Random Surfaces and Liouville Quantum Gravity

Ewain Gwynne

What is the most natural way of choosing a random surface (two-dimensional Riemannian manifold)? If we are given a finite set X , the easiest way to choose a random element of X is uniformly, i.e., by assigning equal probability to each element of X . More generally, if we are given a set $X \subset \mathbb{R}^n$ with finite, positive Lebesgue measure, the simplest way of choosing a random element of X is by sampling from Lebesgue measure normalized to have total mass one. However, the space of all surfaces is infinite-dimensional for any reasonable notion of dimension, so it is not immediately obvious whether there is a canonical way of choosing a random surface.

Nevertheless, there is a class of canonical models of “random surfaces” called *Liouville quantum gravity (LQG) surfaces*. The reason for the quotation marks is that LQG surfaces are not Riemannian manifolds in the literal sense, since they are too singular to admit a smooth structure.

Ewain Gwynne is a postdoc in mathematics at the University of Cambridge, supported by a Clay Research Fellowship and a Trinity College Junior Research Fellowship. His email address is eg558@cam.ac.uk.

Communicated by Notices Associate Editor Chikako Mese.

*For permission to reprint this article, please contact:
reprint-permission@ams.org.*

DOI: <https://doi.org/10.1090/noti2059>

Instead, LQG surfaces are defined as random topological surfaces equipped with a measure, a metric, and a conformal structure. These surfaces are fractal in the sense that the Hausdorff dimension of an LQG surface, viewed as a metric space equipped with its intrinsic metric, is strictly bigger than two.

LQG surfaces have a rich geometric structure that is still not fully understood (see Section 4). Furthermore, such surfaces are important in statistical mechanics, string theory, and conformal field theory and have deep connections to other mathematical objects such as Schramm–Loewner evolution, random matrix theory, and random planar maps.

Note. Many citations in this article have been omitted since *Notices of the AMS* imposes an upper limit on the number of references in the bibliography. A version of the article with the full reference list can be found on the arXiv; see arXiv:1908.05573.

1. Discrete Random Surfaces

Before we discuss random surfaces, let us first, by way of analogy, consider the simpler problem of finding a canonical way to choose a random curve in the plane. As in the case of surfaces, the space of all planar curves is

infinite-dimensional and does not admit a canonical probability measure in an obvious way. To get around this, we discretize the problem. Let us consider for each $n \in \mathbb{N}$ the set of all nearest-neighbor paths in the integer lattice $\mathbb{Z}^2$ with n steps. This is a finite set (of cardinality 4^n), so we can choose a uniformly random element S_n from it. This discrete random path S_n is called the *simple random walk*.

By linearly interpolating, we may view S_n as a curve from $[0, n]$ to $\mathbb{R}^2$. There is a classical theorem in probability due to Donsker that states that the rescaled random walk paths $t \mapsto n^{-1/2}S_n(nt)$ converge in distribution (with respect to the uniform topology) to a limiting random continuous curve called *Brownian motion*, which can be thought of as the canonical random planar curve. Brownian motion has a much richer structure than an ordinary smooth curve. A Brownian motion curve is nowhere differentiable and crosses itself in every interval of time. It has zero Lebesgue measure but Hausdorff dimension two.

One can define a canonical random surface via a similar approach. Let us first define the discrete random surfaces that we will consider. A *planar map* is a graph (multiple edges and self-loops allowed) embedded into the plane $\mathbb{C}$ in such a way that no two edges cross, viewed modulo orientation-preserving homeomorphisms $\mathbb{C} \rightarrow \mathbb{C}$. Planar maps have vertices, edges, and faces, but these objects do not correspond to particular points/sets in $\mathbb{C}$ since we do not specify a particular embedding. See Figure 1 for an illustration of a planar map.

We can think of a planar map as a discrete random surface where each face is equipped with the flat Riemannian metric coming from a polygon in $\mathbb{C}$ with the appropriate number of edges and unit side length. Equivalently, a planar map is obtained by starting with a collection of polygons with unit side length and identifying pairs of their sides (in such a way that the identification between any two edges is a Euclidean isometry) to produce a surface, subject to the constraint that there are no holes or handles. The surface obtained in this way always has the topology of the sphere, but we can similarly obtain random surfaces with other topologies (we are primarily interested in the local geometry of random surfaces rather than their topology).

For each $n \in \mathbb{N}$, there are only finitely many planar maps with n edges. Hence it makes sense to choose such a map uniformly at random. One can also consider uniform planar maps with local constraints, such as triangulations (resp., quadrangulations), which are required to have three (resp., four) edges on the boundary of each face.

It is also natural to consider planar maps weighted by some sort of additional structure on the map. Indeed, suppose, for example, that we are interested in planar maps M decorated by a *spanning tree*, i.e., a connected subgraph of M that contains every vertex of M and has no cycles. Then

Figure 1. **Left:** A planar map. **Right:** A planar map decorated by a spanning tree.

we might want to sample a uniform random pair (M, T) consisting of a planar map with n edges and a spanning tree on it.¹ In this case, the marginal distribution of M is not uniform; rather, the probability that M is equal to any fixed planar map $\mathfrak{M}$ with n edges is proportional to the number of possible spanning trees of $\mathfrak{M}$. In a similar vein, one might want to look at planar maps sampled with probability proportional to the number of certain types of orientations on the edges of M or the partition function of a statistical mechanics model on M (like the Ising model or the Fortuin–Kasteleyn model).

Since planar maps can be thought of as discrete surfaces, it is natural to expect that random planar maps converge, in some sense, to limiting random surfaces as the total number of edges tends to ∞ . In other words, if we sample a large random planar map and “zoom out” so that we see only its large-scale structure, we should get something that looks like some sort of (continuum) random surface. For a certain class of random planar maps, including uniform planar maps and planar maps sampled with probability proportional to the partition function of a statistical mechanics model with parameters set to the critical values, the particular types of surfaces that arise in this way are LQG surfaces, as discussed above and defined precisely in Section 2. We will discuss the precise sense in which this convergence occurs in Section 3.

2. Liouville Quantum Gravity

2.1. Isothermal coordinates. To define LQG surfaces, we first recall some facts from the theory of deterministic surfaces. Suppose $\mathcal{S}$ is a continuously differentiable surface; i.e., in local coordinates $\mathcal{S}$ can be represented by $E dx^2 + F dx dy + G dy^2$ for some continuously differentiable functions E, F, G of the parameters (x, y) . We will be primarily interested in the local geometry of $\mathcal{S}$, so by possibly replacing $\mathcal{S}$ by an open subset of $\mathcal{S}$, we can assume that $\mathcal{S}$ is homeomorphic to the open unit disk $\mathbb{D} \subset \mathbb{R}^2$. A standard theorem in Riemannian geometry asserts that,

¹There is some ambiguity when counting pairs (M, T) , since there can be automorphisms $M \rightarrow M$ that do not fix T . In practice, this ambiguity is removed by specifying a distinguished edge e of M : there are no nontrivial automorphisms of M that fix e .

at least locally, $\mathcal{S}$ can be parametrized by *isothermal coordinates*. This means that we can parametrize $\mathcal{S}$ by coordinates (x, y) in $\mathbb{D}$ in such a way that the Riemannian metric tensor takes the form $e^{h(z)}(dx^2 + dy^2)$ for some continuously differentiable function $h : \mathbb{D} \rightarrow \mathbb{R}$. Here, $z = x + iy$, and $dx^2 + dy^2$ is the Euclidean Riemannian metric tensor.

It is easy to describe areas and distances with respect to isothermal coordinates. For a Lebesgue measurable set $A \subset \mathbb{D}$, the area of the corresponding subset of the surface $\mathcal{S}$ is given by

$$\int_A e^{h(z)} d^2z, \quad (2.1)$$

where $d^2z = dx dy$ is two-dimensional Lebesgue measure. The $\mathcal{S}$ -distance between any two points in $z, w \in \mathbb{D}$ is given by

$$\inf_{P: z \rightarrow w} \int_a^b e^{h(P(t))/2} |P'(t)| dt, \quad (2.2)$$

where the inf is over all piecewise continuously differentiable paths $P : [a, b] \rightarrow \mathbb{D}$ from z to w .

2.2. The Gaussian free field. We want to define an LQG surface as a random surface parametrized by isothermal coordinates by making a random choice of h . Since LQG surfaces should describe the large-scale behavior of random planar maps, by analogy with the central limit theorem a natural first guess is that h should be a “standard Gaussian random variable” taking values in the space of differentiable functions on $\mathbb{D}$. To explain what this means, suppose that we are given a finite-dimensional Hilbert space $\mathcal{H}$ and let $\{x_1, \dots, x_n\}$ be an orthonormal basis for $\mathcal{H}$. We can define a *standard Gaussian random variable* on $\mathcal{H}$ by

$$\sum_{j=1}^n \alpha_j x_j, \quad (2.3)$$

where the α_j 's are i.i.d. standard Gaussian random variables on $\mathbb{R}$; i.e., they are sampled from the probability measure with density $\frac{1}{\sqrt{2\pi}} e^{-x^2/2}$.

For an open domain $U \subset \mathbb{C}$, consider the infinite-dimensional Hilbert space $\mathcal{H}(U)$ that is the Hilbert space completion of the space of smooth, compactly supported functions on $\mathbb{D}$ with respect to the *Dirichlet inner product*

$$(f, g)_{\nabla} = \frac{1}{2\pi} \int_U \nabla f(z) \cdot \nabla g(z) d^2z, \quad (2.4)$$

where ∇ denotes the gradient and $\cdot$ denotes the dot product. This is sometimes called the *first order Sobolev space* on U , with zero boundary conditions.

The *Gaussian free field (GFF)* is the standard Gaussian random variable on $\mathcal{H}(U)$. That is, let $\{f_j\}_{j \in \mathbb{N}}$ be an orthonormal basis for $\mathcal{H}(U)$. By analogy with (2.3), we

define the GFF h by

$$h := \sum_{j=1}^{\infty} \alpha_j f_j, \quad (2.5)$$

where the α_j 's are i.i.d. standard Gaussian random variables. The sum (2.5) a.s. does not converge pointwise, so the GFF does not have well-defined pointwise values.

However, it is not hard to show that the GFF makes sense as a random distribution (generalized function) on $\mathcal{H}(U)$ [28]. This means that for any $f \in \mathcal{H}(U)$ the Dirichlet inner product $(h, f)_{\nabla}$ and the L^2 inner product $(h, \phi) = \int_U h(z) f(z) d^2z$ are well-defined. The random variables $(h, f)_{\nabla}$ and $(h, g)_{\nabla}$ for $f, g \in \mathcal{H}(U)$ are jointly centered Gaussian with covariances $\text{Cov}((h, f)_{\nabla}, (h, g)_{\nabla}) = (f, g)_{\nabla}$.

Although we will primarily be interested in the two-dimensional case, we remark that the above construction of the GFF also makes sense in other dimensions. In dimension one, one gets a Brownian bridge (a one-dimensional Brownian motion defined on an interval and conditioned to be zero at the endpoints). Hence the GFF can be seen as a generalization of Brownian motion with two time dimensions. In dimension at least three, the GFF is in some sense rougher, i.e., further from being a function, than in dimension two, and hence regularization procedures such as the ones described below do not give interesting objects. Rather, it is more natural to consider so-called *log-correlated* Gaussian fields [10] (the GFF is log-correlated only in dimension two).

2.3. Liouville quantum gravity surfaces. Let h be the GFF on a domain $U \subset \mathbb{C}$ and let $\gamma \in (0, 2)$. Recalling the above discussion on isothermal coordinates, we define the γ -*Liouville quantum gravity (LQG) surface* associated with (U, h) to be the random surface parametrized by U with Riemannian metric tensor “ $e^{\gamma h}(dx^2 + dy^2)$.” This definition does not make literal sense, since h is not a function, so it cannot be exponentiated. However, one can define LQG surfaces rigorously via regularization procedures that we will describe shortly.

The parameter γ controls the “roughness” of the surface: $\gamma = 0$ corresponds to a smooth surface, and larger values of γ correspond to surfaces that are more fractal and less Euclidean-like. The parameter γ is also related to the type of random planar map model under consideration. The case when $\gamma = \sqrt{8/3}$ (sometimes called “pure gravity”) describes the large-scale behavior of uniform planar maps. Other values of γ (“gravity coupled to matter”) correspond to random planar maps sampled with probability proportional to some sort of additional structure on the map, such as the number of spanning trees ($\gamma = \sqrt{2}$) or the critical Ising model partition function ($\gamma = \sqrt{3}$).

LQG surfaces were first defined nonrigorously in the physics literature by Polyakov in the 1980s in the context of bosonic string theory.² We refer to [11] for an extensive list of references to the physics literature on LQG. Certain special LQG surfaces are the topic of study in Liouville conformal field theory, the simplest nonrational conformal field theory. See, e.g., [3, 21] and the references therein for rigorous works on LQG from the conformal field theory perspective.

We now explain how to make rigorous sense of LQG surfaces as random metric measure spaces. The basic idea is to consider a family of continuous functions $\{h_\varepsilon\}_{\varepsilon>0}$ that approximate the GFF as $\varepsilon \rightarrow 0$, define approximate notions of area and distance by replacing h with a multiple of h_ε in (2.1) and (2.2), then send $\varepsilon \rightarrow 0$ and renormalize appropriately to get a measure and metric associated with γ -LQG. One possible choice of h_ε is the convolution with the heat kernel

$$h_\varepsilon(z) = \int_U h(z) p_\varepsilon(z, w) d^2w \text{ for } p_\varepsilon(z, w) := \frac{1}{\pi \varepsilon^2} e^{-|z-w|^2/\varepsilon^2}, \quad (2.6)$$

where the integral is interpreted as the distributional pairing of h with $p_\varepsilon(z, \cdot)$ and d^2w is two-dimensional Lebesgue measure. Note that $p_\varepsilon(z, w) d^2w$ approximates a point mass at z as $\varepsilon \rightarrow 0$, so h_ε is close to h (e.g., in the distributional sense) when ε is small. Other possible choices for h_ε include convolutions with other mollifiers, averages over circles, truncated versions of the orthonormal basis expansion (2.5), etc.

Figure 2. Simulations of the γ -LQG measure on the unit square produced by J. Miller. The square is divided into dyadic subsquares that all have approximately the same γ -LQG mass. Squares are colored according to their Euclidean size. Note that as γ increases, the Euclidean sizes of these squares become more variable.

²Roughly speaking, for $\mathbf{c} \in \mathbb{N}$ an evolving string in $\mathbb{R}^{\mathbf{c}-1}$ traces out a two-dimensional surface embedded in space-time $\mathbb{R}^{\mathbf{c}-1} \times \mathbb{R}$, called a world sheet. Polyakov wanted to develop a theory of integrals over all possible surfaces embedded in $\mathbb{R}^{\mathbf{c}}$ as a string-theoretic generalization of the Feynman path integral (which is an integral over all possible paths). To do this one needs to define a probability measure on surfaces. The most natural way of doing this turns out to work only when the “dimension of the space into which the surface is embedded” (a.k.a. the central charge) $\mathbf{c}$ lies in $(-\infty, 1]$, in which case the desired probability measure is γ -LQG with $\mathbf{c} = 25 - 6(2/\gamma + \gamma/2)^2$.

2.4. The γ -LQG area measure. For several possible choices of $\{h_\varepsilon\}_{\varepsilon>0}$, one can define the γ -LQG area measure as the a.s. limit

$$\mu_h = \lim_{\varepsilon \rightarrow 0} \varepsilon^{\gamma^2/2} e^{\gamma h_\varepsilon(z)} d^2z, \quad (2.7)$$

with respect to the weak topology on measures on U , where d^2z denotes Lebesgue measure as above. Indeed, this construction is a special case of a more general theory of regularized random measures called *Gaussian multiplicative chaos*, which was initiated in the work of Kahane in the 1980s; see [1, 27] for surveys of this theory. Several important properties of the measure (including the convergence of the circle average approximation and the so-called *KPZ formula*) were established by Duplantier and Sheffield [11]. The γ -LQG measure has no point masses and assigns positive mass to every open subset of U , but it is mutually singular with respect to the Lebesgue measure. In fact, it assigns full mass to a set of Hausdorff dimension $2 - \gamma^2/2$. See Figure 2 for simulations of the γ -LQG measure.

In the case when $\gamma \geq 2$, the limit (2.7) is identically zero, which is why we restrict to $\gamma \in (0, 2)$ when defining an LQG surface. For $\gamma = 2$ one can construct a measure with similar properties, but an additional logarithmic correction is needed in the scaling factor. For $\gamma > 2$, one can make sense of the γ -LQG measure as a purely atomic measure (i.e., a countable sum of point masses) that is closely related to γ' -LQG for $\gamma' = 4/\gamma \in (0, 2)$.

2.5. The γ -LQG metric. The γ -LQG metric D_h can be constructed in an analogous way to (2.7), but the proof that the approximations converge is much more involved than in the case of the measure. Intuitively, the reason for this is that if we replace h by a multiple of h_ε in (2.2), then the near-minimal paths could in principle be very different for different values of ε (although one gets a posteriori that this is not the case).

Before describing the construction of the metric we need to introduce an exponent $d_\gamma > 2$ that plays a fundamental role in the study of LQG distances and that can be defined in several equivalent ways. For example, it was shown in [6], building on [7, 12], that for certain random planar maps in the γ -universality class, a graph distance ball of radius $r \in \mathbb{N}$ in the map typically has of order r^{d_γ} vertices. Once the γ -LQG metric D_h is constructed, it is possible to show that d_γ is the Hausdorff dimension of the metric space (U, D_h) .

It can be shown using special symmetries for uniform planar maps or $\sqrt{8/3}$ -LQG that $d_{\sqrt{8/3}} = 4$. However, d_γ is not known (even at a heuristic level) for $\gamma \in (0, 2) \setminus \{\sqrt{8/3}\}$; determining its value is one of the most important open problems in the theory of LQG. See Section 4 for more on d_γ .

To approximate LQG distances, we let $D_h^\varepsilon(z, w)$ for $z, w \in U$ and $\varepsilon > 0$ be the metric defined as in (2.2) with $(2\gamma/d_\gamma)h_\varepsilon$ in place of h , i.e.,

$$D_h^\varepsilon(z, w) = \inf_{P: z \rightarrow w} \int_a^b e^{\frac{\gamma}{d_\gamma} h_\varepsilon(P(t))} |P'(t)| dt. \quad (2.8)$$

The reason why we have γ/d_γ instead of $\gamma/2$ in (2.8) is as follows. Since d_γ is the dimension of the γ -LQG surface, scaling γ -LQG areas by $C > 0$ should correspond to scaling γ -LQG distances by C^{1/d_γ} . By (2.7), scaling areas by C corresponds to adding the constant $\frac{1}{\gamma} \log C$ to h . By (2.8), this scales D_h^ε by C^{1/d_γ} , as desired.

It was shown by Ding–Dubédat–Dunlap–Falconet [4] that there are constants $\{\alpha_\varepsilon\}_{\varepsilon>0}$ such that the rescaled metrics $\alpha_\varepsilon^{-1} D_h^\varepsilon$ are tight with respect to the local uniform topology on U , and every subsequential limit is bi-Hölder continuous with respect to the Euclidean metric on U . Building on this and other works, including [8, 14], Gwynne and Miller [16] showed that in fact $\alpha_\varepsilon^{-1} D_h^\varepsilon$ converges in probability (not just subsequentially) to a limiting metric D_h that is defined to be the γ -LQG metric. They also proved an axiomatic characterization of D_h that implies that it is the only possible metric associated with γ -LQG.

The metric D_h induces the same topology on U as the Euclidean metric, but the Hausdorff dimension of the metric space (U, D_h) is $d_\gamma > 2$. Moreover, many of its geometric properties (e.g., scaling properties and the behavior of geodesics) are quite different from those of the Euclidean metric or indeed any smooth Riemannian metric on U . See Figure 3 for simulations of LQG metric balls.

There is also an earlier construction of the LQG metric in the special case when $\gamma = \sqrt{8/3}$ due to Miller and Sheffield [24, 25]. This construction does not use the direct regularization procedure (2.8) and instead is based on special symmetries for $\sqrt{8/3}$ -LQG. The Miller–Sheffield construction does not work for $\gamma \neq \sqrt{8/3}$, but it gives additional information about the $\sqrt{8/3}$ -LQG metric (e.g., its connection to random planar maps and certain Markov properties of metric balls) that is not apparent from the construction of [16]. The two constructions give the same metric by the characterization theorem from [16].

2.6. Conformal coordinate change. Just as for deterministic surfaces, it is possible to parametrize LQG surfaces in different ways. Suppose $\phi : \tilde{U} \rightarrow U$ is a conformal (i.e., bijective and holomorphic) map. Let h be the GFF on U and let

$$\tilde{h} := h \circ \phi + Q \log |\phi'| \quad \text{where} \quad Q := \frac{2}{\gamma} + \frac{\gamma}{2}. \quad (2.9)$$

Then $\tilde{h}$ is a random distribution on $\tilde{U}$. It is shown in [11, 15] that the γ -LQG area measures and metrics associated with h and $\tilde{h}$ are a.s. related by $\mu_h(\phi(A)) = \mu_{\tilde{h}}(A)$,

Figure 3. Simulations of γ -LQG metric balls w.r.t. the same GFF instance, produced by J. Miller. The colors indicate distances to the center of the ball. Geodesics from points in a grid back to the center point are shown in black. Note that these geodesics have a tree-like structure: unlike geodesics for a smooth Riemannian metric, LQG geodesics with different starting points and targeted at 0 merge into one another before reaching 0. This was proven to be the case in [14].

for each Borel measurable set $A \subset \tilde{U}$, and $D_h(\phi(z), \phi(w)) = D_{\tilde{h}}(z, w)$ for each $z, w \in \tilde{U}$. We think of (U, h) and $(\tilde{U}, \tilde{h})$ as representing different parametrizations of the same LQG surface. The coordinate change relation for μ_h and D_h shows that these objects depend only on the LQG surface, not on the choice of parametrization.

3. LQG as the Limit of Random Planar Maps

We now discuss the senses in which random planar maps should converge to LQG surfaces and the extent to which each type of convergence has been proven.

Figure 4. **Left:** Simulation of a large uniform quadrangulation embedded into $\mathbb{R}^3$ in such a way that the embedding is in some sense as close as possible to being an isometry, made by J. Bettinelli. The Gromov–Hausdorff limit of these triangulations is a $\sqrt{8/3}$ -LQG surface with the topology of the sphere. **Middle:** Simulation of a circle packing of a uniform triangulation made by J. Miller. **Right:** Simulation of the Tutte embedding of an instance of the $\sqrt{2}$ -mated CRT map, made by J. Miller. The counting measure on vertices approximates the $\sqrt{2}$ -LQG measure when the number of vertices is large.

3.1. Gromov–Hausdorff convergence. Let $\mathcal{K}$ be the set of all compact metric spaces. The Gromov–Hausdorff (GH) distance on $\mathcal{K}$ is the metric on $\mathcal{K}$ defined by

$$\mathfrak{d}^{\text{GH}}((X_1, d_1), (X_2, d_2)) := \inf_{(Y, D), t_1, t_2} \mathfrak{d}^{\text{H}}(t_1(X_1), t_2(X_2)), \quad (3.1)$$

where d^H denotes the Hausdorff distance on compact subsets of Y and the infimum is over all compact metric spaces (Y, D) and isometric embeddings $t_1 : (X_1, d_1) \rightarrow (Y, D)$ and $t_2 : (X_2, d_2) \rightarrow (Y, D)$. A planar map can be viewed as a compact metric space equipped with its graph distance (see Figure 4, left). One can then ask whether large random planar maps, with their graph distance rescaled appropriately, converge in distribution to γ -LQG surfaces w.r.t. the GH topology. We can additionally equip the map with its rescaled counting measure and ask for convergence in the *Gromov–Hausdorff–Prokhorov (GHP) topology*, the analog of the GH topology for metric measure spaces.

So far, GH or GHP convergence has been established only for uniform planar maps (including uniform maps with local constraints), which we recall correspond to $\gamma = \sqrt{8/3}$. The first such convergence results were obtained independently by Le Gall [22] and Miermont [23]. They showed that certain types of uniform planar maps (with graph distances rescaled by $n^{-1/4}$ and the counting measure on vertices scaled by $1/n$) converge in the GHP sense to the *Brownian map*, a random metric measure space that can be constructed from a continuum random tree via an explicit metric quotient procedure. Miller and Sheffield [24, 25] showed that a certain special $\sqrt{8/3}$ -LQG surface called the *quantum sphere* has the same distribution as the Brownian map viewed as a metric measure space modulo measure-preserving isometries. Hence uniform planar maps converge to $\sqrt{8/3}$ -LQG surfaces in the metric space sense. However, the proofs in [22–25] rely on special symmetries and combinatorial miracles for uniform planar maps and $\sqrt{8/3}$ -LQG (these are the same sorts of miracles which are needed to show that $d_{\sqrt{8/3}} = 4$). So, the convergence proofs do not extend to nonuniform maps, and $\gamma \neq \sqrt{8/3}$.

3.2. Embedding convergence. Although a planar map is defined only modulo orientation-preserving homeomorphisms of $\mathbb{C}$, there are various ways of *embedding* the map into $\mathbb{C}$, i.e., associating each vertex (resp., edge) with a point (resp., curve) in $\mathbb{C}$ in such a way that no two edges cross. Examples include *circle packing*, where the planar map is realized as the tangency graph of a collection of circles in the plane (Figure 4, middle), and *Tutte embedding* (a.k.a. harmonic embedding or barycentric embedding), which is defined by the condition that the position of each vertex is the average (barycenter) of the positions of its neighbors (Figure 4, right). Once we have chosen an embedding, we can ask, e.g., whether the counting measure on the vertices of the planar map, rescaled by the total number of vertices, converges in distribution (w.r.t. the weak topology on $\mathbb{C}$) to a variant of the γ -LQG area measure.

So far, there are three results establishing this type of convergence for random planar maps. The first result [17] establishes embedding convergence for a one-parameter family of random planar maps called *mated-CRT maps*, one for each $\gamma \in (0, 2)$, embedded via the Tutte embedding. Mated-CRT maps are constructed by gluing together pairs of correlated continuum random trees (CRTs) and are more directly connected to LQG than other types of random planar maps due to the results of [9]. We will not define these maps precisely here.

Embedding convergence toward $\sqrt{8/3}$ -LQG has also been established for uniform triangulations embedded via the so-called *Cardy embedding*, which is defined using crossing probabilities for percolation on the map [20], and for the Poisson–Voronoi approximation of the Brownian map embedded via the Tutte embedding [18]. In the case of uniform triangulations, it is in fact shown in [20] that one has convergence in the GHP sense (as discussed in Section 3.1) and the mating-of-trees sense (as discussed in Section 3.3 below) *simultaneously* with the convergence of the embedded map. It is a major open problem to establish these types of convergence results for other random planar map models and/or other embeddings (see Problem 2).

3.3. Mating-of-trees convergence. There are several combinatorial bijections that encode a random planar map decorated by some additional structure by means of a random walk on $\mathbb{Z}^2$. The simplest example of such a bijection is the Mullin bijection [2, 26], which encodes a planar map decorated by a spanning tree by a nearest-neighbor random walk on $\mathbb{Z}^2$. There are other such bijections, with different walks, that encode planar maps decorated by, e.g., percolation, bipolar orientations, or the critical Fortuin–Kasteleyn model. These bijections are called *mating-of-trees* bijections since the planar map is constructed from the walk by gluing together, or mating, the discrete random trees associated with the two coordinates of the walk. In the case of the Mullin bijection, the two trees are the spanning tree on the map and its corresponding dual spanning tree.

For random planar maps that are related to γ -LQG, it can be shown that the encoding walk converges in distribution to a two-dimensional Brownian motion $Z = (L, R)$ with the correlation between the coordinates L_t and R_t given by $-\cos(\pi\gamma^2/4)$ for each time t . A fundamental theorem of Duplantier, Miller, and Sheffield [9] shows that one can construct a γ -LQG surface from this two-dimensional Brownian motion Z via a continuum analog of a mating-of-trees bijection. Hence the convergence of the encoding walks in the discrete mating-of-trees bijections toward Z can be viewed as a convergence statement for random planar maps toward LQG surfaces in a certain topology: the one where two surfaces are close if their encoding paths

are close. This type of convergence is referred to as *mating-of-trees* or *peanosphere* convergence.

Mating-of-trees convergence is not strictly weaker than Gromov–Hausdorff convergence or embedding convergence, but it is arguably less natural than these other modes of convergence. However, it can still be used to extract a substantial amount of useful information about the random planar map. This includes scaling limits for various functionals of the map and the computation of exponents related to partition functions, graph distances, random curves on the map, etc. Moreover, for many types of random planar maps mating-of-trees convergence is the only scaling limit result available, and it can sometimes be used as an intermediate step in proving one of the other types of convergence (as is done in [20]). See [13] for a survey of mating-of-trees theory and its applications.

4. Open Problems

Here we discuss some of the most important open problems in the theory of Liouville quantum gravity. Additional open problems can be found, e.g., in [16, Section 7]. Our first problem was alluded to in Section 2.3.

Problem 1. What is the Hausdorff dimension d_γ of a γ -LQG surface, viewed as a metric space, for $\gamma \in (0, 2) \setminus \{\sqrt{8/3}\}$?

Recall that the Hausdorff dimension of $\sqrt{8/3}$ -LQG is four. The value of d_γ for general $\gamma \in (0, 2)$ is not known even at a physics level of rigor. The best-known physics prediction for d_γ , due to Watabiki in 1993 [29], was proven to be incorrect, at least for small values of γ , by Ding and Goswami [5]. However, it is known that $\gamma \mapsto d_\gamma$ is strictly increasing [6], and there are reasonably sharp upper and lower bounds for d_γ [6, 19]. For example, one has $3.550408 \leq d_{\sqrt{2}} \leq 3.63299$. Many quantities associated with random planar maps and LQG can be expressed in terms of d_γ , so computing d_γ would yield many additional results.

Problem 2. Show that weighted random planar map models of the type discussed in Section 1 converge in distribution to γ -LQG surfaces with $\gamma \in (0, 2) \setminus \{\sqrt{8/3}\}$ in both the Gromov–Hausdorff sense and under suitable embeddings into $\mathbb{C}$.

As noted in Section 3, both types of convergence have already been established for certain types of uniform planar maps toward $\sqrt{8/3}$ -LQG.

It is possible to construct regularized random measures associated with log-correlated Gaussian fields on n -dimensional manifolds for any $n \in \mathbb{N}$; see [1, 27]. However, the associated metric has been constructed only in dimension two.

Problem 3. Is it possible to construct random metrics associated with log-correlated Gaussian fields on $\mathbb{R}^n$ or, more generally, on n -dimensional manifolds for $n \geq 3$?

It is not clear whether random measures and metrics related to log-correlated Gaussian fields in dimensions greater than two should have any connection to higher-dimensional analogs of random planar maps.

ACKNOWLEDGMENTS. We thank two anonymous referees and also Nina Holden, Jason Miller, and Scott Sheffield for helpful comments. We thank Jérémie Bettinelli and Jason Miller for allowing us to use their beautiful simulations in this article.

References

- [1] Nathanaël Berestycki, *An elementary approach to Gaussian multiplicative chaos*, Electron. Commun. Probab. **22** (2017), Paper No. 27, 12, DOI 10.1214/17-ECP58. MR3652040
- [2] Olivier Bernardi, *Bijjective counting of tree-rooted maps and shuffles of parenthesis systems*, Electron. J. Combin. **14** (2007), no. 1, Research Paper 9, 36. MR2285813
- [3] François David, Antti Kupiainen, Rémi Rhodes, and Vincent Vargas, *Liouville quantum gravity on the Riemann sphere*, Comm. Math. Phys. **342** (2016), no. 3, 869–907, DOI 10.1007/s00220-016-2572-4. MR3465434
- [4] J. Ding, J. Dubédat, A. Dunlap, and H. Falconet, *Tightness of Liouville first passage percolation for $\gamma \in (0, 2)$* , arXiv e-prints (2019).
- [5] Jian Ding and Subhajit Goswami, *Upper bounds on Liouville first-passage percolation and Watabiki’s prediction*, Comm. Pure Appl. Math. **72** (2019), no. 11, 2331–2384, DOI 10.1002/cpa.21846. MR4011862
- [6] J. Ding and E. Gwynne, *The fractal dimension of Liouville quantum gravity: universality, monotonicity, and bounds*, Comm. Math. Phys. (to appear).
- [7] Jian Ding, Ofer Zeitouni, and Fuxi Zhang, *Heat kernel for Liouville Brownian motion and Liouville graph distance*, Comm. Math. Phys. **371** (2019), no. 2, 561–618, DOI 10.1007/s00220-019-03467-8. MR4019914
- [8] J. Dubédat, H. Falconet, E. Gwynne, J. Pfeffer, and X. Sun, *Weak LQG metrics and Liouville first passage percolation*, arXiv e-prints (2019).
- [9] B. Duplantier, J. Miller, and S. Sheffield, *Liouville quantum gravity as a mating of trees*, arXiv e-prints (2014).
- [10] Bertrand Duplantier, Rémi Rhodes, Scott Sheffield, and Vincent Vargas, *Log-correlated Gaussian fields: an overview*, Geometry, analysis and probability, Progr. Math., vol. 310, Birkhäuser/Springer, Cham, 2017, pp. 191–216. MR3821928
- [11] Bertrand Duplantier and Scott Sheffield, *Liouville quantum gravity and KPZ*, Invent. Math. **185** (2011), no. 2, 333–393, DOI 10.1007/s00222-010-0308-1. MR2819163
- [12] E. Gwynne, N. Holden, and X. Sun, *A mating-of-trees approach for graph distances in random planar maps*, arXiv e-prints (2017).

-
- [13] E. Gwynne, N. Holden, and X. Sun, *Mating of trees for random planar maps and Liouville quantum gravity: a survey*, arXiv:1910.04713, 2019.
 - [14] E. Gwynne and J. Miller, *Confluence of geodesics in Liouville quantum gravity for $\gamma \in (0, 2)$* , Ann. Probab. (to appear).
 - [15] E. Gwynne and J. Miller, *Conformal covariance of the Liouville quantum gravity metric for $\gamma \in (0, 2)$* , arXiv e-prints (2019).
 - [16] E. Gwynne and J. Miller, *Existence and uniqueness of the Liouville quantum gravity metric for $\gamma \in (0, 2)$* , arXiv e-prints (2019).
 - [17] E. Gwynne, J. Miller, and S. Sheffield, *The Tutte embedding of the mated-CRT map converges to Liouville quantum gravity*, arXiv e-prints (2017).
 - [18] E. Gwynne, J. Miller, and S. Sheffield, *The Tutte embedding of the Poisson-Voronoi tessellation of the Brownian disk converges to $\sqrt{8/3}$ -Liouville quantum gravity*, Comm. Math. Phys. (to appear).
 - [19] Ewain Gwynne and Joshua Pfeffer, *Bounds for distances and geodesic dimension in Liouville first passage percolation*, Electron. Commun. Probab. **24** (2019), Paper No. 56, 12. MR4003130
 - [20] Nina Holden and Xin Sun, *Convergence of uniform triangulations under the Cardy embedding*, arXiv e-prints (2019).
 - [21] A. Kupiainen, R. Rhodes, and V. Vargas, *Integrability of Liouville theory: proof of the DOZZ Formula*, Ann. of Math. (to appear).
 - [22] Jean-François Le Gall, *Uniqueness and universality of the Brownian map*, Ann. Probab. **41** (2013), no. 4, 2880–2960, DOI 10.1214/12-AOP792. MR3112934
 - [23] Grégory Miermont, *The Brownian map is the scaling limit of uniform random plane quadrangulations*, Acta Math. **210** (2013), no. 2, 319–401, DOI 10.1007/s11511-013-0096-8. MR3070569
 - [24] Jason Miller and Scott Sheffield, *Liouville quantum gravity and the Brownian map I: the QLE(8/3, 0) metric*, Invent. Math. **219** (2020), no. 1, 75–152, DOI 10.1007/s00222-019-00905-1. MR4050102
 - [25] J. Miller and S. Sheffield, *Liouville quantum gravity and the Brownian map II: geodesics and continuity of the embedding*, arXiv e-prints (2016).
 - [26] R. C. Mullin, *On the enumeration of tree-rooted maps*, Canadian J. Math. **19** (1967), 174–183, DOI 10.4153/CJM-1967-010-x. MR205882
 - [27] Rémi Rhodes and Vincent Vargas, *Gaussian multiplicative chaos and applications: a review*, Probab. Surv. **11** (2014), 315–392, DOI 10.1214/13-PS218. MR3274356
 - [28] Scott Sheffield, *Gaussian free fields for mathematicians*, Probab. Theory Related Fields **139** (2007), no. 3-4, 521–541, DOI 10.1007/s00440-006-0050-1. MR2322706
 - [29] Yoshiyuki Watabiki, *Analytic study of fractal structure of quantized surface in two-dimensional quantum gravity*, Quantum gravity (Kyoto, 1992), Progr. Theoret. Phys. Suppl. **114** (1993), 1–17, DOI 10.1143/PTPS.114.1. MR1268911

Ewain Gwynne

Credits

Figure 1 and author photo are courtesy of the author.

Figures 2, 3, and 4 (middle and right) are courtesy of Jason Miller.

Figure 4 (left) is courtesy of Jérémie Bettinelli.