PIMS 2025 – Spectral geometry of Liouville quantum gravity

Lecturer: Professor Nathanaël Berestycki

Notes by: Andrew Lin

June 2025

1 June 24, 2025

This will be a mini-course whose goal is to tell a story without going into too many details in the limited time that we have. We'll begin by telling the story of what "spectral geometry" even means: LQG is a canonical random geometry that's arising naturally in physics, and we want to know what classical theorems of spectral geometry look like in this new context. (That is, we can take any classical theorem and see what happens to it.) So we're curious about what happens to **eigenvalues and eigenfunctions** of LQG, "whether we can **hear the shape** of LQG" (much like Mark Kac's question "can you hear the shape of a drum"), and what connections exist to **quantum chaos**.

Our plan will be to start with some background on Liouville theory (explaining GFF, GMC), then talk about Liouville Brownian motion, which is a canonical diffusion in LQG (defining it, describing its spectrum), do some reminders on spectral geometry (in particular the classical Weyl law) and then move on to the results on spectral geometry. At the end, we'll talk about some conjectures and the potential connections to quantum chaos.

Fact 1

This whole story comes from Polyakov (in 1981) who was quantizing the Liouville Lagrangian, motivated by 2D quantum gravity. Given a 2D Riemannian manifold (Σ, g), we can define for a function $\phi : \Sigma \to \mathbb{R}$ the action

$$S(\phi) = \frac{1}{4\pi} \int_\Sigma \left(|\nabla \phi(z)|^2 + Q R_g(z) \phi(z) + \mu e^{\gamma \phi(z)} \right) d\nu_g(z),$$

where $Q = \frac{\gamma}{2} + \frac{2}{\gamma}$, μ is a cosmlogical constant, and ν_g is the Riemannian volume form.

This integral over the surface has several terms: there's a term which physically corresponds to the Dirichlet energy of ϕ, a term involving the curvature on the surface, and another term which is interesting and problematic involving an exponential of our field. Liouville conformal field theory is then what we get if we use this action to define a measure

$$P(d\phi) = \exp(-S(\phi)) D\phi$$

where $D\phi$ is a "uniform measure on fields." Such a measure does not actually exist – the space is infinite-dimensional – but that's the motivation for the whole area. A rigorous definition of this expression was initiated by Duplantier and Sheffield in 2010, then fully completed by David, Kupiainen, Rhodes, and Vargas in 2016. The initial "toy model" is to let ϕ be the Gaussian free field on a surface Σ or in a domain D, with Dirichlet boundary conditions. ϕ should be thought of here as a "random conformal factor" inducing a new measure: very informally, we can use it to define a distance

$$\text{dist}(a, b) = \inf_\eta \int_0^1 e^{\gamma \phi(\eta(t))} |\eta'(t)| dt$$

where we're minimizing over paths from a to b, and also to define a volume

$$\text{vol}(A) = \int_A e^{\gamma\phi(x)} dx.$$

Here γ should be thought of as a coupling constant, and the picture to keep in mind is that we have a triangulation of a sphere where there are some very small and very big triangles, but all triangles have the same mass (so they look small to Euclidean eyes but for intrinsic geometry they are all the same). Because ϕ is very random and very rough, there will be interesting multi-factor phenomena; in some regions we have very packed behavior where ϕ is quite high, and some are relatively empty where ϕ is quite low (negative).

We'll be a bit more careful now and describe the continuum Gaussian free field, and then we'll describe how to use it to define GMC and Liouville conformal field theory.

Fact 2

Let $D \subset \mathbb{R}^2$ be a bounded domain, and let $G_D(x,y) = \pi \int_0^\infty p_t^D(x,y)dt$ be the continuum Green's function with Dirichlet boundary conditions. (Here $p_t^D(x,y)$ is the transition probability of a Brownian motion from x to y in time t, killed when it leaves the domain, so we end up with "the total amount of time we spend at or near y.) Near the diagonal we have $G_D(x,y) = -\log|x-y| + O(1)$ as $|x-y| \to 0$, but the function is nice away from the diagonal. Thus, we cannot define $(\phi(x))_{x\in D}$ as a centered Gaussian process with $G_D(x,y)$ as covariance function, since $\phi(x)$ must have infinite variance.

Instead, we'll view this as a Gaussian stochastic process, indexed by test functions $f \in \mathcal{D}(D)$ (that is, smooth, compactly supported functions). So we have a collection $((\phi,f))_{f\in\mathcal{D}(D)}$, defined via Kolmogorov's extension theorem and the relation

$$\mathbb{E}[(\phi,f)(\phi,g)] = \int_D \int_D f(x) g_D(x,y) g(y) dxdy.$$

(If ϕ was defined pointwise this is the formula we would naturally write, since we can think formally of having

$$\begin{aligned}\mathbb{E}[(\phi,f)(\phi,g)] &= \mathbb{E}\left[\iint f(x)g(y)\phi(x)\phi(y)dxdy\right] \\ &= \iint f(x)g(y)\mathbb{E}[\phi(x)\phi(y)]dxdy.\end{aligned}$$

So the two-point function of this field, if it were to exist, would be the Green's function, and the point is that the definition makes sense. But now we can extend this definition to rougher functions than test functions: we'll allow ourselves to integrate against $f = f^+ - f^-$ for two nonnegative measures with $\iint G_D(x,y) f^\pm(dx) f^\pm(dy) < \infty$; call the allowed space $\mathcal{M}$. Analytically this is basically saying that the function f must live in the Sobolev space $H^{-1}(D)$.

Example 3

For any test function $f \in \mathcal{D}(D)$, we have $f \in \mathcal{M}$ because

$$\int_D G(x,y) f(x) f(y) dxdy$$

because G blows up logarithmically near $x = y$ and is nice otherwise. And we can integrate the log in two dimensions, so this is okay.

Example 4

Next, let f be the uniform measure on a circle contained in D. We thus want to be able to integrate

$$\int_D G(x,y)\rho(dx)\rho(dy)$$

for ρ uniform on the circle, and again this is allowed because for a fixed x we are trying to integrate log in one dimension, and again this is okay. The same works for a typical one-dimensional smooth curve.

Example 5

However, let $f = \delta_{x_0}(\cdot)$. This cannot be integrated because $G_D(x,x)$ is infinite. So the issue has to do with "whether the measure ρ is supported on something that is sufficiently spread out."

It's good to go back to those circle averages in more detail:

Definition 6

For a domain D, take a point z such that $B(z,\varepsilon) \subset D$. The **circle average of ϕ at distance ε from z** (which is ϕ integrated along the circle) will be denoted $\phi_\varepsilon(z)$.

We should think of this as a regularization of the free field at scale ε. We can compute that $\phi_\varepsilon(z)$ is Gaussian (by definition) and that its variance is $\log(\frac{1}{\varepsilon}) + O(1)$, so that we have logarithmic blowup. Indeed,

$$\mathrm{Var}(\phi_\varepsilon(z)) = \int_D G(x,y)\rho(dx)\rho(dy)$$

involves points x, y uniformly on a circle of radius ε, so the typical distance between them is some constant times ε, and thus the value of $G(x,y)$ will typically be $-\log\varepsilon$. It turns out actually that if we reparametrize at the scale

$$B_t = \phi_{e^{-t}}(z),$$

then $(B_t)_{t\geq t_0}$ will actually be a one-dimensional Brownian motion. Indeed, this comes down to the domain property of the Gaussian free field (so we have independent and stationary increments); another way to see this is to try to get it from the formula directly. Indeed, we want to show that the following quantity is the same as it is for Brownian motion, which is $s \wedge t$:

$$\mathrm{Cov}(B_s, B_t) = \iint G_D(x,y)\rho_{e^{-s}}(dx)\rho_{e^{-t}}(dy).$$

By Fubini's theorem, we can "freeze y" in this expression and integrate over x. This Green's function is a harmonic function of x, so the circle average of $G_D(x,y)$ is equal to the value at the center by the mean value theorem, and thus it's also equal to the average on any other circle around the center, in particular that of e^{-t}: that's how we get the $s \wedge t$. So that's nice and explains why we can't define the GFF pointwise – as distances get closer and closer, the circle average fluctuates more and more between $+\infty$ and $-\infty$. The miraculous thing though is that we can still integrate against test functions and get something sensible with a finite value!

Definition 7

The **thick points** are exceptional points of the GFF; we say a point is **α-thick** if

$$\lim_{\varepsilon\to 0} \frac{\phi_\varepsilon(z)}{\log(1/\varepsilon)} = \alpha.$$

This is exceptional because we know that circle averages should behave like Brownian motion, and thus being α-thick corresponds to having $\lim_{t\to\infty} \frac{B_t}{t} = \alpha$ (since we have $\varepsilon = e^{-t} \iff t = \log\frac{1}{\varepsilon}$). So for any fixed $\alpha > 0$ and any fixed point $z_0 \in D$, the point will not be α-thick almost surely. Nevertheless, there could still be some exceptional points for which the Brownian motion behaves in an unusual manner and has this positive drift:

Fact 8
The Hausdorff dimension of the set of thick points $\mathcal{T}_\alpha$ is $\max(0, 2 - \frac{\alpha^2}{2})$. Furthermore, the set is nonempty if and only if $\alpha \leq 2$.

Remark 9. *If we're concerned about measurability issues, it turns out that there is a version of the GFF such that all circle averages simultaneously exist; much like how the Kolmogorov continuity theorem can be used to do things for Brownian motion, we can do something for the construction here.*

It's useful to quickly explain where this constant for the Hausdorff dimension comes frmo, since it explains a lot about Gaussian multiplicative chaos. We'll just do a general "plausibility argument" for this: suppose our domain is a square, and we break it up in a grid with mesh size ε. For each point on the grid, we then ask the question of "whether $\phi_\varepsilon(z) \approx \alpha \log\frac{1}{\varepsilon}$." We can count the number of such points, and to say that the set of thick points has dimension $2 - \frac{\alpha^2}{2}$ is basically the same as saying that the number of such ε-points is $\varepsilon^{-(2-\alpha^2/2)}$. This is not so difficult to believe, since we can compute the expected number of such points – via the Gaussian tail (which is estimated by its density), each of the ε^{-2} points in the mesh has probability $e^{\alpha^2/2}$ to be this exceptionally big, since

$$\exp\left(-\frac{1}{2}\frac{(\alpha \log 1/\varepsilon)^2}{\log 1/\varepsilon}\right) = \exp\left((\log\varepsilon)\cdot\frac{\alpha^2}{2}\right).$$

So in particular if we pick $\alpha > 2$ we expect an exponentially small number of thick points and it makes sense to not have any of them.

Example 10
Gaussian multiplicative chaos was first introduced by Kahane in the 1980s – he was motivated by ideas of turbulence (coming from work of Kolmogorov and Mandelbrot in the 1970s). We'll take a GFF on the domain D (or more generally a log-correlated Gaussian field on $\mathbb{R}^d$). Let ϕ_ε be some regularization of ϕ such as the circle average or $\phi * \theta_\varepsilon$ for some compactly supported convolution kernel θ. We then consider

$$M_\varepsilon(dx) = \varepsilon^{\gamma^2/2} e^{\gamma\phi_\varepsilon(x)} dx.$$

The reason there's a problem with just taking something like $M(dx) = \exp(\gamma\phi(x))$ dx is that ϕ is a distribution and thus taking nonlinear functions of it doesn't make any sense a priori. And a more probabilistic picture is the following: the GFF is pointwise not defined, and the miracle is that if we integrate against test functions the $+\infty$ and $-\infty$ will cancel. But if we exponentiate, the cancellation now means nothing, and thus we need to do some appropriate regularization and normalization. This factor $\varepsilon^{\gamma^2/2}$ turns out to be the right factor:

Theorem 11 (Kahane '85, Berestycki '17, Shamov '17)
For $0 < \gamma < 2$, $\lim_{\varepsilon\to 0} M_\varepsilon(dx)$ exists in probability with respect to the weak topology, and the limiting GMC is nonzero if and only if $\gamma < 2$ (or more generally $\gamma < \sqrt{2d}$). Furthermore, this limit is universal in that it doesn't depend on the regularization θ up to some multiplicative constant.

Remark 12. *Possibly to explain the name GMC, there's a version of Gaussian multiplicative chaos which takes place on a binary tree where we multiply independent variables along the paths of the tree. And for the infinite dyadic tree encoding the interval* $[0,1]$*, what this does is basically the same as the theory we'll present soon, but it's "multiplicative." Somehow the GFF is a sum over many different scales, and so when we exponentiate we're multiplying those contributions together. It was initially thought that we'd get a theory which would reflect the central limit theorem but in multiplicative form, but in reality this theory is much more complicated.*

Fact 13

At the critical value $\gamma = 2$, if we just do what we described then we end up with a trivial limit. But there is a way of doing a slightly better regularization (multiplying by some additional factor), which yields the critical GMC. But beyond that we do not get a nontrivial limit in probability, only in law (so we can't fix a realization of our field and associate to it a realization of GMC).

The normalization here is actually not so bad to understand: we have

$$M_\varepsilon(A) = \varepsilon^{\gamma^2/2} \int_A e^{\gamma \phi_\varepsilon(z)} dz.$$

We've chosen this so that the expectation can be written

$$\mathbb{E}[M_\varepsilon(A)] = \varepsilon^{\gamma^2/2} \int_A \mathbb{E}\left[e^{\gamma \phi_\varepsilon(z)}\right] dz,$$

and if we want to compute $e^{\gamma \phi_\varepsilon(z)}$ it's like computing the Laplace transform of a Gaussian: it's $\exp\left(\frac{1}{2}\gamma^2 \mathrm{Var}(\phi_\varepsilon(z))\right)$, and the variance is basically $\log \frac{1}{\varepsilon}$ plus a constant. So we end up with something exactly of order 1, and thus expectations actually have a chance of converging this way.

There's a number of intermediate results leading to this – all of them gave existence and uniqueness in some restricted sense. In particular, Professor Berestycki's paper is pretty short and has a pretty quick argument which should be readable.

Fact 14

If we plot the "density of the GMC" as we vary γ (in reality the limiting measure is singular with respect to Lebesgue measure, though), for γ small we should have something close to Lebesgue measure, and the more we increase γ the more the measure peaks on a small number of points. That's somehow because $\varepsilon^{\gamma^2/2}$ becomes smaller, so we need ϕ to be bigger and bigger. So we require more and more exceptional behavior – in fact the behavior of the field required is that we need to have thick points.

We can in fact show that indeed, M is almost surely supported on only the γ-thick points, meaning $M(\mathcal{T}_\gamma^c) = 0$ almost surely. And since sampling from the measure M yields almost surely γ-thick points, and we can only have at most 2-thick points, it makes sense that we can only get this to work up until $\gamma = 2$. (The idea is that if we sample from M, then $\frac{\phi_\varepsilon(z)}{\log 1/\varepsilon}$ will be exactly equal to γ almost surely, and thus these points are atypical for Lebesgue measure but typical for GMC.)

2 June 26, 2025

Last time, we talked about the Gaussian free field and defined GMC and how it's something that is supported on the set of thick points of the GFF. (In fact, we can think of the multiplicative chaos as a measure which is uniform on

those thick points.) We'll now take a moment to explain how we can use this to get a meaning of Polyakov's formula for $S(\phi)$ – recall that Liouville conformal field theory is supposed to define the measure $\exp(-S(\phi))D\phi$. To assign a meaning to this, the idea is to go back to the constructive field theory of Glimm and Jaffe from the 1970s and define

$$\mathbb{P}^{\mathrm{GFF}}(d\phi) = \exp\left(-\int_\Sigma |\nabla\phi|^2 d\nu_g\right) D\phi$$

as the law of the GFF. From here, the rest is a Radon-Nikodym derivative reweighting the GFF – oversimplifying the work of David, Kupiainen, Rhodes, and Vargas, they define it to be the reweighted GFF via the factor $\exp\left(-\int_\Sigma QR_g(z)\phi(Z)\nu_g(dz) - \mu M(\Sigma)\right)$ (which has a curvature term and then an additional term where $M(\Sigma)$ is the total mass of the GMC measure). But if Σ is something like the sphere, then there are no boundary conditions (while we usually want zero on the boundary) and some additional complications come up. But the nice thing is that we can compute the correlation functions by some exact, explicit (but complicated) formulas called the **DOZZ formula** (by Kupiainen, Rhodes, and Vargas). And there are close connections to random planar maps like in a previous mini-course.

Fact 15

We'll stick to the toy model, though – we won't worry about the Polyakov measure and will just work with the toy model $\mathbb{P}^{\mathrm{GFF}}$ for the rest of this course on some domain $D \subset \mathbb{R}^2$. We think of it as endowing the domain D with a random geometry with distance and volume, and GMC is then playing the role of the "uniform volume measure."

Example 16

We can now turn to **Liouville Brownian motion**. Recall that one of the things we want to do in this course is study the spectral geometry of LQG, meaning we care about the eigenvalues and eigenfunctions. These objects are associated to an operator; if we were doing geometry we would want this to be the Laplace-Beltrami operator, but it turns out to be much more natural to take a probabilistic approach and define the diffusion instead.

If we were doing geometry and wanted to define Brownian motion on a manifold, normally we'd say that the manifold structure allows us to differentiate functions, and thus we can define the Laplacian and the Brownian motion will be a stochastic process on the manifold whose generator is that Laplacian. In this random context this will not work because things are too rough to differentiate, and there's something we can do instead.

Theorem 17 (Berestycki '15, Garban-Rhodes-Vargas '15)

There is a nontrivial limiting object $Z(t) = \lim_{\varepsilon\to 0} Z_\varepsilon(t)$, where Z_ε is a time-change of Brownian motion

$$Z_\varepsilon(t) = B_{F_\varepsilon^{-1}(t)}, \quad F_\varepsilon(t) = \varepsilon^{\gamma^2/2}\int_0^t e^{\gamma\phi_\varepsilon(B_s)}ds.$$

This timechange basically measures how much our Brownian motion has seen so far – we can imagine having a rough "mountain" landscape that the GFF gives us, and then the trajectory spends much more time on the top of the mountains because it slows down (not because it's attracted – the process is completely isotropic). And we know that this is the right thing to do by conformal invariance of Brownian motion. (We can equivalently say that Z is a time-change of Brownian motion with PCAF having Revuz measure exactly M.) The Dirichlet form description of this Liouville Brownian motion is

$$\mathcal{E}(f,g) = \int_D \nabla f(x)\cdot\nabla g(x)dx,$$

where we're doing things with respect to $L^2(M)$ rather than the usual measure.

Fact 18

The picture to keep in mind is the scaling limit of simple random walk on a triangulation. The idea is that regions where the field is high has many triangles packed, so it takes a long time to escape those regions. And Professor Berestycki showed with Gwynne in 2020 that such simple random walk does indeed converge to Liouville Brownian motion, so again we do know that we have the correct object.

We do know some important properties of this process:

- The LBM $(Z(t))_{t\leq\tau_D}$ is almost surely continuous and does not stay stuck (if and only if $0 \leq \gamma < 2$). Indeed, we have a weird, but continuous time-change – it could be hypothetically that we slow down at the top of a mountain for so long that we stay there forever, but in fact this does not occur. Likewise, it could be that the bottom of the valley is so deep that the process ends up jumping over in the limit, and again we can show that this does not occur. (In some sense for $\gamma > 2$ that's what happens.)
- Almost surely, we can start this LBM simultaneously from all points given a single ϕ. This is interesting because it yields, given the GFF, a Feller process (since we can start from a random process), and it leaves the GMC measure invariant (for example if we're defining things on the sphere or torus with no boundary). M is essentially taking the role of the invariant measure on SRW of a graph.
- For each $t \geq 0$, almost surely $Z(t)$ will be at a thick point $\mathcal{T}_\gamma$, so the LBM only sees thick points in some sense.
- Perhaps surprisingly, almost surely the map $t \mapsto Z(t)$ is differentiable almost everywhere for $\gamma > \sqrt{2}$, and in fact $Z'(t) = 0$. (The idea is that "the process spends a lot of time being near the top of mountains, so we have a Cantor-type time change.)
- The heat kernel exists – that is, for a starting point x, we have $P_t(x, \cdot) \ll M$, so we have a Radon-Nikodym derivative with respect to the GMC. But do note that this depends on the realization of our GFF and so P is actually random.

Example 19

We'll now move towards spectral geometry, and to do so we need to remind ourselves about the classical story. There's this story asked in the 1960s of "can we hear the shape of a drum?," but the first referene to that question came from a paper of the physicist Schuster in 1882, who wrote that "it would baffle the most skillful mathematician to find out the shape of a bell by means of the sound which it is capable of sending out."

It took a while for the question to really catch on: Lorentz visited Göttingen in 1910, and in a lecture he gave he said the following. "In an enclosure with a perfectly reflecting surface there can form standing electromagnetic waves, analogous to tones of an organ pipe. We shall confine our attention to very high overtones [frequencies]... there arises the mathematical problem to prove that the number of overtones [eigenvalues] which lie between frequencies $\nu, \nu + d\nu$ is independent of the shape of the enclosure and is simply proportional to its volume.. it has been verified for many simple shapes... there is no doubt that the theorem holds in general." Hilbert was in the room when this seminar was given, and he said that it wouldn't be solved in his lifetime. But actually it was solved just two years later by his own student Weyl, using methods introduced by Hilbert:

Theorem 20 (Weyl's law, 1912, in two dimensions)
Let D be a bounded domain, and consider $N(\lambda) = \sum_{n \geq 1} 1\{\lambda_n \leq \lambda\}$, the eigenvalue counting function (for the Laplacian with zero boundary conditions) – that is, we want the solutions to

$$-\frac{1}{2}\Delta f = \lambda f, \quad f = 0 \text{ on } \partial D.$$

Then

$$\frac{N(\lambda)}{\lambda} \to \frac{1}{2\pi}\text{Leb}(D).$$

The same result holds on any 2D Riemannian manifold – we get the same constant again, regardless of the metric we put on the manifold. And this also works with appropriate adjustments for higher dimensions as well.

We're curious what these results become when applied to LQG, but first we'll explain this classical problem first. Kac asked the following in 1966 (motivated by Weyl's law): **given a sequence of eigenvalues, do they determine uniquely the domain up to isometry**? The answer turns out to be no – counterexamples were constructed first by Milnor in high dimensions, then Gordon, Webb, and Wolpert in two dimensions in 1992. That latter paper was important because it constructed counterexamples in interesting ways with connections to Riemann surfaces.

But turning back to LQG, we can in fact say interesting things. First of all, we need to figure out how to define the spectrum in the first place – it might seem like it's not so straightforward to define eigenvalues and eigenfunctions. The answer (mentioned in papers of Andres and Kajino, as well as Maillard, Rhodes, Vargas, and Zeitouni in 2016) is that instead of considering the generator itself, we should work instead with its inverse, the **Green's function**. Formally, our Green's function $G^\gamma(x, y)$ should encode "how much time a Liouville Brownian motion starting at x spends near y before leaving the domain," and it's not so complicated because we have a time-change of Brownian motion, and we just count the visits of a usual Brownian motion, reweighted accordingly. That time-change is always the same:

$$G^\gamma(x, y) = e^{\gamma\phi(y)}G^0(x, y).$$

We do have to understand this appropriately – instead of actually doing this formally, we should instead write the Green operator as

$$G^\gamma(x, dy) = G^0(x, y)M(dy).$$

This is a nice observation to make, because we know everything about the function $G^0(x, y)$ (for example, log-divergence on the diagonal and otherwise very nice). So it turns out that $G(x, dy)$ will be almost surely a nice compact operator on $L^2(M)$, meaning that we can apply the spectral theorem to it. This yields, for free, that we get existence of an orthonormal basis of eigenfunctions $\{f_n\}$ with respect to $L^2(M)$ with associated eigenvalues λ_n, such that

$$Gf_n(x) = \frac{1}{\lambda_n}f_n(x).$$

(Remember that these eigenvalues and eigenfunctions are random.) So there's no essential difficulty – we have a nice function for the Green operator and that lets us do everything – and now we can define again a (random) eigenvalue counting function $N(\lambda) = \sum_n 1\{\lambda_n \leq \lambda\}$.

Theorem 21 (Berestycki, Wong '23)

Take ϕ a Dirichlet GFF in a domain $D \subset \mathbb{R}^2$ and $\gamma \in (0,2)$. Then

$$\frac{N(\lambda)}{\lambda} \to c_\gamma M(D),$$

where $M(D)$ is now a random volume and where

$$c_\gamma = \frac{1}{\pi}\left(\mathbb{E}\left[\int_0^\infty \mathcal{I}(e^{\gamma(B_t - \alpha t)})dt\right] + \mathbb{E}\left[\int_0^\infty \mathcal{I}(e^{-\gamma(\mathcal{B}^\alpha_t)})dt\right]\right), \quad \mathcal{I}(x) = xe^{-x},$$

where $\alpha = \frac{2}{\gamma} - \frac{\gamma}{2}$ and $\mathcal{B}^\alpha$ is a Brownian motion with drift α conditioned to stay positive. In fact this constant can be computed explicitly: we find that $c_\gamma = \frac{1}{\pi(2-\gamma^2/2}$.

So notice that this constant c_γ is equal to the usual $\frac{1}{2\pi}$ at $\gamma = 0$, and $c_\gamma > c_0$ for all $\gamma \in (0,2)$. (It's not clear if there's a good conceptual reason for this, though.) Interestingly, in numerical simulation we do see that the number of eigenvalues grows linearly at a higher slope than what we'd expect in Riemannian asymptotics, but if we zoom in towards the "beginning" it looks like Riemannian does better at first. Intuitively, the explanation may be that when we look at relatively low-level eigenfunctions and eigenvalues, they're determined by the low-level modes of the GFF, where it's not yet very multifractal.

We won't do the full proof, but we would get started on something like this by starting with the **trace formula**: we write down the spectral decomposition of the Liouville heat kernel

$$p_t(x,y) = \sum_{n=1}^\infty e^{-\lambda_n t} f_n(x) f_n(y)$$

(this is true for any Markov chain on a finite set, for example, and in much more generality). We can now observe that this involves the parameters x, y, t and is in fact almost surely jointly continuous in all three arguments (which is not obvious from the Radon-Nikodym derivative definition before!). Because we have a continuous version, it does mean something to plug in a particular value of y, and we'll plug in $y = x$ and then integrate over $x \in D$. We get that

$$H(t) = \boxed{\int_D p_t(x,x)\mu(dx)} = \sum_{n=1}^\infty e^{-\lambda_n t} = \boxed{\int_0^\infty e^{-\lambda t} dN(\lambda)}$$

(since f_n was orthonormal in L^2 by choice, so $\int f_n^2 = 1$). The left expression is called the **heat trace** (it's like the trace of the heat kernel viewed as a matrix), and the right expression involves the eigenvalue counting function for LQG $N(\lambda)$.

So the heat trace is the Laplace transform of the eigenvalue counting function, which is super helpful – recalling Karamata's Tauberian theorem, understanding asymptotics of $N(\lambda)$ means we just need asymptotics of $H(t)$! So what we need to do is show that for any open set $A \subset D$, we have convergence in probability

$$t\int_A p_t(x,x)M(dx) \overset{t\to 0}{\to} c_\gamma M(A).$$

Normally it would be very difficult to work with the heat kernel $p_t(x,y)$ – the tools developed in that area really only give us two-sided Gaussian-type bounds, but we really need the constant in front. Luckily, special to LQG we have the **bridge decomposition**

$$\int_0^\infty g(t)p_t(x,y)dt = \int_0^\infty \mathbb{E}_{x\to y;t}\left[g(F(t))\right]\tilde{p}_t(x,y)dt$$

for $F(t)$ the random timechange (this is the PCAF, also called the **quantum clock**) and $\tilde{p}_t$ the ordinary heat kernel

(transition probability of ordinary Brownian motion killed at D). Here, the expectation $\mathbb{E}_{x\to y;t}$ means we're looking at a Brownian bridge from x to y of length t, hence the name of the decomposition. So even if we can't understand p_t at a fixed time t, we can take care of integrals over time, since by nature we have a timechange of ordinary Brownian motion.

So once we establish that $t\int_A p_t(x,x)M(dx) \overset{t\to 0}{\to} c_\gamma M(A)$ for any open set A, we might ask whether we need to integrate at all – is it true pointwise that $p_t(x,x) \sim \frac{c_\gamma}{t}$ as $t \to 0$, M-almost-everywhere? The "con" of framing the question this way is that $p_t(x,x)$ is a random variable which depends on how the Gaussian free field looks in the vicinity of x, and that varies a lot based on the point and so we might have some multifractal geometry story going on here – it's not clear why we'd have the same constant c_γ everywhere. But the "pro" is that framing the question this way is that it lets us restrict to typical points, and so maybe we tend to see the same type of landscape around us.

The answer turns out to be **no**, even when we restrict to typical points. Indeed, what we do is sample x from $M(dx)$ first, then use the limiting theorem for $p_T(x,x)$ when viewed in "annealed asymptotics" (that is, averaged over only the randomness of the GFF).

Theorem 22 (Berestycki, Klein '25)
We have $tp_t(x,x)$ converging in law (annealed) to a nontrivial random variable X as $t \to 0$. In fact, for any two points x, y, we have $(tp_t(x,x), tp_t(y,y)) \to (X,Y)$ independent realizations of this random variable.

So the picture here is subtle and quite interesting: the theorem says that the law of the random variable has a nontrivial limit, but if we fix a realization of the GFF and fix the point x and ask what $p_t(x,x)$ actually looks like, this picture suggests that it will not look like $\frac{c_\gamma}{t}$; instead we get some fluctuations where it'll look like $\frac{\log t}{t}$ or $\frac{1}{t\log t}$. But when we look at many points at the same time and do integrated asymptotics, all of those fluctuations will then self-average out and we will only see the $\frac{1}{t}$ overall dependence, with c_γ basically being the expectation of the random variable! So it's rather fortunate that this cancellation turns out to happen so that the Weyl law works.

3 June 27, 2025

We'll start today by discussing some ongoing work with Jacob Klein about the "second term" of Weyl's law. Weyl conjectured in the classical setting back in 1912 that the eigenvalue counting function actually satisfies

$$N(\lambda) = c_0\lambda \mathrm{Leb}(D) - c_0'\sqrt{\lambda}|\partial D| + o(\sqrt{\lambda})$$

for a smooth domain. Amazingly, this is actually still open – Ivrii found in 1981 that this is true subject to an ergodic assumption about periodic geodesics (that is, if billiards don't have too many of them), but that's hard to verify in practice. What is known is the corresponding heat trace expansion, though: we have

$$H(t) = \int_D p_t(x,x)dx = c_0\frac{|D|}{t} - c_0'\frac{|\partial D|}{\sqrt{t}} + \cdots,$$

so in particular the size of the boundary $|\partial D|$ is "spectrally determined" and many more terms are actually known. (Unfortunately Tauberian theorems don't get us second-order asymptotics of one quantity from the other, since there is no monotonicity after subtracting off the first-order term.)

We're interested in understanding how this translates over to the LQG world: we define the **anomalous heat trace expansion** in the following way. Let ϕ be Sheffield's γ-quantum cone restricted to some bounded smooth domain D.

Informally, we can think of this as being

$$\phi = \mathrm{GFF}_{\mathbb{C}}(z) + \gamma \log(1/|z|).$$

(The idea is that just having a zero-boundary GFF would not be very interesting, and this way we get something nontrivial.) Such objects determine the scaling limit of whole-plane models, or equivalently local limits when we zoom in. We then want to understand the heat trace expansion:

Theorem 23 (Berestycki, Klein '25+)
Let $H(t) = \int_D p_t^D(x,x)M(dx)$ be the LQG heat trace on D for some smooth bounded domain. Then the average over the realization of the GFF is

$$\mathbb{E}[H(t)] = \frac{c_\gamma \mathbb{E}[M(D)]}{t} - t^{-1+b(\gamma)+o(1)}, \quad b(\gamma) = \frac{1}{2} + \frac{2}{\gamma^2}\left(\sqrt{1+\frac{\gamma^4}{16}} - 1\right).$$

In particular, this is different from what's conjectured in the standard Euclidean case. We have this expression, and thus it makes sense to conjecture that the analogous expansion will also hold for the eigenvalue counting function as well. (And simulations do suggest that fluctuations between the eigenvalue counting function and the linear slope are much smaller than what we see in the Euclidean case, which is constant with $b(\gamma) > \frac{1}{2}$.)

To get started on proving something like this, we write

$$H(t) = \int_D p_t^D(x,x)M(dx) = \int_D p_t^{\mathbb{C}}(x,x)M(dx) - \int_D p_t^{\mathbb{C}\setminus D}(x,x)M(dx),$$

since as $t \to 0$ and x is a point well inside the domain, the Brownian motion from x to x will generally not realize that it's killed when leaving at D. Now what's nice is that the first term is exactly the area term, since we have exact scale-invariance of the heat kernel on Sheffield's quantum cone – the second term then comes from terms near the boundary, since the Brownian motion has to "manage to come out of the domain" and therefore only x right near the boundary have a chance of contributing. It turns out that the dominant behavior actually occurs from points with atypical thickness (even for LQG)

$$\alpha = Q - \sqrt{Q^2 - 2}, \quad Q = \frac{2}{\gamma} + \frac{\gamma}{2}.$$

instead of γ. (Here $\alpha < \gamma$.) There is in fact a very good reason for this because we have connections with the KPZ scaling: this is an intrinsic fact perhaps because of the following:

Conjecture 24
With high probability we have (we've removed the expectation now) $H(t) = \frac{c_\gamma M(D)}{t} - t^{-1+\Delta+o(1)}$, for $\Delta \in [0,1]$ the **quantum scaling exponent** of ∂D.

Intuitively, the quantum scaling exponent is such that the Hausdroff dimensions satisfy $\frac{\dim_\gamma(\partial D)}{\dim_\gamma(D)} = 1 - \Delta$, and the KPZ scaling relation (this is a different KPZ from Kardar-Parisi-Zhang) says that the Euclidean scaling exponent is

$$x = \frac{\gamma^2}{4}\Delta^2 + \left(1 - \frac{\gamma^4}{4}\right)\Delta.$$

So if our set A is the Euclidean boundary of our domain D, then x should be $\frac{1}{2}$ (since the Hausdorff dimension of a smooth domain is 1 and the interior is 2, so $\Delta = \frac{1}{2}$, and then $\gamma = 0$.) So the value we had above is exactly solving the quadratic equation for Δ with $x = \frac{1}{2}$ in general, and it predicts the second term in the expansion not just for smooth

domains but also more complicated objects. We can think of the KPZ equation here as "relating the random and deterministic dimensions."

Example 25

It's natural now to consider the KPZ relation on quantum disks. We can think of the quantum disk as a scaling limit of planar maps with the disk topology – it turns out that the boundary of such objects is like SLE_κ with $\kappa = \gamma^2 \in (0,4)$ instead of a smooth curve.

If the conjecture from before is to be believed, we can use the known Hausdorff dimension of such a curve (which is $1 + \frac{\kappa}{8}$, due to Beffara). So then we can compute the associated scaling exponent, solve the KPZ equation, and then get a prediction for the second term in the Weyl law in general: we get $x = \frac{1}{2} - \frac{\kappa}{16}$, and then actually we get $\Delta = \frac{1}{2}$ independently of γ:

Theorem 26 (Berestycki, Klein '25+)

For a quantum disk (D, ϕ), we have

$$\mathbb{E}[H(t)] = \frac{c_\gamma \mathbb{E}[M(D)]}{t} - t^{-1/2+o(1)}.$$

In particular, we're very sensitive to the boundary conditions and geometry once we start looking at this second term.

To understand where this KPZ conjecture comes from in more detail, we can let d_γ denote the Hausdorff dimension for the metric space associated to LQG_γ. (Recall that we do have a random metric associated to this surface..) Even heuristically it's actually unknown what d_γ is apart for $\gamma = \sqrt{8/3}$ – there we believe and know in some sense that the dimension corresponds to the the volume growth of planar maps, so $d_\gamma = 4$. For a long time there was Watabiki's prediction (yielding a formula for any γ), but unfortunately it's been disproved now. It's expected that we have for the Liouville Brownian motion Z_t that

$$\mathrm{dist}_\gamma(Z_t, Z_0) \approx t^{1/d_\gamma} \text{ as } t \to 0.$$

We can thus cover the boundary of our domain with balls of radius $r = t^{1/d_\gamma}$; this will require $r^{-\dim_\gamma(\partial D)} = t^{-(1-\Delta)}$ balls in total by definition of Δ. (These points within the balls are the ones that should contribute significantly to the second term of the heat trace, since they're the only ones with a good chance of escaping.) So the integral we want to estimate looks like

$$\int p_t^{\mathbb{C}\setminus D}(x,x) M(dx) = \int_{\text{balls}} \frac{1}{t} M(dx) = \frac{1}{t} \cdot t^{-(1-\Delta)} \cdot (\text{mass of each ball});$$

because the Hausdorff dimension is d_γ, the mass will be (this is also how the volume grows) $r^{d_\gamma} = t$. Therefore this whole expression indeed simplifies to $t^{-(1-\Delta)}$, and what's remarkable is that it's all in terms of ratios of Hausdorff dimensions (the dimensions themselves don't actually show up anywhere).

There are natural questions we can ask from here, but we'll start with some conjectures instead.

Conjecture 27 (Berestycki, Wong '23)

One "can hear the shape of LQG," meaning that given the sequence of eigenvalues and being told γ, we can recover the field ϕ almost surely (so we have a measurable function of the $\{\lambda_n\}$s).

Implicitly when we make such a conjecture, we're assuming that we need the domain. But maybe we can ask whether even the domain can be recovered as well – we know that it's not possible for $\gamma = 0$, and of course we can

only hope to recover up to the equivalence relation of isometry. But the conjecture is actually that we can determine (D, ϕ) up to this equivalence! (And we've been thinking that we implicitly are given γ, but it's also interesting to think about whether we can recover γ from the eigenvalues.)

Problem 28

Next, we can ask about localization: should eigenfunctions look like in Brownian motion where they are delocalized, or should they be like the Anderson model in which they are localized?

For example, the eigenfunctions of the Laplacian in the unit square are sine waves, and a high eigenvalue corresponds to fast-oscillating waves. Thus the L^2 mass is spread out all over the domain and we have a delocalized behavior. On the other hand, if the L^2 energy comes from a tiny portion of the domain, then we have localized behavior, and we know this is happening in the Anderson model in which we look not at the Laplacian but at $-\Delta + \varepsilon W$ for W white noise. It's a famous physics prediction that the eigenfunctions of this latter object are localized no matter how small the noise – intuitively, there's a lot of impurities and thus the randomness makes it very hard for a function to be an eigenfunction unless we make an effort in some tiny region and then just let the function be basically zero everywhere else.

Simulations suggest that the lower eigenfunctions look somewhat localized, but the higher eigenfunctions are more delocalized.

Conjecture 29

Eigenfunctions of LQG are delocalized, and in fact $|f_n(x)|^2 M(dx)$ approximates the GMC measure $M(dx)$.

The connection here has to do with **quantum chaos** – this is a story that makes sense for a deterministic surface or manifold as well. The idea is to manifest ergodicity of the geodesic flow at a quantum level: assuming that we have a hyperbolic surface with a geodesic flow (given position and an initial velocity, follow a geodesic in that direction), an initial error in the velocity or position will result in geodesics that diverge very quickly. Thus the flow will be ergodic, meaning that the position of the particle will be very quickly uniformly distributed. So what we're saying is that we have the manifestation of this at the quantum level, in which we take a quantum particle whose energy is associated to the nth eigenvalue. The position of that particle will be given precisely by $|f_n(x)|^2$ by quantum mechanics, and so ergodicity means that at high levels we should be equally likely to be anywhere:

$$|f_n(x)|^2 \overset{n\to\infty}{\to} \nu_g(dx).$$

This has already been proven for hyperbolic surfaces up to a dense subsequence (due to Shnirelman in 1974 and Zelditch 1987, then generalized by Colin de Verdière 1985), meaning that over the first N integers the proportion of terms in the subsequence converges to 1. There's a famous conjecture of Rudnick and Sarnak from 1994 that we have quantum **unique** ergodicity, meaning that we don't need to restrict to a subsequence, and Lindenstrauss got the Fields medal for proving this for a certain class of **arithmetic surfaces** (related to number theory). So this is a very important topic!

The conjecture for this problem is that this is exactly what happens in LQG, and that's because Polyakov's action $S(\phi)$ (defining LQG) is actually minimized for ϕ having constant negative curvature. Thus we're actually a hyperbolic surface to first order (with some perturbation around it), so we should expect quantum chaos to be observed for LQG too. (And when we send $\gamma \to 0$ we can indeed prove something about this kind of behavior, due to Lacoin, Rhodes, and Vargas.)

Problem 30

Relatedly, we also have the question of **random waves** – in the usual geometric sense (even for a deterministic domain and the usual Laplacian), the local behavior of eigenfunctions is predicted to converge to Berry's random wave model, which is a Gaussian random field b in the plane with covariance function

$$\mathbb{E}[b(x)b(y)] = J_0(||x-y||), \quad J_0 \text{ the Bessel function of the first kind.}$$

(So the randomness in the limit is just from choosing our random point.)

The conjecture is that the same will happen locally for LQG eigenfunctions (to Berry's random wave, up to some local scale change, since it seems like at some points we are always more dense than others). And there's an interesting "filament structure" (which is poorly understood even heuristically) – the points where the eigenfunctions are reasonably big are being carried by lines. There's been various definitions that have been proposed, but often the formal definitions that are put forward have been shown to not occur.

Finally, one more aspect of quantum chaos has to do with eigenvalue spacing:

Conjecture 31

Under delocalization, the eigenvalues will repel each other, and in fact the eigenvalues satisfy

$$\text{eigenvalue spacing} \to F_{\text{GOE}}(x).$$

(Meanwhile if we have localized eigenfunctions we would expect eigenvalues to form basically a Poisson point process, since they're "probing different parts of the space.") This can be checked numerically for instances of LQG – indeed, the histograms line up quite well with the Wigner surmise rather than an exponential distribution.

In principle, eigenvalues are something that we can define for graphs, and we don't need to know how they're embedded at all. So given a random planar map of size N, we can look at the eigenvalues and likely the spacings should also be given by GOE. Again, the questions about localization do not depend on the embedding at all, and we can ask whether the mass of eigenfunctions for a planar map is concentrated on a small set of vertices (which is a purely combinatorial question). So all of these conjectures have analogs and it's an interesting question to study!