Unimodular random graphs and quantum percolation

Tasmin Chu • 3 Jun 2026

https://functor.network/user/3440/entry/1842

Editor's note: This blogpost is roughly based on an in-class presentation I gave for a class on Anderson localization (Ma 191c) at Caltech, which was taught by Lingfu Zhang. The presentation was based on the 2013 paper **Mean quantum percolation** by Charles Bordenave, Arnab Sen, and Bálint Virág. I make no claim to originality in this blogpost.

Motivating question: What features of random graphs are captured by their adjacency matrices (which are necessarily *random* matrices) and their associated spectral measures?

I'm going to first give you *results* without formally defining things as I should, and then I'll formally define what I mean by things like: expected spectral measure, spectral measure of the adjacency matrix, local convergence, unimodular random graph, etc. I'll also remark that by training, I use the word Bienaymé tree to refer to what many other probabilists call Galton-Watson trees (or less precisely, branching processes.)

Let's take some case studies.

Model 1: Bernoulli(p) percolation on $\mathbb{Z}^d$

Consider Bernoulli(p) bond percolation on $\mathbb{Z}^d$, so independently for each edge of $\mathbb{Z}^d$, we delete an edge with probability $1-p$ and retain it with probability p. Call the resulting random graph $G[p]$. For $p=1$, $G[p]$ is just $(\mathbb{Z}^d, o)$, which has absolutely continuous spectral measure. In fact,

$$\mathrm{Spec}(A) = [-2d, 2d]$$

where Spec denotes the spectrum of the adjacency operator A of $\mathbb{Z}^d$. (Recall that for a bounded linear operator $A : \mathcal{H} \to \mathcal{H}$ on some Hilbert space H, we have that

$$\mathrm{Spec}(A) = \{\lambda \in \mathbb{C} : A - \lambda I \text{ is not bijective}\}$$

In fact, the corresponding spectral measure μ_o (called the Plancherel measure) is exactly the d-fold convolution of the measure μ defined by the density

$$d\mu = \frac{1}{\pi\sqrt{4-\lambda^2}} 1_{[-2,2]}(\lambda) d\lambda$$

On the other hand, for $p < p_c$, because all clusters are finite, the expected spectral measure of $A_{G[p]}$ is some countable mixture of atomic measures. In other words, it is purely atomic.

For the two-dimensional case, at least, the expected spectral measure does indeed capture something about the phase diagram:

Theorem 1. *For bond percolation on $\mathbb{Z}^2$, the expected spectral measure has a continuous part iff $p > p_c$.*
Remark: That they could show this for $\mathbb{Z}^2$ and not $d \geq 3$ is probably not surprising to many in the field. Indeed, the proof uses the planar self-duality of $\mathbb{Z}^2$, which obviously cannot be leveraged in high dimensions.

Model 2: The Erdős–Rényi graph $G(n, p)$

Consider the complete graph K_n on n vertices (so there is an edge between any two vertices u, v in the ambient graph). An extremely well-studied random graph model is Bernoulli(p_n) percolation on K_n. We call the graph $G(n, p)$ the Erdős–Rényi graph $G(n, p)$. Of course, this is often more interesting when we let p_n be some function on n.

What does the adjacency matrix look like? We can describe it quite well: because this graph has no self-loops, it's 0 on the diagonal ($A_{ii} = 0$), while off the diagonal $A_{ij} = A_{ji} = \mathrm{Ber}(p_n)$, with $\mathbb{E}[A_{ij}] = p_n$. We should expect the operator norm of A_n (the adjacency matrix of $G(n, p_n)$ to be roughly $\sqrt{n(p_n)(1-p_n)}$.

So, if we look at the appropriately centred/normalized and rescaled matrices,

$$\widetilde{A_n} = \frac{A_n - p_n 1 1^T}{\sqrt{np_n(1-p_n)}}$$

so long as $\sqrt{np_n(1-p_n)} \nearrow \infty$, we can use a universality result (Wigner's semicircle law) and conclude that the distribution of the eigenvalues converges to the Wigner semicircular distribution which has density

$$d\mu_{sc} = \frac{1}{2\pi}(4 - |x|^2)_+^{1/2} dx$$

What about the sparse regime? That is, the regime where say $p_n = c/n$ and thus $\mathbb{E}[np_n] \to c \in (0, \infty)$? Here the connection probability c/n

decays in n. In this case, the limiting expected spectral measure μ_c was already known to be quite different. It was known the limiting spectral measure was supported on all of $\mathbb{R}$. Recall by the Lebesgue-Radon-Nikodym theorem that any measure ν on $\mathbb{R}$ decomposes into a pure-point part, a part which is absolutely continuous with respect to the Lebesgue measure, and a part which is atomless but singular with respect to the Lebesgue measure (I'll call this the singular continuous part of the measure.) It was known that the pure-point of this spectral measure μ^{pp} was supported on a dense set of atoms. Since the closure of a dense set of atoms is all of $\mathbb{R}$, it's totally consistent for this measure μ_c to have no continuous part, but it's also possible for this measure to have a continuous part.

Question 2. When does the measure μ_c have a continuous part?
It turns out the answer is exactly what you would hope it would be:

Theorem 3. *The measure μ_c has a continuous part iff $c > 1$.*
To most discrete probabilists, this isn't too surprising. Indeed, if you study the largest component $C_{\max}(G_n)$ in $G_n = G(n, c/n)$, a classical theorem says that

$$\frac{|C_{\max}(G_n)|}{n} \xrightarrow{p} \theta(c)$$

where $\theta(c)$ is the survival probability of an $\mathrm{Pois}(c)$ Bienaymé tree.

(More can be shown: $c = 1$ is exactly the critical theshold. For $c < 1$, there are multiple $\log(n)$ components; for $c > 1$, there is a unique giant component of linear size, meaning of size $\theta(c)n$, while the other components are $O(\log(n))$. At $c = 1$, the largest component is of size $n^{2/3}$.)

In other words, μ_c, the expected spectral measure captures the phase transition that we observe in $G(n, c/n)$. But why is this true? From a local standpoint, for $G(n, c/n)$, a uniformly random root sees an $\mathrm{Pois}(c)$ Bienaymé tree. Intuitively, if we do breadth-first search of a component, what we see looks like a (unimodular) Bienaymé tree with offspring distribution $\mathrm{Binomial}(n, c/n)$. But $\mathrm{Binomial}\, n, c/n$ random variables converge in total variation distance to $\mathrm{Pois}(c)$ random variables. In fact, you can show that there exists a coupling of $X \sim \mathrm{Pois}(c)$, $Y \sim \mathrm{Bin}(n, c/n)$ such that

$$\mathbb{P}(X \neq Y) \leq \frac{\lambda^2}{n}.$$

It turns out that this analysis will follow by one last case study: (unimodular) Bienaymé trees.

Model 3: (Unimodular) Bienaymé trees

Fix a probability measure Z on $\mathbb{N}_{\geq 0}$ with finite mean. The unimodular Bienaymé tree with offspring distribution Z can be described as follows: the root $\varnothing$ produces $Z_\varnothing$ children according to Z, while every other vertex independently of the root and of each other reproduces according to $\hat{Z}$, the size-biasing of Z. Here, the size-biasing of Z is defined by

$$\hat{Z}(k) = \frac{(k+1)\hat{Z}(k+1)}{\mathbb{E}Z}$$

What does the expected spectral measure look like?

Theorem 4. *The expected spectral measure for a unimodular Bienaymé(Z) with $\hat{Z} \neq \delta_2$ has a non-trivial continuous part iff*

$$\mu(\hat{Z}) > 1.$$

Again we see that some kind of phase transition is captured by the spectral measure. Indeed, by the extinction criterion for (typical) Bienaymé trees (throwing away the degenerate-case that produces 1 child almost surely), a Bienaymé(Z) tree is infinite with positive probability if and only if $\mu(\hat{Z}) > 1$.

But why are all of these results true, and how are they related?

Philosophy

The philosophy is that the spectral measure of a random graph is *local*. The reason why is that you can solve for the eigenvalues of the adjacency of a finite graph locally.

I've been thinking a lot about local and universal results recently. A notion of convergence is basically as useful as what properties are preserved under it. Sometimes, you want to understand things by understanding what universality class they live in. You might want to use Donsker's theorem to understand centred random walks, Aldous' theorem to understand the convergence of critical Bienaymé trees (conditioned to have size n) to the Brownian continuum random tree. You might want to understand the conformally-invariant scaling limits of 2-dimensional percolation interfaces. What connects these type of results is that that the local particularities wash out: for example, in the last example, it shouldn't matter whether you study percolation on the triangular lattice or the square lattice. Or, for example, everything looks "degree 3" in the Brownian continuum random tree, but you can take an offspring distribution Z with mean 1 that has $\mathbb{P}(Z = 2) = 0$ and still the sequence (T_n) will converge to the Brownian continuum random tree.

In other cases, though, you really do care about local information, and you should study what is preserved under *local convergence*. Some examples:

1. Take a sequence of vertex-transitive graphs G_n which converge locally to a vertex-transitive graph G. Suppose that $p_c(G_n) < 1$, so these graphs G_n aren't one-dimensional. Then a 2023 theorem of Philip Easo and Tom Hutchcroft says that
$$p_c(G_n) \to p_c(G)$$
as $n \to \infty$. This was Schramm's locality conjecture.

2. Eigenvalues and expected spectral measure. We'll make sense of the following statement: let $(G_n)_{n \geq 1}$ be a sequence of unimodular random graphs (G_n, o_n) converging to a unimodular random graph (G, o). Then the expected spectral measure $\mu_{G_n} \to \mu_G$ as measures, and if you additionally assume that (G_n, o_n) are sofic, then
$$\lim_{n \to \infty} \mu_{G_n}(\{x\}) = \mu_G(\{x\}),$$
that is, you get pointwise convergence of atoms.

What is this local convergence that I'm talking about? Let $\mathcal{G}_*$ be the rooted graphs (G, o), up to rooted isomorphisms. (In other words, I quotient out.) Say that $d((G, o), (H, \rho)) = 2^{-k}$ where k is the largest integer such that $B_G(o, k) \cong B_H(\rho, k)$, meaning the k-neighbourhoods of the root are isomorphic. It turns out that this topology makes $\mathcal{G}_*$ a nice Polish space, and once you have a notion of convergence on any Polish space, you can always define what it means for measures valued on $\mathcal{G}_*$ to converge. To make it more explicit, say a sequence of random rooted graphs (G_n, o_n) converges to (G, o) iff for any finite graphs F,
$$\mathbb{P}(B_{G_n}(o_n, k) \cong F) \to \mathbb{P}(B_G(o, k) \cong F)$$

One of the main insights of this paper is to study things at the level of unimodular random graphs (or equivalently, Borel pmp graphs). These objects arise naturally from many different models in discrete probability. Correctly normalized Bienaymé trees are examples of unimodular random graphs; the cluster of the origin in Bernoulli(p) percolation on a transitive graph also gives rise to a URG. Unimodular random graphs, even non-sofic ones, have a well-defined expected spectral measure; in fact this is not obvious, and one of the main technical results of *Mean quantum percolation* (see Proposition 1.4).