extended Hermite kernel) The spacetime point process of n non-intersecting Ornstein–Uhlenbeck processes is determinantal. Up to a determinantal gauge transformation, its correlation kernel is

$$K_n\big((s,x),(t,y)\big) = \begin{cases} \displaystyle\sum_{k=0}^{n-1} e^{k(s-t)/2}\varphi_k(x)\varphi_k(y), & s \geq t, \\ \displaystyle -\sum_{k=n}^{\infty} e^{k(s-t)/2}\varphi_k(x)\varphi_k(y), & s < t. \end{cases}$$

At a fixed time, the kernel in Theorem 20.1 becomes

$$K_n(x,y) = \sum_{k=0}^{n-1} \varphi_k(x)\varphi_k(y).$$

This is the integral kernel of the orthogonal projection onto $\operatorname{span}\{\varphi_0,\dots,\varphi_{n-1}\}$. After the corresponding variance normalization, it is also the correlation kernel of the eigenvalues of an $n \times n$ Gaussian unitary ensemble matrix.

Proposition 20.2 Let Ξ be a determinantal point process with locally trace-class Hermitian kernel K, and let A be a Borel set on which the restriction K_A is trace class. If $(\mu_j)_{j\geq 1}$ are the eigenvalues of K_A, then

$$\#(\Xi \cap A) \overset{\mathrm{d}}{=} \sum_{j\geq 1} Y_j,$$

where the Y_j are independent Bernoulli random variables with $\mathbb{P}(Y_j = 1) = \mu_j$.

Proof. For $z \in [0,1]$, the generating functional of a determinantal process gives

$$\mathbb{E}\Big[z^{\#(\Xi\cap A)}\Big] = \det(I - (1-z)K_A) = \prod_{j\geq 1}\big(1 - \mu_j + z\mu_j\big).$$

The last expression is the probability generating function of $\sum_{j\geq 1} Y_j$, which proves the claim. □

We now study the upper spectral edge as n tends to infinity. The spacing between the top few curves is of order $n^{-1/6}$, and their nontrivial fluctuations occur on that scale; see Figure 88.

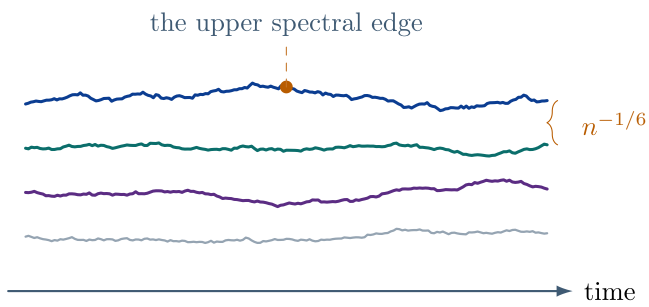

Figure 88

The Hermite-to-Airy Limit

The upper edge of the Gaussian unitary ensemble lies near $\sqrt{2n}$ in the present normalization. Introduce the edge coordinate

$$x = \sqrt{2n} + \frac{\xi}{\sqrt{2}\, n^{1/6}}.$$

The appearance of the Airy function can already be seen at the level of differential operators. Define the **harmonic oscillator operator**

$$D \coloneqq \frac{1}{2}\left(-\frac{\mathrm{d}^2}{\mathrm{d}x^2} + x^2 - 1\right).$$

The Hermite functions satisfy $D\varphi_k = k\varphi_k$. In the edge coordinate ξ, a direct calculation gives

$$n^{-1/3}(D - n) = -\frac{\mathrm{d}^2}{\mathrm{d}\xi^2} + \xi + \frac{\xi^2}{4n^{2/3}} - \frac{1}{2n^{1/3}}.$$

Consequently, the edge-scaled harmonic oscillator formally converges to the **Airy operator**

$$\mathcal{H}_{\mathrm{Ai}} \coloneqq -\frac{\mathrm{d}^2}{\mathrm{d}\xi^2} + \xi.$$

This operator limit explains why edge-scaled Hermite functions converge to translates of the Airy function. A rigorous proof requires uniform asymptotics for the Hermite functions, which may be obtained by contour integration.

Definition 20.3 The **Airy function** Ai is the unique real solution of

$$f''(x) = xf(x)$$

that tends to 0 as x tends to $+\infty$. It admits the oscillatory integral representation

$$\mathrm{Ai}(x) = \frac{1}{\pi}\int_0^\infty \cos\left(\frac{t^3}{3} + xt\right) \mathrm{d}t.$$

Its exponential decay on the positive half-line and oscillation on the negative half-line are shown in Figure 89. For comparison with the stochastic Airy operator, consider $\mathcal{H}_{\mathrm{Ai}}$ on $L^2(0,\infty)$ with

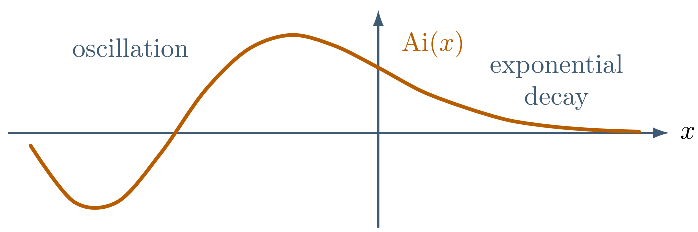

FIGURE 89

the Dirichlet boundary condition $f(0) = 0$. Let

$$a_1 > a_2 > \cdots$$

be the negative zeros of Ai, and set $\lambda_j = -a_j$. Then

$$f_j(x) \coloneqq \mathrm{Ai}(x - \lambda_j) \qquad \text{and} \qquad \mathcal{H}_{\mathrm{Ai}} f_j = \lambda_j f_j.$$

Indeed, the Airy differential equation verifies the eigenvalue equation, and $f_j(0) = \mathrm{Ai}(a_j) = 0$. This is the zero-noise case of the stochastic Airy operator, in which the Brownian white noise potential has been removed.

After the corresponding spacetime edge scaling, the extended Hermite kernel in Theorem 20.1 converges to the extended Airy kernel

$$K_{\mathrm{Ai}}\big((x,s),(y,t)\big) = \begin{cases} \displaystyle\int_0^\infty e^{-\lambda(s-t)}\,\mathrm{Ai}(x+\lambda)\,\mathrm{Ai}(y+\lambda)\,\mathrm{d}\lambda, & s \ge t, \\ \displaystyle -\int_{-\infty}^0 e^{-\lambda(s-t)}\,\mathrm{Ai}(x+\lambda)\,\mathrm{Ai}(y+\lambda)\,\mathrm{d}\lambda, & s < t. \end{cases}$$

This is precisely the kernel used to define the stationary Airy line ensemble $\mathcal{A}$ in Definition 15.1. The first few curves are represented in Figure 90. The stationary ensemble $\mathcal{A}$ and its parabolically shifted version

$$\mathcal{P}_i(t) \coloneqq \mathcal{A}_i(t) - t^2$$

must be distinguished. The parabolic correction arises because centring at a tangent line removes

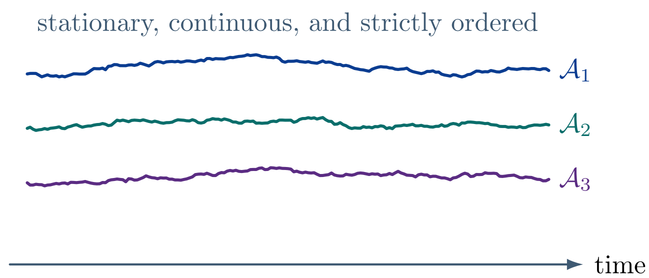

Figure 90

the constant and linear terms of the macroscopic edge profile, while its quadratic curvature remains. This geometry is shown in Figure 91. At every fixed time, the values of the Airy curves form the

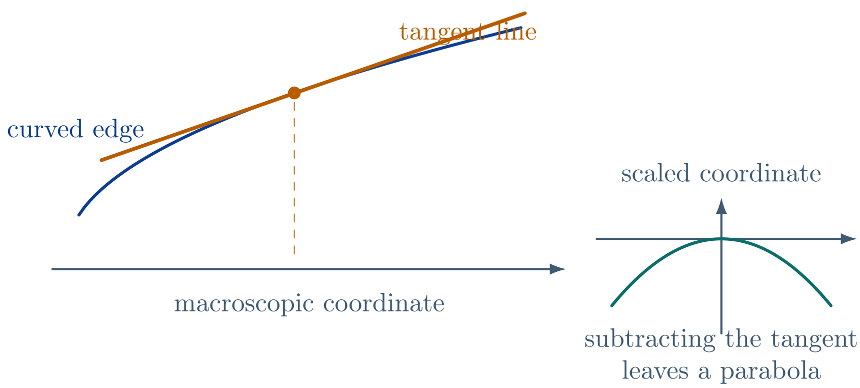

Figure 91

Airy point process. A representative slice is depicted in Figure 92. The lower points have the classical asymptotic behaviour

$$\frac{\mathcal{A}_k(0)}{-\left(\frac{3\pi k}{2}\right)^{2/3}} \longrightarrow 1 \qquad \text{almost surely as } k \longrightarrow \infty.$$

Thus the magnitude of the kth point grows on the $k^{2/3}$ scale. Three complementary structures characterize the ensemble:

1. The Brownian normalization $2^{-1/2}(\mathcal{A}_i(t) - t^2)$ has the Brownian Gibbs property, as stated in Theorem 15.3.
2. At a fixed time, the point process is determinantal with the projection kernel

$$K_{\mathrm{Ai}}(x, y) = \int_0^\infty \mathrm{Ai}(x + \lambda)\,\mathrm{Ai}(y + \lambda)\,\mathrm{d}\lambda.$$

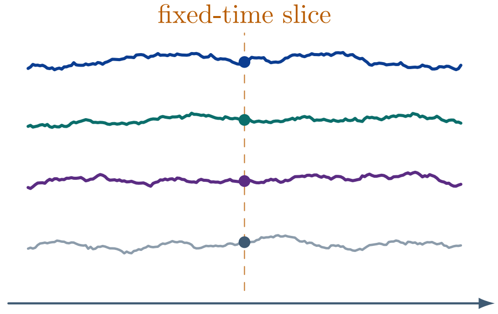

FIGURE 92

3. The top curve has the Fredholm determinant distribution

$$\mathbb{P}\big(\mathcal{A}_1(0) \le r\big) = \det(I - K_{\mathrm{Ai}})_{L^2(r,\infty)}\,.$$

To verify (3), restrict the fixed-time kernel in (2) to $L^2(r,\infty)$, and denote its eigenvalues by $(\mu_j)_{j\ge 1}$. Proposition 20.2 gives independent Bernoulli random variables Y_j with parameters μ_j such that

$$\#\{i : \mathcal{A}_i(0) > r\} \overset{\mathrm{d}}{=} \sum_{j\ge 1} Y_j.$$

It follows that

$$\mathbb{P}\big(\mathcal{A}_1(0) \le r\big) = \mathbb{P}(Y_j = 0 \text{ for all } j \ge 1) = \prod_{j\ge 1}(1-\mu_j) = \det(I - K_{\mathrm{Ai}})_{L^2(r,\infty)}\,.$$

Brownian Last-Passage Percolation

Let $B_1, \dots, B_n$ be independent two-sided standard Brownian motions. For $x \le y$, the Brownian last-passage value from (x,n) to $(y,1)$ is

$$d_B\big((x,n),(y,1)\big) = \sup_{x=t_n\le\cdots\le t_0=y} \sum_{i=1}^{n}\big(B_i(t_{i-1}) - B_i(t_i)\big).$$

The staircase path records which Brownian motion contributes to the weight at each time, as shown in Figure 93. For $k \in \{1,\dots,n\}$, let

$$G_{n,k}(t) := d_B\big((0,n)^k,(t,1)^k\big) \qquad \text{and} \qquad G_{n,0}(t) := 0,$$

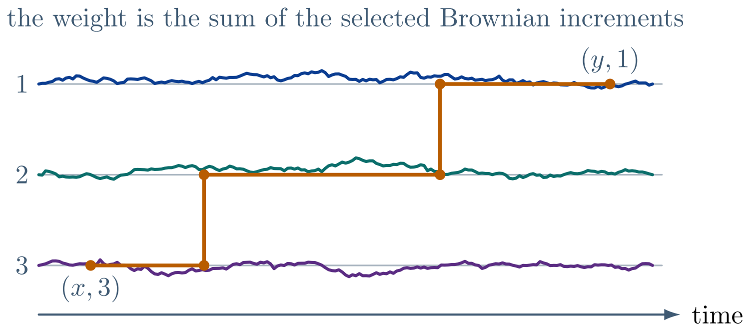

FIGURE 93

where $G_{n,k}(t)$ is the maximum total weight of k mutually disjoint paths. Define the increments

$$\Delta_{n,k}(t) \coloneqq G_{n,k}(t) - G_{n,k-1}(t).$$

By the Greene identity (3), the increments are precisely the lines of the Brownian melon:

$$\Delta_{n,k}(t) = (WB)_k(t).$$

The deterministic shape at time t is $2\sqrt{nt}$. To observe the upper edge near $t = 1$, write

$$t = 1 + \frac{2u}{n^{1/3}}.$$

Taylor expansion of the shape function gives

$$2\sqrt{nt} = 2\sqrt{n} + 2un^{1/6} - u^2 n^{-1/6} + O(n^{-1/2}).$$

The edge limit of Brownian last-passage percolation is therefore

$$\Delta_{n,k}\left(1 + \frac{2u}{n^{1/3}}\right) = 2\sqrt{n} + 2un^{1/6} + n^{-1/6}\left(\mathcal{A}_k(u) - u^2 + o_n(1)\right). \tag{20.1}$$

The convergence in (20.1) holds jointly for every fixed number of lines and, after establishing tightness, locally uniformly in u. Thus subtracting only the tangent line terms produces the parabolically shifted Airy ensemble, whereas subtracting the complete shape produces the stationary Airy ensemble. This limit is the one-dimensional boundary of the two-parameter directed landscape constructed next.

Lecture 21 — Wednesday, April 07, 2021

Directed Landscape

The **directed landscape** is the universal continuum limit of directed first-passage and last-passage metrics in the Kardar–Parisi–Zhang universality class. It is a random directed metric on spacetime, with spatial coordinate x and time coordinate s.

To describe the general scaling, fix a macroscopic direction $\boldsymbol{v}$ and a transverse direction $\boldsymbol{u}$. For a directed metric d, one expects

$$\begin{aligned} &d\left(\boldsymbol{v}\sigma^3 s+\boldsymbol{u}\sigma^2 x,\ \boldsymbol{v}\sigma^3 t+\boldsymbol{u}\sigma^2 y\right) \\ &\quad=\alpha\sigma^3(t-s)+\beta\sigma^2(y-x)+\chi\sigma\left(\mathcal{L}(x,s;y,t)+o_\sigma(1)\right). \end{aligned} \tag{21.1}$$

The first two terms are deterministic shape terms. The longitudinal, transverse, and fluctuation scales are respectively σ^3, σ^2, and σ. The constants α,β,χ, as well as the choice of $\boldsymbol{u}$, depend on the microscopic model and the direction $\boldsymbol{v}$. The error in (21.1) tends to 0 in probability after division by $\chi\sigma$. The longitudinal displacement and transverse perturbation are illustrated in Figure 94. For Brownian last-passage percolation in direction $\boldsymbol{v}=(\rho,-1)$, take $\boldsymbol{u}=(\tau,0)$. If

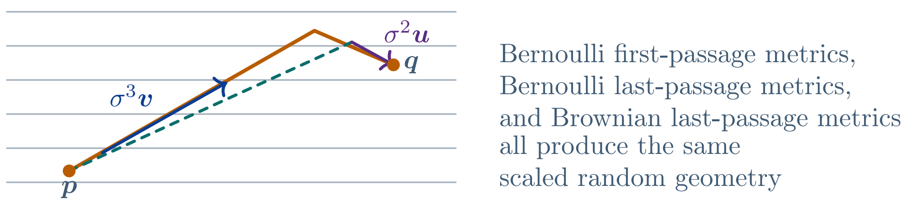

FIGURE 94

$\Delta=t-s$ and $\delta=y-x$, expansion of the deterministic shape $2\sqrt{ab}$ gives

$$\begin{aligned} 2\sqrt{\left(\rho\sigma^3\Delta+\tau\sigma^2\delta\right)\sigma^3\Delta}=2\sqrt{\rho}\,\sigma^3\Delta+\frac{\tau}{\sqrt{\rho}}\,\sigma^2\delta \\ -\frac{\tau^2}{4\rho^{3/2}}\,\sigma\frac{\delta^2}{\Delta}+O(1). \end{aligned}$$

The fluctuation coefficient is $\chi=\sqrt{\rho}$. Matching the quadratic coefficient to the standard parabola requires

$$\frac{\chi}{\tau^2}=\frac{1}{4\rho^{3/2}},$$

and hence $\tau = 2\rho$. Thus we may take

$$\alpha = 2\sqrt{\rho}, \qquad \beta = \frac{\tau}{\sqrt{\rho}} = 2\sqrt{\rho} \qquad \text{and} \qquad \chi = \sqrt{\rho}.$$

For a single pair of spacetime points, the directed landscape has the distribution

$$\mathcal{L}(x,s;y,t) \overset{\mathrm{d}}{=} \begin{cases} -\dfrac{(x-y)^2}{t-s} + (t-s)^{1/3}\,\mathcal{T}_2, & s<t, \\ 0, & (x,s)=(y,t), \\ -\infty, & s>t, \text{ or } s=t \text{ and } x \neq y, \end{cases} \tag{21.2}$$

where $\mathcal{T}_2$ has the parameter-2 Tracy–Widom distribution.

Set

$$\mathbb{R}^4_{\uparrow} \coloneqq \{(x,s;y,t) \in \mathbb{R}^4 : s<t\}.$$

The directed landscape is characterized by the following four properties.

1. The directed landscape has independent increments: its restrictions to disjoint time intervals are independent. More precisely, if $I_1, \dots, I_k$ are pairwise disjoint intervals, then the random functions formed from pairs whose two time coordinates lie in $I_1, \dots, I_k$, respectively, are independent.
2. The directed landscape is continuous in the sense that the function $\mathcal{L} \colon \mathbb{R}^4_{\uparrow} \to \mathbb{R}$ is almost surely continuous.
3. The directed landscape satisfies metric composition: for $r<t<s$,
$$\mathcal{L}(x,r;z,s) = \max_{y\in\mathbb{R}} \{\mathcal{L}(x,r;y,t) + \mathcal{L}(y,t;z,s)\}.$$
4. The directed landscape has Airy sheet increments: for $a>0$, the two-parameter function $(x,y) \mapsto \mathcal{L}(x,t;y,t+a^3)$ has the law of an Airy sheet at scale a:
$$\mathcal{S}_a(x,y) \overset{\mathrm{d}}{=} a\,\mathcal{S}_1\left(\frac{x}{a^2}, \frac{y}{a^2}\right).$$

The independent time slabs in (1) and the composition rule in (3) are depicted in Figure 95.

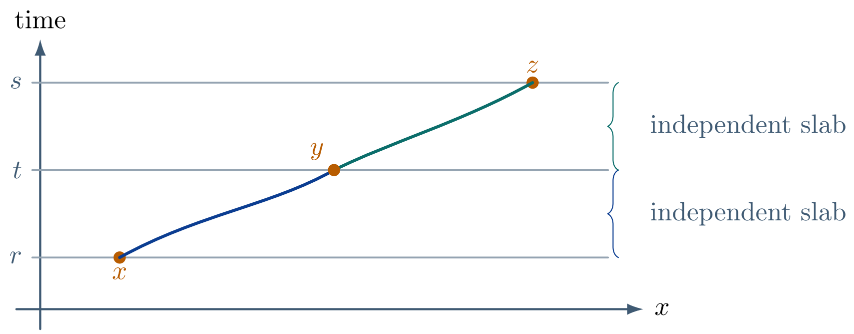

Figure 95

Theorem 21.1 There is a unique law on continuous functions $\mathbb{R}^4_\uparrow \to \mathbb{R}$ satisfying (1)–(4).

The Airy sheet scaling and Theorem 21.1 imply the exact **1:2:3 scale invariance**

$$\left\{\mathcal{L}(a^2x, a^3s; a^2y, a^3t)\right\} \overset{\mathrm{d}}{=} \{a\,\mathcal{L}(x,s;y,t)\}\,, \qquad a>0, \tag{21.3}$$

as random functions indexed by $\mathbb{R}^4_\uparrow$.

Geodesics and Multiplicity

For each deterministic $p=(x,s)$ and $q=(y,t)$ with $s<t$, a geodesic from p to q exists almost surely and is almost surely unique. The null event can depend on the pair (p,q); consequently, random exceptional endpoints can admit several geodesics.

Theorem 21.2 Almost surely, every directed pair p,q is connected by at least one and at most nine geodesics.

Theorem 21.3 Almost surely,

$$\{\#\{\text{geodesics from } (x,0) \text{ to } (y,1)\} : x,y\in\mathbb{R}\} = \{1,2,3,4\}.$$

Representative branching topologies for the multiplicities in Theorem 21.3 are shown in Figure 96. The informal names distinguish whether splitting occurs near the lower endpoint, near the upper endpoint, or in both portions of the network.

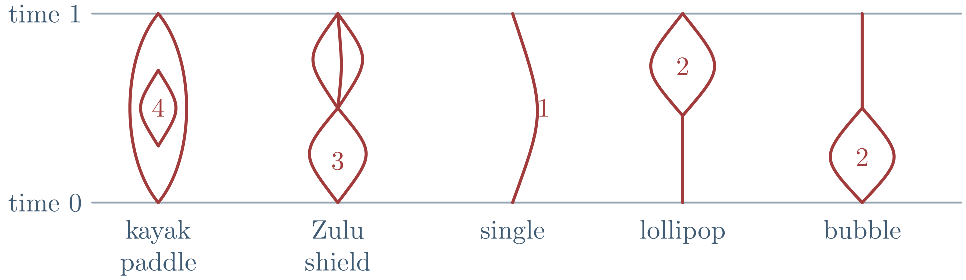

FIGURE 96

Definition 21.4 A **directed path** from (x, s) to (y, t), where $s < t$, is a continuous function $\gamma\colon [s, t] \to \mathbb{R}$ satisfying $\gamma(s) = x$ and $\gamma(t) = y$. It is identified with its spacetime graph

$$\{(\gamma(r), r) : r \in [s, t]\}.$$

The **landscape length** of γ is

$$\|\gamma\|_{\mathcal{L}} \coloneqq \inf_{\substack{k \geq 1 \\ s = t_0 < t_1 < \cdots < t_k = t}} \sum_{i=1}^{k} \mathcal{L}\big(\gamma(t_{i-1}), t_{i-1}; \gamma(t_i), t_i\big).$$

This value belongs to $\mathbb{R} \cup \{-\infty\}$. The path is a **geodesic** if

$$\|\gamma\|_{\mathcal{L}} = \mathcal{L}\big(\gamma(s), s; \gamma(t), t\big).$$

The infimum over partitions is natural because metric composition shows that every partition sum is bounded above by the landscape value between the endpoints. Refining a partition can only decrease this upper bound. Most deterministic paths have landscape length $-\infty$, even though the directed landscape is finite between every pair of strictly time-ordered points.

Example 21.5 Consider the path $\gamma(r) = 0$ from $(0,0)$ to $(0,1)$, partitioned at $t_i = i/n$. The construction is shown in Figure 97.

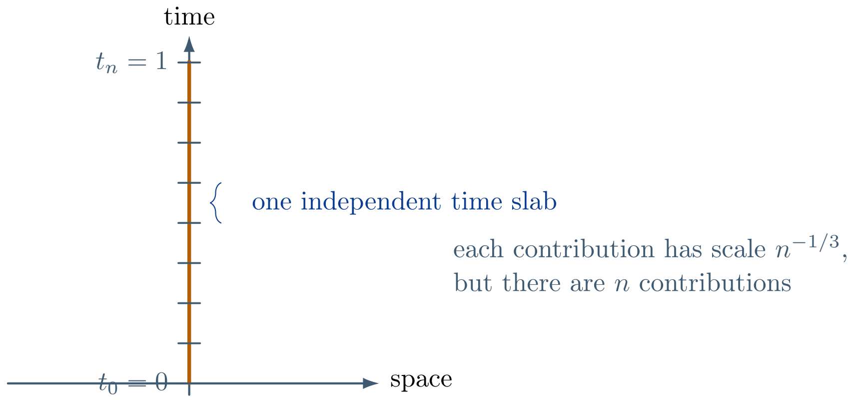

FIGURE 97

By (21.2) and independence across disjoint time slabs,

$$L_n := \sum_{i=1}^{n} \mathcal{L}\left(0, \frac{i-1}{n}; 0, \frac{i}{n}\right) = n^{-1/3} \sum_{i=1}^{n} Y_i,$$

where $Y_1, Y_2, \ldots$ are independent, identically distributed type 2 Tracy–Widom random variables. Their common mean $m = \mathbb{E}[Y_1]$ is negative. The strong law of large numbers gives

$$\frac{1}{n} \sum_{i=1}^{n} Y_i \longrightarrow m \qquad \text{almost surely.}$$

Therefore

$$L_n = n^{2/3} \left(\frac{1}{n} \sum_{i=1}^{n} Y_i \right) \longrightarrow -\infty \qquad \text{almost surely.}$$

Since the landscape length is the infimum over all partition sums, the vertical path has length $-\infty$.

Geodesic Trees

Fix $p = (0, 0)$ and take the union of all geodesics from p to points on the time-one line. Geodesics coalesce for part of their journey and then branch toward their terminal points, producing a random planar tree. This geometry is illustrated in Figure 98. The branch points are countable. One way to see why countability is plausible is to encode a branch point by rational times immediately below and above it together with two rational intervals that separate its outgoing branches. Planar ordering also implies that distinct branches cannot cross and later separate again. Both statements require genuine geometric arguments; they do not follow from path continuity alone.

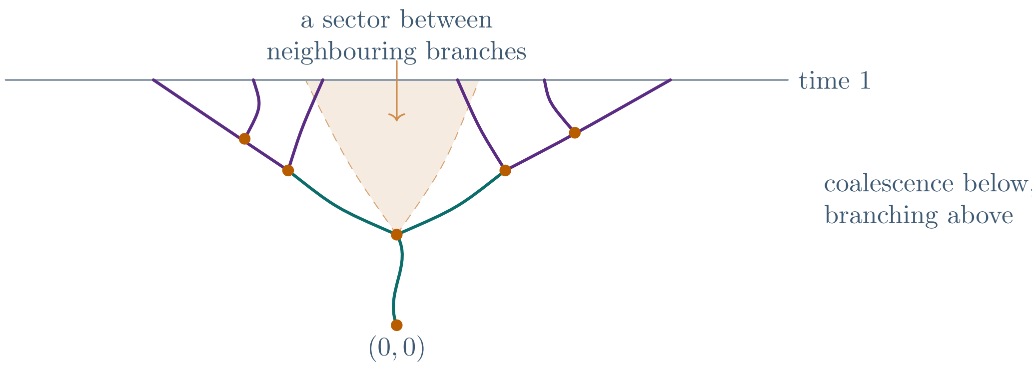

Figure 98

At a branch point, two geodesics may share an initial segment and then separate, forming the lollipop topology from Figure 96. This does not contradict uniqueness for deterministic endpoints, because the branch point and the associated endpoint are random. The contrast with an ordinary pair of ordered geodesics is summarized in Figure 99. In particular, any fixed deterministic spacetime point almost surely does not lie on the exceptional branch locus.

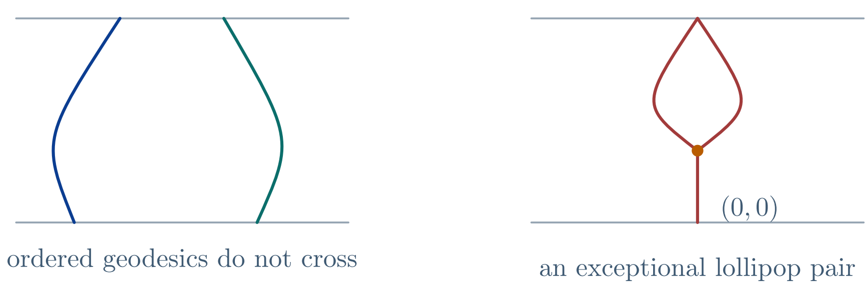

Figure 99

Theorem 21.6 Almost surely, the set

$$\{x \in \mathbb{R} : \text{there are at least two geodesics from } (0,0) \text{ to } (x,1)\}$$

is exactly the set of local minima of $\mathcal{A}_1 - \mathcal{A}_2$.

Remark 21.7 Under the stationary convention of Definition 15.1,

$$\mathcal{L}(0,0;x,1) = \mathcal{A}_1(x) - x^2 = \mathcal{P}_1(x).$$

The same parabola is subtracted from every line, so $\mathcal{P}_1 - \mathcal{P}_2 = \mathcal{A}_1 - \mathcal{A}_2$. The second line has a more involved multipath interpretation in terms of the directed landscape.

Airy Sheet

Definition 21.8 The **Airy sheet** is the two-parameter random function obtained by taking last-passage values from infinity through the parabolically shifted Airy line ensemble and fixing the remaining additive constant by stationarity.

The functional Robinson–Schensted–Knuth correspondence converts the line ensemble last-passage problem into an equivalent transformed line ensemble. This viewpoint is shown schematically in Figure 100. The normalization is chosen so that a one-dimensional slice satisfies

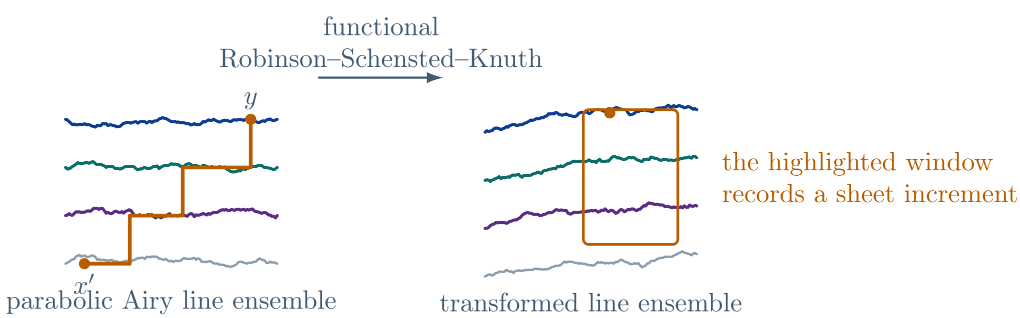

Figure 100

$$\mathcal{S}(0,y) \stackrel{\mathrm{d}}{=} \mathcal{P}_1(y) = \mathcal{A}_1(y) - y^2.$$

We next describe the last-passage-from-infinity construction. For $x > 0$ and $k \geq 1$, set

$$\boldsymbol{x}_k(x) := \left(-\sqrt{\frac{k}{2x}},\, k\right).$$

For points (a, i) and (b, j) in the parabolically shifted Airy line ensemble, write $d_{\mathcal{P}}((a, i), (b, j))$ for the corresponding line-ensemble last-passage value. For rational $x > 0$, y, and z, define

$$\begin{aligned} D_x(z, y) := d_{\mathcal{P}}\big(&\boldsymbol{x}_k(x), (y, 1)\big) \\ &- d_{\mathcal{P}}\big(\boldsymbol{x}_k(x), (z, 1)\big) \end{aligned} \tag{21.4}$$

for every sufficiently large k. The right-hand side eventually stabilizes almost surely. In geometric terms, the two optimizing staircases coalesce far enough in the line ensemble that moving their common entrance point farther toward infinity changes both passage values by the same amount. The construction appears in Figure 101. The stabilized difference determines only differences of

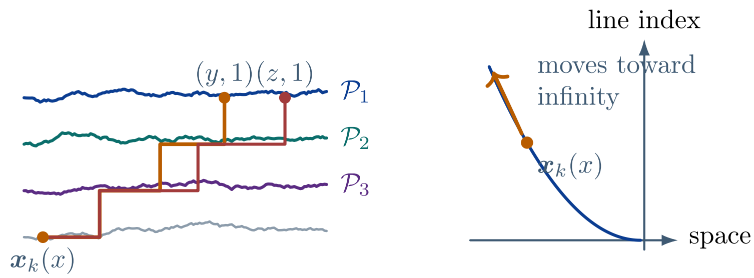

Figure 101

sheet values:

$$D_x(z, y) = \mathcal{S}(x, y) - \mathcal{S}(x, z).$$

The remaining additive constant is recovered by an ergodic average. More precisely,

$$\mathcal{S}(x, y) = \mathbb{E}[\mathcal{A}_1(0)] + \lim_{n\to\infty} \frac{1}{2n+1} \sum_{z=-n}^{n} \left(D_x(z, y) - (z - x)^2\right). \tag{21.5}$$

To understand the correction in (21.5), observe that $\mathcal{S}(x, z) + (z - x)^2$ is stationary in z and has mean $\mathbb{E}[\mathcal{A}_1(0)]$. The ergodic theorem therefore makes the average on the right converge to $\mathcal{S}(x, y) - \mathbb{E}[\mathcal{A}_1(0)]$.

The construction is first carried out simultaneously for rational $x > 0$, y, and z. Almost sure continuity and the symmetries of the model then give a unique continuous extension to all of $\mathbb{R}^2$.

Geodesics

The metric composition law gives a useful way to locate a geodesic. Fix endpoints $(0, 0)$ and $(0, 1)$, and let $r \in (0, 1)$. Define the profile

$$F_r(z) \coloneqq \mathcal{L}(0, 0; z, r) + \mathcal{L}(z, r; 0, 1).$$

By (3),

$$\mathcal{L}(0, 0; 0, 1) = \max_{z \in \mathbb{R}} F_r(z).$$

If γ is a geodesic between the two endpoints, then $\gamma(r)$ must be a maximizer of F_r. The two summands in F_r depend on disjoint time slabs and are therefore independent. This decomposition is shown in Figure 102.

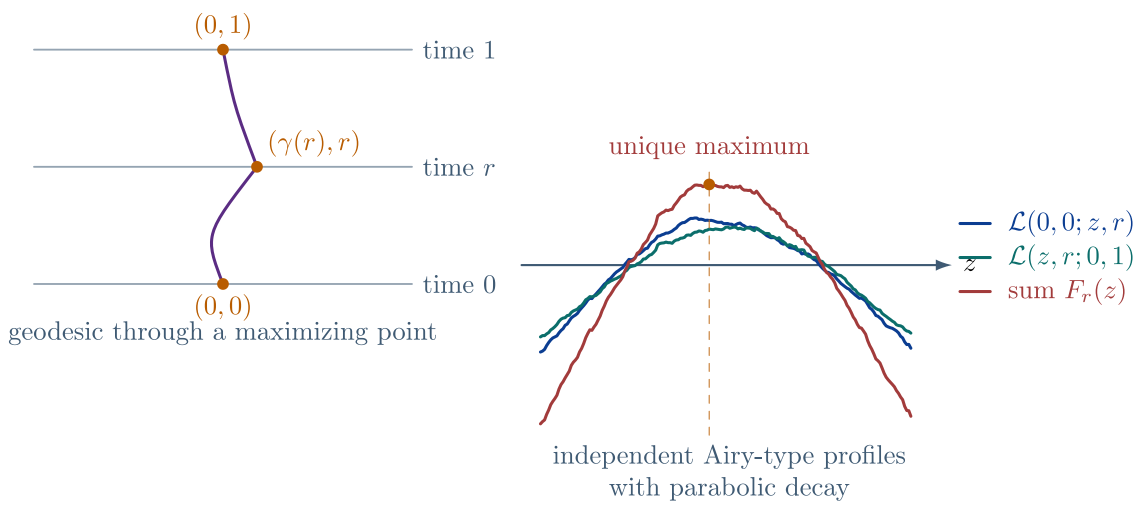

Figure 102

After their deterministic parabolic terms are removed, both summands are locally absolutely continuous with respect to Brownian motion. Their sum therefore has a unique maximizer on every compact interval almost surely. The negative parabolic drift controls the profile at spatial infinity, so the global maximizer is also unique. Applying this argument at every rational intermediate time and then using path continuity determines a unique geodesic. This is the basic mechanism behind the almost sure uniqueness assertion for deterministic endpoints.

Lecture 22 — Friday, April 09, 2021

5. Random Matrix Limits and Directed Landscape Applications

Wigner Semicircle Law

This treatment follows Todd Kemp's exposition of random matrix theory.

One motivation comes from image denoising. A greyscale image may be represented by a matrix whose entries record pixel intensities. After centring a local image block, we can model part of its noise by a random matrix. Spectral information then separates large structured modes from the bulk of the noise. The Wigner semicircle law identifies the limiting eigenvalue distribution of that random component.

Definition 22.1 For each n, let $\mathbf{Y}_n = (Y_{ij})_{1\le i,j\le n}$ be a real symmetric random matrix. Assume that the entry distributions below do not depend on n and that

1. The upper-triangular entries $\{Y_{ij} : 1 \le i \le j \le n\}$ are independent;
2. The off-diagonal entries are identically distributed, centred, and satisfy
$$\mathbb{E}[Y_{ij}^2] = t \in (0,\infty), \qquad i<j;$$
3. The diagonal entries are identically distributed, centred, and have finite variance.

The normalized matrix
$$\mathbf{X}_n \coloneqq \frac{1}{\sqrt{n}}\mathbf{Y}_n$$
is called a **Wigner matrix**.

Because $\mathbf{X}_n$ is real symmetric, it has n real eigenvalues, counted with multiplicity.

Definition 22.2 The **empirical spectral measure** of $\mathbf{X}_n$ is the random probability measure

$$\mu_{\mathbf{X}_n} := \frac{1}{n}\sum_{j=1}^{n} \delta_{\lambda_j(\mathbf{X}_n)}.$$

Definition 22.3 For $t > 0$, the **Wigner semicircle distribution** of variance t, denoted by σ_t, has density

$$\frac{\mathrm{d}\sigma_t}{\mathrm{d}x}(x) = \begin{cases} \frac{1}{2\pi t}\sqrt{4t - x^2}, & -2\sqrt{t} \le x \le 2\sqrt{t}, \\ 0, & \text{otherwise.} \end{cases}$$

The density for $t = 1$ is shown in Figure 103.

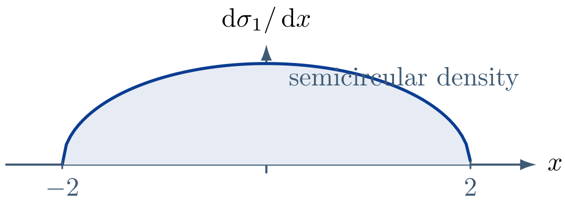

FIGURE 103

Theorem 22.4 (Wigner semicircle law) Let $(\mathbf{X}_n)_{n\ge 1}$ be Wigner matrices as in Definition 22.1. For every bounded continuous function $f : \mathbb{R} \to \mathbb{R}$ and every $\varepsilon > 0$,

$$\lim_{n\to\infty} \mathbb{P}\left(\left| \int_{\mathbb{R}} f \,\mathrm{d}\mu_{\mathbf{X}_n} - \int_{\mathbb{R}} f \,\mathrm{d}\sigma_t \right| > \varepsilon \right) = 0.$$

Equivalently, $\mu_{\mathbf{X}_n}$ converges weakly to σ_t in probability.

Remark 22.5 If $\mathbf{X}_n = \mathbf{U}_n \mathbf{\Lambda}_n \mathbf{U}_n^*$ is a spectral decomposition, then functional calculus and cyclicity of the trace give

$$\begin{aligned} \int_{\mathbb{R}} f \,\mathrm{d}\mu_{\mathbf{X}_n} &= \frac{1}{n}\sum_{j=1}^{n} f(\lambda_j(\mathbf{X}_n)) \\ &= \frac{1}{n}\operatorname{Tr}(f(\mathbf{\Lambda}_n)) \end{aligned}$$

$$= \frac{1}{n} \operatorname{Tr}(\mathbf{U}_n f(\mathbf{\Lambda}_n) \mathbf{U}_n^*)$$
$$= \frac{1}{n} \operatorname{Tr}(f(\mathbf{X}_n)).$$

Thus the empirical measure may be studied through normalized traces.

The Moment Method

For the elementary moment proof, assume initially that the entries have moments of every order. The general finite-second-moment statement in Theorem 22.4 follows by truncating large entries, recentering, and then removing the truncation.

Theorem 22.6 Fix $k \geq 0$. For every $\varepsilon > 0$,

$$\lim_{n\to\infty} \mathbb{P}\left(\left| \frac{1}{n} \operatorname{Tr}\left(\mathbf{X}_n^k\right) - \int_{\mathbb{R}} x^k \, \mathrm{d}\sigma_t(x) \right| > \varepsilon \right) = 0.$$

The substitution $x = \sqrt{t}\, u$ shows that

$$\int_{\mathbb{R}} x^k \, \mathrm{d}\sigma_t(x) = t^{k/2} \int_{\mathbb{R}} u^k \, \mathrm{d}\sigma_1(u).$$

Write

$$m_k := \int_{\mathbb{R}} x^k \, \mathrm{d}\sigma_1(x).$$

Symmetry gives $m_{2r+1} = 0$ for every $r \geq 0$, and $m_0 = 1$. For the even moments,

$$m_{2r} = C_r = \frac{1}{r+1} \binom{2r}{r},$$

where C_r is the rth Catalan number. Equivalently,

$$C_r = \frac{2(2r-1)}{r+1} C_{r-1}, \qquad r \geq 1.$$

Consequently,

$$\int_{\mathbb{R}} x^k \, \mathrm{d}\sigma_t(x) = \begin{cases} t^{k/2} C_{k/2}, & k \text{ is even}, \\ 0, & k \text{ is odd}. \end{cases} \tag{22.1}$$

By Chebyshev's inequality, it is enough to prove the following two statements:

1. The expectations converge:

$$\lim_{n\to\infty} \mathbb{E}\left[\frac{1}{n}\operatorname{Tr}\left(\mathbf{X}_n^k\right)\right] = \int_{\mathbb{R}} x^k \, \mathrm{d}\sigma_t(x).$$

2. The variances vanish:

$$\lim_{n\to\infty} \operatorname{Var}\left(\frac{1}{n}\operatorname{Tr}\left(\mathbf{X}_n^k\right)\right) = 0.$$

The combinatorial interpretation comes from matrix multiplication. For every square matrix $\mathbf{Y}$,

$$(\mathbf{Y}^k)_{ij} = \sum_{1\le i_2,\ldots,i_k\le n} Y_{ii_2}Y_{i_2i_3}\cdots Y_{i_kj}.$$

Each summand is encoded by a walk through its sequence of indices. In a trace, the walk is closed. For example, the sequence

$$i_1 = 1, \qquad i_2 = i_3 = 4, \qquad i_4 = 6, \qquad i_5 = 2 \qquad \text{and} \qquad i_6 = i_1$$

gives the five-step closed walk in Figure 104. The loop at vertex 4 records the step from i_2 to i_3.

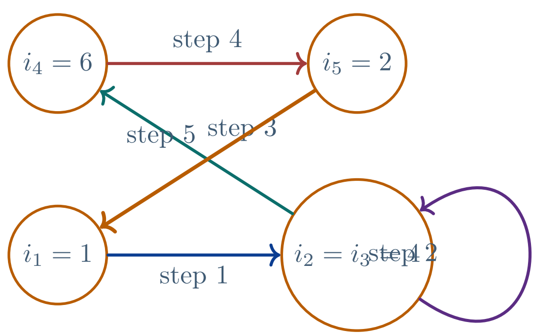

FIGURE 104

Proof of Theorem 22.6. Expanding the normalized trace gives

$$\mathbb{E}\left[\frac{1}{n}\operatorname{Tr}\left(\mathbf{X}_n^k\right)\right] = n^{-1-k/2} \sum_{i_1,\ldots,i_k=1}^{n} \mathbb{E}[Y_{i_1i_2}Y_{i_2i_3}\cdots Y_{i_ki_1}]\,. \tag{22.2}$$

Associate an unoriented graph to each index sequence: its vertices are the distinct indices appearing in the sequence, and its edges are the unordered pairs traversed by the closed walk. Suppose that the graph has v vertices and e distinct edges.

If some edge is traversed exactly once, the corresponding expectation in (22.2) vanishes by centring and independence. Hence every contributing edge is traversed at least twice, so

$$e \leq \frac{k}{2}.$$

The graph traced by one walk is connected, and therefore $v \leq e+1$. There are $O(n^v)$ ways to assign labels from $\{1,\dots,n\}$ to a fixed unlabeled walk. After multiplication by the factor $n^{-1-k/2}$, its total order is at most

$$n^{v-1-k/2} \leq n^{e-k/2}.$$

This tends to 0 unless $k=2r$, $e=r$, and $v=r+1$. In the surviving case, the graph is a tree and every edge is traversed exactly twice. The expectation of the associated product is then t^r.

After choosing a starting corner and recording the order in which edges are first explored, the surviving walks are in bijection with rooted plane trees with r edges. These trees are counted by the Catalan number C_r. It follows that

$$\lim_{n\to\infty} \mathbb{E}\left[\frac{1}{n}\operatorname{Tr}\Big(\mathbf{X}_n^k\Big)\right] = \begin{cases} t^r C_r, & k=2r, \\ 0, & k \text{ is odd.} \end{cases}$$

Together with (22.1), this proves (1).

For (2), expand the square of the trace as a sum over pairs of closed k-step walks. If the two walks use disjoint sets of matrix entries, independence makes their covariance zero. Every remaining pair shares an edge. Its union is connected, every edge that survives centring appears at least twice, and the same vertex-counting argument shows that the total contribution is $O(n^{-1})$. Hence

$$\operatorname{Var}\left(\frac{1}{n}\operatorname{Tr}\Big(\mathbf{X}_n^k\Big)\right) \longrightarrow 0.$$

Chebyshev's inequality now proves the asserted convergence in probability. □

The moments in (22.1) determine the semicircle distribution uniquely. Moreover, the second-moment bound gives tightness of the empirical spectral measures. Polynomial approximation on compact sets then upgrades the matrix moment convergence to weak convergence against bounded continuous functions. Finally, truncation and recentering remove the temporary higher-moment assumption and yield Theorem 22.4.

Lecture 23 — Monday, April 12, 2021

The Decorated Airy Point Process

This section follows Gorin and Sodin's work on the Kardar–Parisi–Zhang equation and moments of random matrices.

Let $\mathbf{M}_N$ be an unnormalized $N \times N$ matrix from the Gaussian orthogonal ensemble when $\beta = 1$, or from the Gaussian unitary ensemble when $\beta = 2$. Order its eigenvalues as

$$\lambda_{1,N} \geq \lambda_{2,N} \geq \cdots \geq \lambda_{N,N},$$

and choose orthonormal eigenvectors $\boldsymbol{V}_{1,N}, \ldots, \boldsymbol{V}_{N,N}$. The spectral theorem gives

$$\left(\frac{\mathbf{M}_N}{2\sqrt{N}}\right)^m = \sum_{j=1}^{N} \left(\frac{\lambda_{j,N}}{2\sqrt{N}}\right)^m \boldsymbol{V}_{j,N} \boldsymbol{V}_{j,N}^*.$$

When m is of order $N^{2/3}$, eigenvalues in the bulk are exponentially suppressed and only the spectral edges remain. Thus high matrix powers simultaneously retain edge eigenvalues and the rank-one projections onto their eigenvectors.

Two asymptotic facts explain the limiting object:

1. By orthogonal or unitary invariance, the eigenbasis is Haar distributed and independent of the eigenvalues. For every fixed collection of coordinate and eigenvector indices, the variables $\sqrt{N}\, \boldsymbol{V}_{j,N}(a)$ converge jointly to independent standard real Gaussians when $\beta = 1$, and to standard complex Gaussians when $\beta = 2$.
2. For every fixed k,
$$\left(N^{1/6}(\lambda_{j,N} - 2\sqrt{N})\right)_{j=1}^{k} \xrightarrow{\mathrm{d}} (\lambda_j)_{j=1}^{k},$$
where $\lambda_1 > \lambda_2 > \cdots$ are the ordered points of the Airy$_\beta$ point process.

The eigenvector limits in (1) decorate each limiting Airy point with a random rank-one matrix.

Definition 23.1 Let

$$\{u_a^n, v_a^n : a, n \geq 1\}$$

be independent standard real Gaussian random variables, independent of an Airy$_\beta$ point process $\lambda_1 > \lambda_2 > \cdots$. For each $n \geq 1$, define the infinite Hermitian rank-one matrix $\mathbf{W}_n$ by

$$[\mathbf{W}_n]_{a,b} = \begin{cases} u_a^n u_b^n, & \beta = 1, \\ \frac{1}{2}(u_a^n + \mathrm{i}v_a^n)(u_b^n - \mathrm{i}v_b^n), & \beta = 2. \end{cases}$$

Equivalently, define the formal column vector $\boldsymbol{w}_n$ by

$$[\boldsymbol{w}_n]_a = \begin{cases} u_a^n, & \beta = 1, \\ 2^{-1/2}(u_a^n + \mathrm{i}v_a^n), & \beta = 2. \end{cases}$$

Then $\mathbf{W}_n = \boldsymbol{w}_n \boldsymbol{w}_n^*$, which makes its rank-one structure explicit. The **decorated Airy$_\beta$ point process** is the matrix-valued random measure

$$\mathrm{Ai}_\beta^{\mathrm{dec}} := \sum_{n=1}^{\infty} \mathbf{W}_n \delta_{\lambda_n}.$$

The construction is represented in Figure 105: every Airy point carries an independent rank-one matrix recording the limiting coordinates of its eigenvector. Let $f\colon \mathbb{R} \to \mathbb{C}$ be continuous and

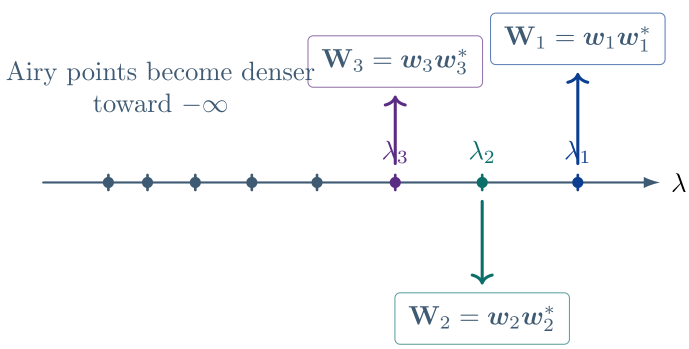

Figure 105

satisfy

$$\int_{-\infty}^{0} |f(\lambda)|\sqrt{|\lambda|}\,\mathrm{d}\lambda < \infty.$$

The integral of f against the decorated process is the infinite Hermitian matrix defined entrywise by

$$\left[\int_{\mathbb{R}} f(\lambda)\,\mathrm{Ai}_\beta^{\mathrm{dec}}(\mathrm{d}\lambda)\right]_{a,b} = \sum_{n=1}^{\infty} [\mathbf{W}_n]_{a,b} f(\lambda_n). \tag{23.1}$$

The integrability condition compensates for the increasing density of Airy points at $-\infty$. In particular, it holds for $f(\lambda) = e^{\alpha\lambda}$ with $\alpha > 0$.

Theorem 23.2 Let $\mathbf{M}_N$ be an unnormalized real symmetric or complex Hermitian Wigner matrix. Assume that the mixed moments of each entry through total degree 4 agree with those of the corresponding Gaussian orthogonal or Gaussian unitary ensemble, and that, for the absolute constant C_0 in the edge-universality estimate,

$$\sup_N \max_{1\le i,j\le N} \mathbb{E}\big[|[\mathbf{M}_N]_{i,j}|^{C_0}\big] < \infty.$$

For $\alpha > 0$, set $m_N = \lfloor \alpha N^{2/3} \rfloor$. Then

$$\frac{N}{2}\left[\left(\frac{\mathbf{M}_N}{2\sqrt{N}}\right)^{2m_N} + \left(\frac{\mathbf{M}_N}{2\sqrt{N}}\right)^{2m_N+1}\right] \xrightarrow{\mathrm{d}} \int_{\mathbb{R}} e^{\alpha\lambda}\,\mathrm{Ai}_\beta^{\mathrm{dec}}(\mathrm{d}\lambda). \tag{23.2}$$

The convergence holds jointly for every fixed finite collection of matrix entries and every fixed finite collection of positive values of α.

To see the exponential weight in the limit, suppose

$$\lambda_{j,N} = 2\sqrt{N} + N^{-1/6}\xi_{j,N}.$$

Then

$$\left(\frac{\lambda_{j,N}}{2\sqrt{N}}\right)^{2m_N} = \left(1 + \frac{\xi_{j,N}}{2N^{2/3}}\right)^{2\lfloor \alpha N^{2/3}\rfloor} \longrightarrow e^{\alpha\xi_j}.$$

The sum of the even and odd powers in (23.2) doubles the contribution from the upper edge and cancels the contribution from the lower edge. For the Gaussian ensembles, the result then follows from the independent limits (1) and (2). The four-moment comparison extends the conclusion to the stated Wigner class, while the uniform C_0-moment bound controls the contribution of rare large entries.

Lecture 24 — Wednesday, April 14, 2021

Eden Growth and Directed Landscape Models

The discrete Eden growth model connects naturally to exponential last-passage percolation and, after scaling, to the directed landscape. This connection explains why subtracting the deterministic limit shape from a simulated growth profile reveals complicated random contours.

Eden Growth and Exponential Last-Passage Percolation

Begin with a finite seed $S \subseteq \mathbb{Z}^2$. In the vertex version of the discrete Eden growth model, set $B_0(S) = S$, and at each step add a uniformly selected site from the outer vertex boundary of the current cluster. This produces an increasing chain

$$B_0(S) \subseteq B_1(S) \subseteq \cdots \subseteq B_m(S) \subseteq \cdots .$$

A directed analogue is the **exponential corner growth model**. Place independent exponential random variables of mean 1,

$$w_{\boldsymbol{v}} \sim \mathrm{Exp}(1), \qquad \boldsymbol{v} = (m, n) \in \mathbb{Z}^2_{\geq 0},$$

at the lattice vertices. For $\boldsymbol{p} \leq \boldsymbol{q}$ coordinatewise, define

$$d(\boldsymbol{p}, \boldsymbol{q}) := \max_{\pi : \boldsymbol{p} \nearrow \boldsymbol{q}} \sum_{\boldsymbol{v} \in \pi} w_{\boldsymbol{v}},$$

where the maximum is over up-right lattice paths from $\boldsymbol{p}$ to $\boldsymbol{q}$. The sets on which $d(\mathbf{0}, \boldsymbol{q})$ is below a prescribed threshold form a random corner growth cluster.

The deterministic shape function is

$$g(a, b) := (\sqrt{a} + \sqrt{b})^2 = a + b + 2\sqrt{ab}, \qquad a, b \geq 0. \tag{24.1}$$

Along the diagonal $a = b$, the tangent plane to g is $2(a + b)$. For $\boldsymbol{q} = (a, b)$, the centring

$$d(\mathbf{0}, \boldsymbol{q}) - 2(a + b)$$

removes the macroscopic linear profile near the diagonal. What remains contains a negative quadratic curvature term and random fluctuations of order $(a+b)^{1/3}$. Consequently, dense contours of the uncentred profile look nearly parallel, while centred contours develop branching arms, narrow peninsulas, and occasional closed components. The transition is represented schematically in Figure 106. The fine geometry of these contours raises natural questions about fractal and

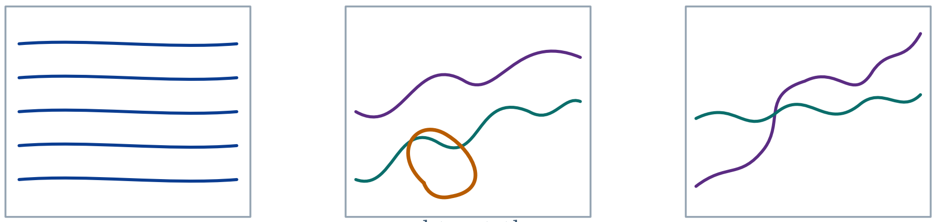

uncentred contours
macroscopic shape
dominates

subtract the
tangent plane
random geometry
is exposed

Kardar–Parisi–Zhang
scaling
directed landscape
contours

Figure 106

Hausdorff dimensions. One-point fluctuation limits do not answer such questions; a joint metric limit is needed.

Airy Objects Revisited

Let
$$C^{\mathbb{N}} \coloneqq C(\mathbb{R}\times\mathbb{N},\mathbb{R})$$
with the product of the uniform-on-compact topologies. Let WB^n be a Brownian n-melon and define
$$P_i^n(y) \coloneqq n^{1/6}\left((WB^n)_i(1+2yn^{-1/3}) - 2\sqrt{n} - 2yn^{1/6}\right).$$
For every fixed number of lines,
$$(P_1^n, P_2^n, \dots) \xrightarrow{\mathrm{d}} (\mathcal{P}_1, \mathcal{P}_2, \dots) \quad \text{in } C^{\mathbb{N}},$$
where $\mathcal{P}_i(y) = \mathcal{A}_i(y) - y^2$ is the parabolically shifted Airy line ensemble. Its curves are strictly ordered:
$$\mathcal{P}_i(y) > \mathcal{P}_j(y) \qquad \text{when } i<j.$$
The Brownian normalization $2^{-1/2}\mathcal{P}$ has the Gibbs property from Theorem 15.3.

For completeness, the Airy sheet from Definition 21.8 may be characterized by the following two properties:

1. It is diagonally stationary:
$$\mathcal{S}(\,\cdot,\cdot\,) \overset{\mathrm{d}}{=} \mathcal{S}(\,\cdot+h,\cdot+h\,) \qquad \text{for every } h\in\mathbb{R}.$$

2. It can be coupled with $\mathcal{P}$ so that $\mathcal{S}(0,\cdot)=\mathcal{P}_1(\cdot)$, and, for every $x\in\mathbb{Q}_{>0}$ and $y,z\in\mathbb{Q}$, there is a random integer $K_{x,y,z}$ such that, when $k\geq K_{x,y,z}$,
$$\begin{aligned} &d_{\mathcal{P}}\left(\left(-\sqrt{\frac{k}{2x}},k\right),(z,1)\right)\\ &\quad - d_{\mathcal{P}}\left(\left(-\sqrt{\frac{k}{2x}},k\right),(y,1)\right) = \mathcal{S}(x,z)-\mathcal{S}(x,y). \end{aligned}$$

For $a>0$, the Airy sheet of scale a is
$$\mathcal{S}_a(x,y) \coloneqq a\,\mathcal{S}\left(\frac{x}{a^2},\frac{y}{a^2}\right).$$

If $\mathcal{S}_a$ and $\mathcal{S}_b$ are independent sheets, then their max-plus composition is a sheet of scale c, where $c^3=a^3+b^3$:
$$\sup_{y\in\mathbb{R}}\left\{\mathcal{S}_a(x,y)+\mathcal{S}_b(y,z)\right\} \overset{\mathrm{d}}{=} \mathcal{S}_c(x,z). \tag{24.2}$$

The directed landscape is the spacetime process that realizes (24.2) simultaneously. Equivalently, it is the continuous random function on $\mathbb{R}^4_{\uparrow}$ satisfying metric composition
$$\mathcal{L}(x,r;y,t)=\sup_{z\in\mathbb{R}}\left\{\mathcal{L}(x,r;z,s)+\mathcal{L}(z,s;y,t)\right\}, \qquad r<s<t,$$
whose restrictions to disjoint time intervals are independent Airy sheets with the corresponding scales. This agrees with Theorem 21.1.

Scaling Exponential Last-Passage Percolation

For compatibility with the time orientation of the directed landscape, index exponential last-passage points by $(x,-y)$. Fix a direction $\rho>0$, and define positive constants α,β,χ,τ by
$$\chi^3=\frac{(\sqrt{\rho}+1)^4}{\sqrt{\rho}}, \qquad \alpha=(\sqrt{\rho}+1)^2, \qquad \beta=1+\frac{1}{\sqrt{\rho}}, \qquad \text{and} \qquad \frac{\chi}{\tau^2}=\frac{1}{4\rho^{3/2}}. \tag{24.3}$$

Set the direction vectors
$$\boldsymbol{v}=(\rho,-1) \qquad \text{and} \qquad \boldsymbol{u}=(\tau,0).$$

For the diagonal direction $\rho = 1$, these constants become

$$\alpha = 4, \qquad \beta = 2, \qquad \chi = 2^{4/3} \qquad \text{and} \qquad \tau = 2^{5/3}.$$

Theorem 24.1 There is a coupling of copies d_σ of exponential last-passage percolation with a directed landscape $\mathcal{L}$ such that

$$\begin{aligned} & d_\sigma \left(x\sigma^2 \boldsymbol{u} + s\sigma^3 \boldsymbol{v},\, y\sigma^2 \boldsymbol{u} + t\sigma^3 \boldsymbol{v}\right) \\ &\quad = \alpha\sigma^3(t-s) + \beta\tau\sigma^2(y-x) \\ &\qquad + \chi\sigma\left(\mathcal{L}(x,s;y,t) + o_\sigma(x,s;y,t)\right). \end{aligned} \tag{24.4}$$

The lattice metric is linearly interpolated when the displayed arguments are not lattice points. The error tends to 0 uniformly on compact subsets of $\mathbb{R}^4_\uparrow$ in probability. More quantitatively, for every such compact set K, we may choose $a > 1$ so that

$$\mathbb{E}\left[a^{\sup_K |o_\sigma|^{3/4}}\right] \longrightarrow 1.$$

The constants in (24.3) come from the shape and its curvature. If $\Delta = t - s$ and $\delta = y - x$, then (24.1) gives

$$\begin{aligned} & g\left(\rho\sigma^3\Delta + \tau\sigma^2\delta,\, \sigma^3\Delta\right) \\ &\quad = (\sqrt{\rho}+1)^2\sigma^3\Delta + \tau\left(1 + \frac{1}{\sqrt{\rho}}\right)\sigma^2\delta \\ &\qquad - \frac{\tau^2}{4\rho^{3/2}}\sigma\frac{\delta^2}{\Delta} + O(1). \end{aligned}$$

The first two coefficients are α and $\beta\tau$. The relation $\chi/\tau^2 = 1/(4\rho^{3/2})$ normalizes the quadratic term to the parabola in (21.2), while the exact one-point fluctuation scale determines χ.

At the law-of-large-numbers scale, the same computation says

$$\frac{1}{\sigma^3} d(\sigma^3\boldsymbol{p}, \sigma^3\boldsymbol{q}) \longrightarrow \|\boldsymbol{q} - \boldsymbol{p}\|_{1/2},$$

where, for a directed displacement $\boldsymbol{r} = (a, b)$,

$$\|\boldsymbol{r}\|_{1/2} \coloneqq (\sqrt{a} + \sqrt{b})^2.$$

Thus the unusual centred contours are not artefacts of plotting. They are the discrete precursors of the Airy sheet and directed landscape geometry revealed after the deterministic tangent plane is removed.

Lecture 25 — Friday, April 16, 2021

Sampling Determinantal Point Processes

The correlation kernel of a determinantal point process also gives an exact sampling algorithm. The construction has two stages: a collection of independent Bernoulli choices first reduces a general kernel to a projection kernel, and a sequence of orthogonal projections then samples the points of the projection process.

Definition 25.1 Let (Λ, μ) be a locally compact measured space. A simple point process $\mathcal{X}$ on Λ is a **determinantal point process** with kernel K if its correlation functions satisfy

$$\rho_k(x_1, \ldots, x_k) = \det\big[K(x_i, x_j)\big]_{i,j=1}^{k} \tag{25.1}$$

for every positive integer k and for μ^k-almost every $(x_1, \ldots, x_k) \in \Lambda^k$.

We view K both as a kernel and as its integral operator on $L^2(\Lambda, \mu)$. A Hermitian, locally trace class kernel defines a determinantal point process precisely when

$$0 \le K \le I \tag{25.2}$$

as an operator. Thus, positive definiteness alone is not sufficient: the upper bound in (25.2) is also necessary.

Definition 25.2 Let H be an n-dimensional subspace of $L^2(\Lambda, \mu)$, and let $\varphi_1, \ldots, \varphi_n$ be an orthonormal basis of H. The orthogonal projection onto H has kernel

$$K_H(x, y) = \sum_{j=1}^{n} \varphi_j(x)\overline{\varphi_j(y)}. \tag{25.3}$$

The determinantal point process with kernel K_H is called a **projection determinantal point process**.

The number of points in this process equals n almost surely. Indeed, its probability generating function is

$$\mathbb{E}\big[z^{|\mathcal{X}|}\big] = \det\big(I + (z-1)K_H\big) = z^n.$$

Consequently, its unordered n-point configuration has density

$$\frac{1}{n!}\det\big[K_H(x_i,x_j)\big]_{i,j=1}^n \prod_{i=1}^n \mu(\mathrm{d}x_i). \tag{25.4}$$

Theorem 25.3 (Spectral Bernoulli decomposition) Suppose that K is a positive, trace class contraction with spectral decomposition

$$K(x,y) = \sum_{j\ge 1} \lambda_j \varphi_j(x)\overline{\varphi_j(y)}, \qquad 0 \le \lambda_j \le 1, \tag{25.5}$$

where the eigenfunctions are orthonormal. Let $B_1, B_2, \dots$ be independent random variables such that

$$\mathbb{P}(B_j = 1) = \lambda_j \qquad \text{and} \qquad \mathbb{P}(B_j = 0) = 1 - \lambda_j.$$

Define the random subspace and its projection kernel by

$$H_B \coloneqq \overline{\operatorname{span}}\{\varphi_j : B_j = 1\} \qquad \text{and} \qquad K_B(x,y) \coloneqq \sum_{j\ge 1} B_j \varphi_j(x)\overline{\varphi_j(y)}.$$

Conditional on the variables $(B_j)_{j\ge 1}$, sample the projection determinantal point process with kernel K_B. After averaging over the Bernoulli variables, the resulting point process is the determinantal point process with kernel K.

Proof. Since K is trace class,

$$\sum_{j\ge 1} \lambda_j = \operatorname{Tr} K < \infty.$$

It follows that only finitely many of the variables B_j equal 1 almost surely, so H_B is almost surely finite dimensional.

Let g be bounded and compactly supported, and let M_g denote multiplication by g. The determinantal generating functional gives

$$\mathbb{E}\left[\prod_{x\in\mathcal{X}} (1 + g(x)) \,\middle|\, (B_j)_{j\ge 1}\right] = \det(I + M_g K_B).$$

Expanding this Fredholm determinant in the eigenfunction basis makes it multilinear in the variables B_j. Independence therefore yields

$$\mathbb{E}_B\left[\det(I + M_g K_B)\right] = \det(I + M_g K),$$

because $\mathbb{E}[B_j] = \lambda_j$. The right-hand side is the generating functional of the determinantal point process with kernel K, which proves the claim. □

It remains to sample a projection determinantal point process. For an r-dimensional subspace H_r, the diagonal of its projection kernel defines the probability measure

$$\nu_{H_r}(\mathrm{d}x) \coloneqq \frac{1}{r} K_{H_r}(x,x)\,\mu(\mathrm{d}x). \tag{25.6}$$

This is normalized because

$$\int_\Lambda K_{H_r}(x,x)\,\mu(\mathrm{d}x) = \operatorname{Tr} K_{H_r} = r.$$

Moreover, for μ-almost every x with $K_{H_r}(x,x) > 0$, the vector $K_{H_r}(\,\cdot\,,x)$ in H_r represents evaluation at x, and

$$\left\|K_{H_r}(\,\cdot\,,x)\right\|^2_{L^2(\mu)} = K_{H_r}(x,x). \tag{25.7}$$

On a discrete state space, $K_{H_r}(\,\cdot\,,x) = K_{H_r}\delta_x$, so (25.7) is the familiar norm of the projected coordinate vector.

Theorem 25.4 Let H be an n-dimensional subspace of $L^2(\Lambda,\mu)$. Initialize $H_n = H$. For $r = n, n-1, \ldots, 1$, perform the following steps:

1. Sample X_r from the probability measure ν_{H_r} in (25.6).
2. Form the normalized evaluation vector

$$\boldsymbol{v}_r \coloneqq \frac{K_{H_r}(\,\cdot\,,X_r)}{\sqrt{K_{H_r}(X_r,X_r)}}$$

and remove its direction by setting

$$H_{r-1} \coloneqq H_r \cap \boldsymbol{v}_r^\perp = \{f \in H_r : f(X_r) = 0\}. \tag{25.8}$$

Then the unordered set $\{X_1,\ldots,X_n\}$ has the projection determinantal point process with kernel K_H.

Proof. Step (2) lowers the dimension by one. The projection kernel of the new subspace is the rank-one update

$$K_{H_{r-1}}(x,y) = K_{H_r}(x,y) - \frac{K_{H_r}(x,X_r)K_{H_r}(X_r,y)}{K_{H_r}(X_r,X_r)}. \tag{25.9}$$

Applying the Schur complement identity to the Gram matrix of the evaluation vectors gives

$$\det\big[K_{H_r}(x_i,x_j)\big]_{i,j=1}^{r} = K_{H_r}(x_r,x_r)\det\big[K_{H_{r-1}}(x_i,x_j)\big]_{i,j=1}^{r-1}. \tag{25.10}$$

The density contributed by step (1) is $K_{H_r}(x_r,x_r)/r$. Iterating (25.10) therefore shows that the ordered sample has density

$$\prod_{r=1}^{n}\frac{K_{H_r}(x_r,x_r)}{r} = \frac{1}{n!}\det\big[K_H(x_i,x_j)\big]_{i,j=1}^{n}.$$

This is precisely the symmetrized density (25.4). □

Combining Theorem 25.3 with Theorem 25.4 gives an exact sampler for every trace class determinantal kernel, as summarized in Figure 107.

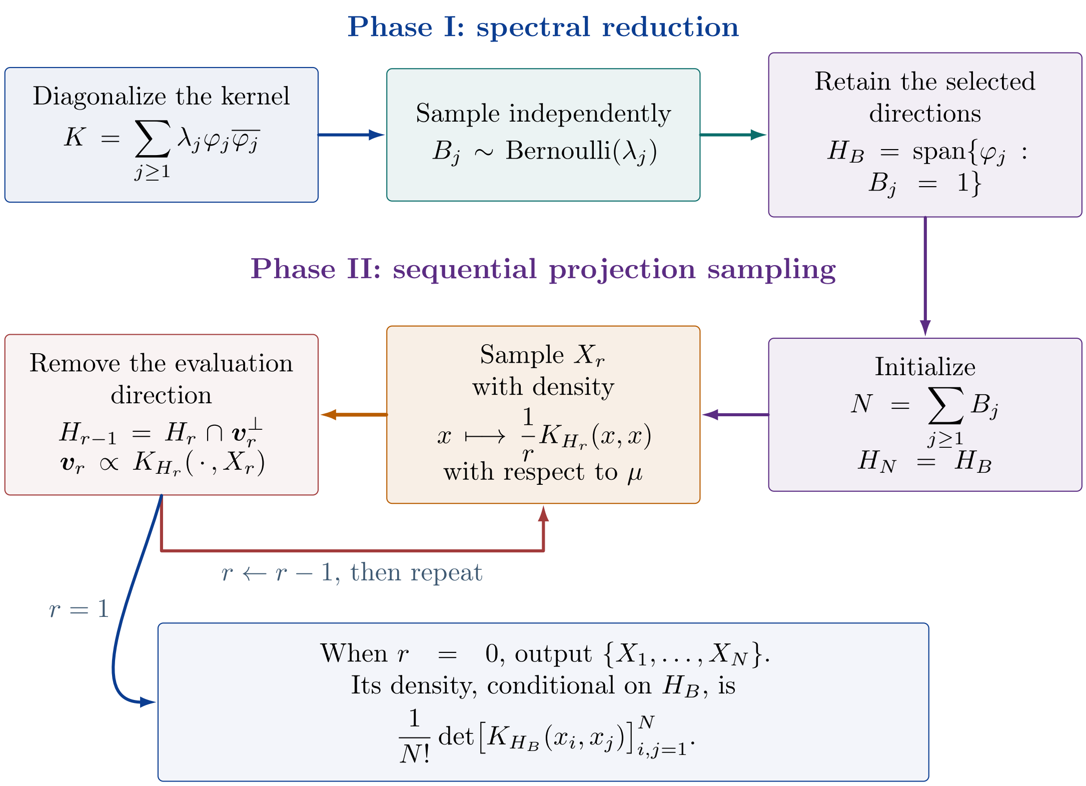
Phase I: spectral reduction
Diagonalize the kernel $K = \sum_{j\geq 1} \lambda_j \varphi_j \overline{\varphi_j}$
Sample independently $B_j \sim \text{Bernoulli}(\lambda_j)$
Retain the selected directions $H_B = \text{span}\{\varphi_j : B_j = 1\}$
Phase II: sequential projection sampling
Initialize $N = \sum_{j\geq 1} B_j$ $H_N = H_B$
Sample X_r with density $x \longmapsto \frac{1}{r} K_{H_r}(x,x)$ with respect to μ
Remove the evaluation direction $H_{r-1} = H_r \cap \boldsymbol{v}_r^{\perp}$ $\boldsymbol{v}_r \propto K_{H_r}(\cdot, X_r)$
$r \leftarrow r-1$, then repeat
$r = 1$
When $r = 0$, output $\{X_1, \ldots, X_N\}$. Its density, conditional on H_B, is $\frac{1}{N!} \det\left[K_{H_B}(x_i, x_j)\right]_{i,j=1}^{N}$.

FIGURE 107

Appendix: Lecture and Tutorial Index