

WHAT IS A RANDOM SURFACE?

SCOTT SHEFFIELD

ABSTRACT

Given $2n$ unit equilateral triangles, there are finitely many ways to glue each edge to a partner. We obtain a random *sphere-homeomorphic* surface by sampling uniformly from the gluings that produce a topological sphere. As $n \rightarrow \infty$, these random surfaces (appropriately scaled) converge in law. The limit is a “canonical” sphere-homeomorphic random surface, much the way Brownian motion is a canonical random path.

Depending on how the surface space and convergence topology are specified, the limit is the *Brownian sphere*, the *peanosphere*, the *pure Liouville quantum gravity sphere*, or a certain *conformal field theory*. All of these objects have concise definitions, and are all in some sense equivalent, but the equivalence is highly nontrivial, building on hundreds of math and physics papers over the past half-century.

More generally, the “continuum random surface embedded in d -dimensional Euclidean space” makes a kind of sense for $d \in (-\infty, 25)$ even when d is not a positive integer; and this can be extended to higher genus surfaces, surfaces with boundary, and surfaces with marked points or other decoration.

These constructions have deep roots in both mathematics and physics, drawing from classical graph theory, complex analysis, probability, and representation theory, as well as string theory, planar statistical physics, random matrix theory, and a simple model for two-dimensional quantum gravity.

We present here an informal, colloquium-level overview of the subject, which we hope will be accessible to both newcomers and experts. We aim to answer, as cleanly as possible, the fundamental question. What is a random surface?

MATHEMATICS SUBJECT CLASSIFICATION 2020


Primary 60G60; Secondary 60G57, 60J67, 60K35

KEYWORDS

Brownian map, peanosphere, Liouville quantum gravity, conformal field theory, Gaussian free field, Schramm–Loewner evolution, continuum random tree, Brownian snake, random planar map, random surface

As a random path, Brownian motion is *canonical* in the sense that it is uniquely characterized by certain symmetries, and *universal* in the sense that it is a limit of many kinds of discrete random walks. This paper will describe a similarly canonical and universal *random surface*. This random surface has several different formulations. To tell the story in a fanciful way, imagine a dialog, first about random paths, then about random surfaces.

INSTRUCTOR: Consider the simple random walk on \mathbb{Z} . At each time step, a coin toss decides whether position goes up or down. If you shrink the graph horizontally by a factor of C and vertically by a factor of \sqrt{C} , then the $C \rightarrow \infty$ limit is *Brownian motion*.



STUDENT: Great! But can you define Brownian motion directly in the continuum?

INSTRUCTOR: Sure! Fix $0 = t_0 < t_1 < \dots < t_n$. Specify the joint law of $B(t_1), \dots, B(t_n)$ by making increments $B(t_k) - B(t_{k-1})$ independent normals with mean 0, variance $t_k - t_{k-1}$. Extend to countable dense set (Kolmogorov extension), then all t (Kolmogorov–Čentsov).

STUDENT: Are there other natural ways to characterize Brownian motion?


INSTRUCTOR: Brownian motion is *canonical* in that it is the only random path with certain symmetries (like stationarity/independence of increments). It is *universal* in that (per central limit theorem) it is a limit of many discrete walks. It comes up everywhere.

STUDENT: What if I want a random path embedded in \mathbb{R}^d ?

INSTRUCTOR: Use a vector $(B_1(t), B_2(t), \dots, B_d(t))$ of independent Brownian motions.


The student is happy. Now imagine a similar dialog for random surfaces.

INSTRUCTOR: Take a uniformly random triangulation of the sphere with n triangles. The picture below is an example with 30,000 triangles by Budzinski, using some software to embed the surface in three dimensions and give us a view. The $n \rightarrow \infty$ limit of this object is a random fractal surface called the *Brownian sphere*. It is also a *peanosphere* and a *pure Liouville quantum gravity sphere*, and a *conformal field theory*.



STUDENT: You just listed four things! Which is the $n \rightarrow \infty$ limit of Budzinski's picture?

INSTRUCTOR: They all are! The difference comes down to the topology of convergence and the features of the limit presumed to be measurable. Think of them as different aspects of the same universal object. Four blind mathematicians feel the surface of an elephant and describe four different things:



Brownian sphere

Peanosphere

LQG sphere

Multipoint formalism

- *Brownian sphere*: a random metric measure space constructed from the so-called *Brownian snake* as explained in surveys by Le Gall, Miermont and Baez [27, 171, 175, 184, 186].
- *Peanosphere*: a mating of continuum random trees that encodes both a surface and an extra tree and/or collection of loops drawn on top of it, as explained in surveys by Gwynne, Holden, and Sun, and by Biane [51, 128].
- *Liouville quantum gravity sphere*: a random fractal Riemannian surface in which areas, lengths, and other measures are given by exponentials of a Gaussian free field ϕ , as explained in surveys by Berestycki, Ding, Dubédat, Duplantier, Garban, Gwynne, Miller, Powell, and Werner [33, 37, 91, 101, 113, 125, 188, 259].
- *Conformal field theory*: a collection of multipoint functions representing (regularized) integrals of products of the form $\prod e^{\alpha_i \phi(x_i)}$ with respect to a certain infinite measure, as explained mathematically by David, Guillarmou, Kupiainen, Rhodes,

and Vargas [87, 120, 121, 160, 254] who also survey the physics literature. The infinite measure is the *Polyakov measure* which is the product of an unrestricted-area measure on LQG spheres (with defining field ϕ) and Haar measure on the Möbius group $\mathrm{PSL}(2, \mathbb{C})$ (to select an embedding in \mathbb{C}).

Technical point: The *unit area* Brownian/Peano/LQG-sphere is a sample from a probability measure dS on a space of unit area surfaces. There is a natural infinite measure on *unrestricted-area* surfaces (with $k \geq 0$ marked points) given by $A^{-7/2+k} dAdS$ where (A, S) is the rescaling of S with area A . This is natural because the number of triangulations (quadrangulations, etc.) with n faces and k marked points scales like $C\beta^n n^{-7/2+k}$ for model-dependent constants C and β . Weighting the counting measure by β^{-n} , we obtain a discrete measure that (appropriately rescaled) converges to the measure above as the area-per-triangle ε goes to zero. If we replace β by the “off-critical” $\beta(1 + \varepsilon\mu)$ then the limit is $A^{-7/2+k} e^{-\mu A} dAdS$, which is finite if $\mu > 0$ and $k \geq 3$. Polyakov included the $e^{-\mu A}$ factor and motivated it in a different way, using the *Liouville equation* as we explain later.

STUDENT: Do I really have to learn all four viewpoints?

INSTRUCTOR: A lot of good work has been done by people fluent in only one of the four. But all four have important applications. For example, in the Brownian sphere construction, *distances* are easy to define, and you can show that the surface has fractal dimension 4. So to cover the surface with metric balls of radius $1/n$, you need $\approx n^4$ balls. In the peanosphere construction, one sees the random trees and loops that naturally live on top of a random surface, and most easily makes contact with loop-decorated discrete models. The LQG approach allows one to define *Brownian motion on a random surface* as well as other conformally defined objects (like SLE curves). Conformal field theory is a staple of physics. It ties into quantum field theory (if one analytically continues from Euclidean to Minkowski space and replaces probability measures with quantum wave functions).

STUDENT: Are these surfaces characterized by simple axioms like Brownian motion is?

INSTRUCTOR: Yes. But each viewpoint comes with its own axioms, its own history, its own motivation, its own surveys. The proofs that they agree are long and involved.

STUDENT: Can any of the viewpoints describe a random surface embedded in \mathbb{R}^d ?

INSTRUCTOR: Sure. You basically weight the law of the surface by the “number of ways” to embed it in \mathbb{R}^d . This amounts to weighting by the d th power of a certain “partition function” and this changes the law of the surface itself (even ignoring the embedding). You can apply such a weighting even when d is not a positive integer, so there is actually a one parameter family of random surfaces parameterized by d , with the scaling limit of Budzinski’s picture corresponding to $d = 0$. These surfaces are “rougher” when d is large and “smoother” when d is small, converging to the Euclidean sphere as $d \rightarrow -\infty$. They are defined as random

metric spaces for any $d < 25$, but are only finite-diameter and finite-volume if $d \leq 1$. The family can equivalently be parameterized by related quantities that come up in LQG theory like $Q > 0$ (where $d = 25 - 6Q^2$) or γ (where $Q = 2/\gamma + \gamma/2$ – note that γ is only real if $d \leq 1$ so that $Q \geq 2$) or by the peanosphere correlation coefficient.

STUDENT: I'm getting a bit lost. Can you give me the four definitions you promised?

The remainder of this paper will aim to do just that. We present a narrative-style introduction to each of the four viewpoints above, emphasizing the distinctive intellectual heritage behind each approach, as well as the relationships between them and the relevant recent developments.

Along the way, we will introduce several other natural objects: random trees (like the continuum random tree), random distributions (like the Gaussian free field) and random non-self-crossing curves (like the Schramm–Loewner evolution). We aim to make the exposition accessible to outsiders and newcomers, as well as to researchers who have expertise in one or more of the viewpoints but who might appreciate a high-level overview of the others. Readers interested in a still less technical account might take a look at the recent Quanta Magazine articles on this subject [143–145].

It is ambitious to try to tell four stories in one article, but we will do our best to convey at least the main ideas. The style is informal—no proofs—but we provide accurate definitions and the surveys mentioned above contain more detail. This paper is targeted primarily at a mathematical audience and is meant to be broadly accessible. This subject has deep roots in physics (especially string theory, quantum field theory, and planar statistical physics), but we will do our best to avoid any terminology that would be difficult for non-physicists to understand. Let us also stress that our reference list is long but very far from complete, biased by both the limits of the author's knowledge and the narrative focus of the paper. As such, it is more of a sampling than an exhaustive survey. Many highlights of the subject are not mentioned here.

For concreteness and simplicity, we will focus mostly on *sphere-homeomorphic* random surfaces. But all of the constructions in this paper have analogs that look the same locally but have different global topology. For example, in addition to the Brownian sphere one has a Brownian disk, a Brownian plane, a Brownian torus, Brownian surfaces of arbitrary genus and/or arbitrarily many boundary components, etc. [50,65,84,174,187,222]. Variants like these can be defined within all four viewpoints. One would expect that these definitions are consistent from one viewpoint to another (so that, e.g., a Brownian three-holed torus is somehow equivalent to a corresponding LQG three-holed torus or Peano three-holed torus). Significant work along these lines has been done, but the set of questions one can ask is very large and the program has not been completed yet.

1. BROWNIAN SPHERE: A RANDOM METRIC MEASURE SPACE

Our first construction is a random surface called the *Brownian sphere* (also known as *Brownian map*). The Brownian sphere is a random sphere-homeomorphic metric measure

space. It has been described in longer survey articles by Le Gall and Miermont [171, 175, 184, 186, 187] and in a shorter overview by Baez [27].


Both the peanosphere and the Brownian sphere can be constructed by “gluing together” a pair of “continuum trees” along their outer boundaries, which produces a sphere decorated by a space-filling curve (that somehow snakes in-between the two trees)—but the law of the pair of trees is different in the two settings. The idea of gluing together trees to obtain a sphere may seem counterintuitive, but we will see that it is well motivated by the discrete models.

Readers familiar with complex dynamics may also recall that the “mating” of two dendritic Julia sets (these are tree-like sets with empty interior) is a topological sphere, see, e.g., [25, 203, 247, 261], or look up the online video animations by Arnaud Chéritat. The Brownian sphere and peanosphere constructions are random versions of this phenomenon. Let us begin by defining *continuum random trees* which are sometimes also called *Brownian trees*.

1.1. Continuum random trees (Brownian trees)

Brownian motion is a random function $B(t)$ defined for $t \in [0, \infty)$. But Brownian motion (like the Brownian sphere) has many *variants* with similar local behavior. For example, the *Brownian bridge* is a Brownian motion on $t \in [0, 1]$ somehow *conditioned* on the (zero probability) event that $B(1) = 0$. The *Brownian excursion* is a Brownian bridge with $B(0) = B(1) = 0$ that is further conditioned to stay positive on $[0, 1]$.

If one starts with the graph of a Brownian excursion B , and identifies two points whenever they are connected by a horizontal chord under the graph (see below), one obtains a random metric space \mathcal{T} called the *continuum random tree*, which was first constructed by Aldous in 1991 [9–11], and which plays a role in the construction of both the Brownian sphere and the peanosphere.




Note that the function taking t to the graph point $(t, B(t))$ induces a function f from $[0, 1]$ to \mathcal{T} that “traces the boundary of the tree clockwise.” The pushforward of Lebesgue measure on $[0, 1]$ endows \mathcal{T} with a natural measure. Moreover, one can define the *distance* between $f(s)$ and $f(t)$ to be

$$\left(B(s) - \inf_{r \in [s, t]} B(r) \right) + \left(B(t) - \inf_{r \in [s, t]} B(r) \right).$$

This measures how far “down and up” one has to travel within \mathcal{T} to get from $f(s)$ and $f(t)$. Since \mathcal{T} comes with a natural measure and a natural metric, we call it a *random metric measure space*.

1.2. Planar map bijections that motivate Brownian sphere and peanosphere

The Brownian sphere has its historical roots in the study of planar maps. A *planar map* is a finite graph together with an embedding in the complex sphere $\mathbb{C} \cup \{\infty\}$, where two embeddings are considered equivalent if there is an orientation-preserving homeomorphism of $\mathbb{C} \cup \{\infty\}$ taking one to the other. For example, the two figures below are isomorphic as graphs but represent different planar maps, as there is no orientation-preserving homeomorphism of $\mathbb{C} \cup \{\infty\}$ mapping one to the other.



Combinatorially, a planar map is fully determined by the graph together with the “clockwise cyclic ordering” of the edges surrounding each vertex. In particular, the number of distinct planar maps with n edges is finite. To eliminate the ambiguity that comes from having nontrivial automorphisms, it is sometimes helpful to specify a “root” by fixing an oriented edge. If an orientation-preserving homeomorphism of $\mathbb{C} \cup \{\infty\}$ takes a planar map to itself (mapping edges to edges and vertices to vertices) and fixes the oriented edge, one can show that it must induce the identity map on the whole set of vertices and edges. To see why, imagine that your car starts driving along the oriented edge. From there you can describe how to get to any other vertex or edge with a set of directions like “Drive until you reach a vertex, then take the furthest road to your left, then drive until you reach another vertex, then take the third road from the left,” and so forth. This method of specifying vertices and edges is preserved by homeomorphisms of $\mathbb{C} \cup \{\infty\}$. In other words, fixing a root gives us a way to uniquely specify all other vertices (just as a Cartesian coordinate system gives us a way to uniquely specify points on a planar lattice). This implies that fixing the root eliminates all non-trivial planar map automorphisms: for example, a triangle has three non-trivial orientation-preserving automorphisms—corresponding to the three rotations—but if the triangle is assigned a root, then only the identity automorphism would fix the root.


Perhaps the most famous planar map problem of all is the “4 color conjecture” (also known as Guthrie’s problem) which captivated mathematicians worldwide from its formulation in 1852 to its computer-assisted proof by Appel and Haken in 1976. See, e.g., [81, 112, 146, 223] for historical accounts. The problem is to show, given any finite planar map, that it is possible to color each vertex one of four colors in such a way that no two neighbors have the same color. Note that if one has a map of South America, say, one can create a *planar map* by putting a vertex in the center of each country and drawing an edge connecting two countries wherever they share a nonzero-length border. The four color theorem allows one to color the countries in such a way that no two countries sharing a nonzero-length border have the same color. One of the key contributors to research on this problem, and to graph theory in general, was William Tutte.

By Tutte’s own account [253] (see also [211]), he was motivated by the four color theorem when he began thinking about enumerating *all* planar maps. He wrote his famous “census of” series in 1962 and 1963 [249–252] which, among other things, included a remarkable formula for the number of rooted planar maps with n edges, namely $\frac{2}{n+2} \cdot \frac{3^n}{n+1} \binom{2n}{n}$. Using, e.g., Stirling’s formula (or the local central limit theorem), one can show that $\binom{2n}{n} = \frac{4^n}{\sqrt{\pi n^{1/2}}} (1 + O(\frac{1}{n}))$ which implies that Tutte’s formula grows asymptotically like $C\beta^n n^{-7/2+k}$ (which is the formula from the *technical point* in our introductory dialog) with $C = \frac{2}{\sqrt{\pi}}$, $\beta = 12$, and $k = 1$. This is consistent with the formula from the introduction because the root effectively plays the role of a single marked point, and we recall that k represents the number of marked points.

These papers were followed by another series by Mullin [206–209] which included a formula for the number of rooted planar maps decorated by a distinguished spanning tree, namely $\frac{(2n)!(2n+2)!}{n![(n+1)!]^2(n+2)!}$. This grows asymptotically like $C\beta^n n^{-4+k}$ with $k = 1$ and $\beta = 16$. The addition of the “spanning-tree decoration” somehow changes the growth rate in a fundamental way, changing not only β and C but also the power of n (effectively replacing the $7/2$ with 4).

Both of these formulas have interesting bijective proofs: the former, in its first form due to Cori and Vauquelin in 1981 [80], was later advanced and popularized by Schaeffer [226]. The latter was essentially due to Mullin in 1967 [209], see also the explanation by Bernardi [40], and it reduces the problem of counting spanning-tree-decorated rooted planar maps to a problem about walks in \mathbb{Z}_+^2 (which can be separately addressed, e.g., with reflection arguments). These bijections motivate the definition of the Brownian map and the peanosphere, respectively. For now let us describe the bijections side by side, beginning with the Mullin bijection.

Suppose that (X_n, Y_n) is a simple walk in \mathbb{Z}_+^2 starting and ending at the origin. Then if we fix a large enough value of C , the graphs of X_n and $C - Y_n$ (linearly interpolated to \mathbb{R}) will not intersect; see the figure below. Draw a vertical red line at each time increment. Then declare two points in the graph of X_n (resp. $C - Y_n$) to be equivalent if one can draw a horizontal chord connecting them that does not go above the X_n graph (resp. below the $C - Y_n$ graph). In other words, we “identify” each pair of points connected by a horizontal (blue, black, or green) line segment. We also glue together the leftmost and rightmost red lines. After this gluing is done, the region below X_n collapses to become a tree (shown with black edges and green vertices) as does the region above $C - Y_n$ (shown with black edges and blue vertices) and we are left with a planar map with black/red edges and blue/green vertices. (The left and right red curves are glued together, which makes the topological disk they enclose into a sphere.) This planar map is a *triangulation* in which each triangle contains two red edges and one black edge.




If the black edges are erased, one is left with a red *quadrangulation* \mathcal{Q} which is bipartite: the blue and green vertices are the two partite classes. Each quadrilateral in \mathcal{Q} has one blue-to-blue diagonal. These diagonals together form a planar map M . Conversely, given the planar map M , it is not hard to show that one can recover \mathcal{Q} (by adding a vertex in the center of each face of M and connecting it to all of the boundary vertices of that face). The blue tree T is a spanning tree of M . The green-to-green diagonals of \mathcal{Q} form the dual graph M^* , and the green tree T^* is the dual spanning tree. The story above gives a bijection between:


- Simple walks (X_n, Y_n) in \mathbb{Z}_+^2 of length $2N$ that start and end at the origin.
- Pairs (M, T) where M is a rooted planar map and T is a spanning tree of M . Here *rooted* means that a distinguished vertex of M and a distinguished incident vertex of M^* are fixed (to be the roots of T and T^* , i.e., the two vertices attached to the leftmost red line, or the equivalent rightmost red line).

Choosing a uniformly *random* walk (X_n, Y_n) of the type described above produces a uniformly *random* (M, T) pair. The probability of a given map M is proportional to the number of spanning trees M has.


But what if we want to simply count the total number of rooted maps M (or the total number of quadrangulations \mathcal{Q}) instead of the number of (M, T) pairs? This is what the Cori–Vauquelin–Schaeffer bijection does. It can be seen as similar to the Mullin bijection but with a few key differences. First, instead of requiring (X_n, Y_n) to traverse lattice edges, we at each step allow Y_n to change by ± 1 and X_n by either 0 or ± 1 . Second, instead of perfectly horizontal green chords, we draw chords that are one unit higher on the right than on the left. We draw one such chord leftward starting at each vertex on the graph of X_n , which means that we have to add an extra vertex of minimal height as shown.



Third, we consider only (X_n, Y_n) pairs for which the above picture has a special property: namely, whenever two red vertical lines are incident to the same blue chord, their lower endpoints have the *same height*.



If we glue together the edges on the $C - Y_n$ graph connected by blue chords (the same way as in the Mullin bijection) we obtain a tree, and the third condition above is equivalent to the condition that if two red lines start at the same blue vertex on that tree then they terminate at green vertices of the same height. Thus we can label each vertex in the tree by the height of the green vertex (or vertices) it connects to. We then obtain a planar map by *gluing* the black and green trees on the right to one another: precisely, we glue each green vertex in the tree on the right to the blue vertex it is connected to by a red edge, effectively shrinking each red edge to a point. (Also, the outermost two green edges are understood to be glued/identified with each other, so that the topological disk they surround becomes a sphere.)



We can think of the graph on the left above as obtained by starting with the region below the X_n graph, then gluing together two black edges of the X_n graph whenever the edges of the $C - Y_n$ graph just above them correspond to the same edge of the upper tree. Whenever two of these black edges are glued together, they are either both horizontal (each part of a triangle with two green edges) or one increasing and one decreasing (respectively part of a two-gon with one green edge and a quadrilateral with three green edges). In each case, once the two black edges are glued together (and erased) we are left with a green quadrilateral.

The construction above yields a bijection between

- Well-labeled rooted planar trees (T, ℓ) (here ℓ maps vertices of T to positive integers, where the root has label one and adjacent vertices differ by 0 or ± 1) and
- Rooted quadrangulations \mathcal{Q} .


To prove that the above is a bijection, one has to show that one can reconstruct the green and black trees given just the green-edge quadrangulation \mathcal{Q} . Roughly speaking, one reconstructs the original green tree as follows. First, we shorten the green edges slightly so that they are no longer connected to each other at the vertices, but then we reconnect some of them: precisely, we connect each edge (directed “downward” towards its lower-label endpoint) to the successor downward-pointing edge that involves “turning maximally to the left.” Once this is done, the green tree is a tree of “leftmost geodesics” built from the directed green edges (each directed in the direction of decreasing distance to the root). Given the green tree, one can construct the X_n in the figure above and it is not too hard to check that one can reconstruct the upper tree as well (since one knows which green vertices have to be equivalent—and hence connected by red–blue–red paths).

We remark that once the black and green trees are glued together, they have vertices (but no edges) in common (this is different from the Mullin bijection) but in a sense the two trees still do not “cross” each other. To try to see why, it is interesting to imagine a path that crosses through the red edges in order from left to right (but does not intersect the green or black trees). Once we take the quotient (with respect to the equivalence class where all points on the same red edge are considered equivalent) this becomes a path that winds between the green and black trees, intersecting the trees at vertices but never crossing a branch of either tree.

This bijection (like the Mullin bijection) constructs a quadrangulation \mathcal{Q} by gluing together a pair of trees. The microscopic details of the gluing are a bit different: in the Mullin bijection the red edges form \mathcal{Q} and in the Cori-Vauquelin-Schaeffer bijection the green edges form \mathcal{Q} . There are many variants and analogs of the CVS bijection, see for instance [8, 12, 44, 49, 55, 226] and the references therein. At the global level, the main difference between the CVS bijections (and its variants) and the Mullin bijection (and its variants) is this: in the CVS bijection, one of the trees (the green one) is the so-called *leftmost geodesic tree of edges* and is determined by \mathcal{Q} itself. The second tree (the black one) is in some sense “dual” to the geodesic tree (and is also determined by \mathcal{Q} itself). By contrast, in the Mullin bijection the spanning tree is not determined by \mathcal{Q} and \mathcal{Q} is not uniform among all quadrangulations.

1.3. Unconstrained variant: when \mathcal{Q} is both pointed and rooted

There is a variant of the Mullin bijection in which we relax the restriction that X_n is nonnegative, see below.




Here we imagine that the left and right sides of the above rectangle are glued to one another (so that both X_n and Y_n then become indexed by a circle). We can then identify points on the same blue chord or the same green (possibly wrapping around) chord, just as before, and we obtain a pair of trees with red edges between them. A key difference from the original construction is that in this construction the root of the lower tree (corresponding to the minimum of X_n) and the root of the upper tree (corresponding to the minimum of Y_n) are no longer required to be adjacent. The minima of X_n and Y_n may occur at different places, and we can think of the vertex corresponding to the X_n minimum as an extra marked point. (In principle, one could allow both X_n and Y_n to take negative values, in which case the two tree roots—and the point described by the leftmost/rightmost red line—would effectively describe three marked points.)


It turns out that the CVS construction also has a similar variant, which corresponds to relaxing the constraint that all of the blue labels are positive. If we relax this constraint, then the corresponding X_n process can be positive and negative, and in particular no longer has to have a minimum at the same place that Y_n does. This means that the root of the upper tree and the root of the geodesic tree are no longer required to be adjacent. With this constraint relaxed, the set of possible labeled upper trees is easier to count: one has a Catalan number $\frac{1}{n+1} \binom{2n}{n}$ of rooted planar trees, and given the tree, there are exactly 3^n ways to choose the labels if we fix the root label to be zero. (We are free to do this since adding a constant to the labels does not affect the map construction, and hence only the labels modulo additive constant are relevant.) The formula $\frac{3^n}{n+1} \binom{2n}{n}$ can also be obtained by starting with Tutte's formula and multiplying by $n + 2$ (there are $n + 2$ possible to choose an extra marked vertex) and dividing by a factor of 2 (which is related to the choice of rootorientation—this is explained in many places, see e.g. [226] or the short explanation in Section 3 of [68]).

There is another way to think about the trajectory (X_n, Y_n) obtained from a labeled tree. For every point on a labeled tree, there is a labeled path from that point back to the root. As one traces the tree clockwise, one obtains a corresponding sequence of labeled paths; the labeled path seen at time n can be drawn as a “snake” (a vertical-to-horizontal function defined on $[0, Y_n]$) with the horizontal coordinate indicating the label.

The figures below (read from left to right) are steps in a Markov process on a space of “snakes.” Each figure is a sequence of edges, each of which goes up one unit and $-1, 0,$ or 1 units to the right. To move from one figure to the next, one first tosses a coin to decide whether to delete the top edge (with probability $1/2$) or to add an up-left edge (probability $1/6$), up-right edge (probability $1/6$) or up edge (probability $1/6$).




If the head of the snake starts at height zero—and we condition on the head height staying nonnegative and returning to zero at time $2n$ —then we can let the head height represent the process Y_n and the head horizontal location represent the process X_n , and it is not hard to see that this corresponds precisely to the (X_n, Y_n) construction from the labeled tree. If one rescales this process vertically by a factor of C and horizontally by \sqrt{C} , one obtains a continuum *Brownian snake* model in the limit. The Brownian snake was introduced by Le Gall in the 1990s (see, e.g., [167]) not long after Aldous introduced the continuum random tree (also known as Brownian tree). The terms *Brownian snake* and *Markov snake* were coined by Dynkin and Kuznetsov (who credit Le Gall for the construction) in 1995 [106].



Now the law of the continuum pair (X_t, Y_t) can be described as follows: one first samples Y_t as a Brownian excursion on $[0, 1]$ and then takes a quotient of the graph of Y_t to produce a Brownian tree \mathcal{T} , together with the natural map f taking t to \mathcal{T} (which traces around the boundary of \mathcal{T} clockwise as t varies from 0 to 1). One then generates a Brownian motion ζ indexed by \mathcal{T} (taken to be zero at the root of \mathcal{T}). This in turn determines a process X_t by $X_t = \zeta(f(t))$. The minimum of this process corresponds to the location of the root of the geodesic tree. The picture of the whole snake at time t (viewed as a vertical-to-horizontal function as above) is described by the restriction of ζ to the path in \mathcal{T} from $f(0)$ and $f(t)$.

1.4. Passing to the continuum

One would expect both the Mullin and the Cori–Vauquelin–Schaeffer bijections to have limits given by the gluing of two fractal trees defined by continuum processes X_t and Y_t , something like the figure below:



However, we would expect the law of (X_t, Y_t) to be rather different in the two scenarios. In the Mullin limit, X_t and Y_t are independent Brownian excursions, and the picture is a topological quotient of two independent continuum random trees (with respect to the equivalence that identifies two points if they are connected by a red curve). The topological sphere constructed this way is a special case of the *peanosphere* that we will discuss in the next section. In the CVS limit, X_t and Y_t correspond to the horizontal and vertical components of the head location in the Brownian snake process. The *Brownian sphere* (also known as *Brownian map*) is defined to be the metric space obtained by starting with the tree generated by X_t (viewed as a metric space) and then taking the metric quotient with respect to the equivalence that identifies points if they are connected (by the red curves) to the same point on the upper tree defined by Y .

The Brownian map was originally defined a bit differently. Marckert and Mokkadem [180], building on Schaeffer’s work [226] (also citing, e.g., the early work of Krikun and of Angel and Schramm [21, 157]), argued that the limit of the discrete models could be defined in some sense (since the (X_n, Y_n) process had a limit) and coined the term *Brownian map* to describe the limiting object; however, they did not show convergence in the space of random metric spaces. Le Gall showed that subsequential limits of random quadrangulations exist w.r.t. to the topology obtained from the *Gromov–Hausdorff metric* (a natural “metric on the space of all compact metric spaces”) and are a.s. homeomorphic to a metric quotient of the Brownian tree [168]. Le Gall and Paulin subsequently showed that all such subsequential limits are a.s. homeomorphic to the 2-sphere [177], see also the sphericity proof by Miermont [183]. The term *Brownian map* was sometimes used (see e.g. [176]) to describe any one of the subsequential scaling limits of the discrete models, and in this language, the question “Do random quadrangulations have a (non-subsequential) Gromov–Hausdorff scaling limit?” was formulated as the question: “Is the Brownian map unique?” This problem was solved independently by Le Gall and by Miermont [170, 185], who also showed that this limit exists in law (without passing to a subsequence) and that the limiting random metric space is indeed the Brownian-snake-based metric quotient defined above.

Since this foundational work, Brownian surfaces have been the subject of a sizable literature. Many different kinds of random planar maps (quadrangulations, triangulations, various variants, etc.) have been shown to have the Brownian sphere as a scaling limit, strengthening the idea that the Brownian map (like Brownian motion) is something “universal” [2–7, 31, 49, 69, 181]. There has also been significant work done on the continuum model. For example, the behavior of the geodesics in the Brownian sphere is quite interesting (geodesics “merge into” one another in ways that are very unlike what one sees for Euclidean spheres) and much has been done to understand their behavior [20, 48, 54, 57, 58, 169, 189].

1.5. An axiomatic approach

Like Brownian motion, the Brownian sphere is uniquely characterized by certain axioms, as shown in a long paper by the author and Miller [196] which draws on related ideas by Bertoin, Budd, Curien, Kortchemski, Krikun, Le Gall, Miermont and others [21, 46, 47, 83, 85, 86, 157, 172, 173]. We won’t give a fully precise description here, but let us summarize the rough idea. For Brownian motion, there is a Markov property, which states that *given* $B(t)$ the conditional laws of B restricted to $[0, t]$ and B restricted to $[t, \infty)$ are independent. Now consider an unrestricted-area Brownian sphere with two marked points x and y . The natural Markov property in this setting states that *given* the boundary length of the filled metric ball of radius r centered at x (here the “filled metric ball” is obtained by starting with the ordinary metric ball and adding the components of its complement that don’t contain y) the conditional laws of the ball and its complement are independent and depend on the boundary length in a scale invariant way. Furthermore, if we cut a filled metric ball into slices—using geodesics from evenly spaced points around the boundary to the center—then the slices are conditionally independent of each other given the length of their intersections with the metric ball boundary. The main result of [196] (very roughly speaking) is that if a measure on sphere-homeomorphic metric measure spaces satisfies these properties, then it must be the Brownian sphere.

2. PEANOSPHERE: A RANDOM MATING OF TREES

2.1. Basic definition

In Section 1 we saw that the Mullin bijection gives a bijection between lattice walks in \mathbb{Z}_+^2 (starting and ending at zero) and pairs (M, T) where M is a rooted planar map and T is a spanning tree on M . The X and Y coordinates of the walk encode, respectively, a tree and a dual tree, which are somehow stitched together to create M . This is actually a special case of a much more general idea. The idea of considering planar maps *together with* an extra structure on the map (a spanning tree, a collection of loops, a distinguished edge subset, a bipolar orientation, etc.) has long been a staple of this subject, both on the mathematics side and the physics side. The extra structure is sometimes called a *decoration* or (on the physics side) a *statistical physics model* or a *matter field*.

It turns out that *many* kinds of planar map decorations can be used to produce a spanning-tree/dual-spanning-tree pair in some way. And the resulting trees are often encoded

by *some kind of* lattice walk conditioned to stay in \mathbb{Z}_+^2 , so that in the fine mesh scaling limit, one obtains *some kind of* Brownian motion conditioned to stay in \mathbb{R}_+^2 (starting and ending at the origin). However, in general, the kind of Brownian motion involved (before imposing the quadrant constraint) may be one in which X_t and Y_t are *correlated*. That is, the “diffusion matrix” may be such that $\text{Var}(X_t) = \text{Var}(Y_t) = t$ and $\text{Cov}(X_t, Y_t) = \rho t$ for some possibly nonzero correlation coefficient ρ . (There is no loss in assuming $\text{Var}(X_t) = \text{Var}(Y_t) = t$, since multiplying X or Y by a constant does not change the topological surface construction.)

The limit of the Mullin construction corresponds to $\rho = 0$, but in general one might consider any ρ strictly between -1 (where the Brownian trees are perfectly negatively correlated) and 1 (where the Brownian trees are perfectly correlated). For each ρ there is a natural way to make sense of this Brownian motion *conditioned* to stay in \mathbb{R}_+^2 (starting at the origin at time 0 and returning there at time 1). One can then use the corresponding (X_t, Y_t) pair to generate a pair of trees that can be “mated together” in the manner described in the previous section.

Formally, then, the peanosphere is just a pair of random measure-endowed metric planar trees that, when glued together along their boundaries, make a topological sphere, as described in the author’s work with Duplantier and Miller [102, 194] and in more recent surveys [51, 128]. The sphere comes equipped with a non-self-crossing, space-filling path (also known as *Peano curve*) which in some sense traces the interface between the two continuum trees. (This is the motivation behind the term *peanosphere* which was originally proposed by Richard Kenyon in private communication.) It is also equipped with a measure, which is the pushforward of Lebesgue measure on the parameterizing interval $[0, 1]$. Note that no matter the value of ρ , the law of (X_t, Y_t) is unchanged if we swap the roles of X_t and Y_t , which (by contrast) is very much not the case for the (X_t, Y_t) pair used to define the Brownian map.

Moore in 1925 gave a very general criterion for determining when a topological quotient space is topologically a sphere [205] and as explained in [102, 194] one can verify directly that these criteria are satisfied in our setting, so that the object obtained by “gluing the two trees together along their outer boundaries” is indeed a topological sphere. The *peanosphere* does not *a priori* come with a simple metric space structure like the Brownian sphere does, because neither the tree nor the dual tree is a tree of geodesics. On the other hand, both the tree and the dual tree (as determined by X_t and Y_t) can be viewed as metric measure spaces themselves, so lengths of arcs *within* these trees are well defined.


2.2. Percolation and other simple kinds of decoration

A major reason to study the *peanosphere* is that it helps one understand random (M, T) pairs, where M is a random planar map and T is some kind *extra decoration* on M . Indeed, most of the scaling limit results known for decorated random planar maps make use of the peanosphere construction in some way. But before we discuss that topic, let us say a few words about randomly decorated *deterministic* planar maps, i.e., let us suppose that M is fixed to be a grid or lattice, and we are choosing some way to “decorate” M , e.g., with a coloring or spanning tree.

In statistical physics, it is often interesting to try to find the *simplest possible models* that exhibit behaviors (phase transitions, correlation decays, fractal patterns, etc.) that one might see in complicated real world systems. By understanding these models thoroughly and mathematically, one can hope to get a glimpse of *why* similar behaviors appear in more complex systems. We will describe a few of these very simple models here.


First, let us discuss *lattice percolation* which was introduced to the mathematics literature by Broadbent and Hammersley in [61]. In this model, one starts with a set of edges—or vertices or faces—and tosses an independent coin for each one to decide whether it is “open” or “closed.” This is easiest to visualize by a computer. For example, the following Mathematica code generates the picture below it:

```
n=40;Graphics[{Table[If[RandomInteger[1]==1,Line[{{i,j},{i,j+1}}]],{i,0,n},{j,0,n-1}],
Table[If[RandomInteger[1]==1,Line[{{i,j},{i+1,j}}]],{i,0,n-1},{j,0,n}]]}
```



The above is a 40×40 grid where one tosses a fair coin independently for each edge to decide whether to display it or not. The displayed edges are the open edges; the edges not displayed are closed. A connected component of open edges is called a *cluster*. One can think of this as a simplistic model for a porous medium—with open edges representing pathways that current or fluid can flow through, and closed edges representing obstructions. The figure above is called *bond percolation* (with parameter $p = 1/2$) because it is the edges (“bonds”) that are assigned to be open or closed. An alternative is *site percolation* which independently assigns each vertex—or face if one uses the “dual perspective”—to be open with probability p . By way of illustration, the following code generates the figure below it:

```
n=40; Graphics[Table[{If[(i-n)(j-n)==0,Blue, If[i j==0, Yellow, If[RandomInteger[1]==1,
Yellow,Blue]}], RegularPolygon[i{-Sqrt[3], -1}+j{-Sqrt[3], 1}, {1, 0}, 6]}, {i, 0, n}, {j, 0, n}]]
```



This figure represents percolation on the faces of a hexagonal grid. For each interior face, an independent coin toss decides whether it is open (blue) or closed (yellow). Here we have also imposed some deterministic boundary conditions, fixing the colors to be blue along one boundary arc and yellow on the complementary arc. The “blue–yellow” edges (i.e., the edges on the boundary between blue and yellow) form many finite length cycles and one long path, which starts at the lower corner and ends at the upper corner. The long path is the boundary between the outermost yellow cluster (which contains half of the boundary) and the outermost blue cluster (which contains the other half of the boundary). In the fine mesh limit, it converges in law to a random fractal curve called the chordal *Schramm–Loewner evolution with parameter $\kappa = 6$* (SLE_6) as proved in the breakthrough work of Smirnov [239] and Camia and Newman [67], see also the expository paper by Sun [243]. We will say more about SLE curves in Section 2.5.

Both bond percolation and site percolation have natural *variants* in which the edge or vertex subsets are sampled from different (nonuniform) probability measures. One of the simplest such models is the *Fortuin–Kasteleyn (FK) random cluster model*, which was introduced by Fortuin and Kasteleyn in 1972 [109–111]. It has a “partition function” (defined below) that is equivalent to the famous *Tutte polynomial* introduced by Tutte in 1954 [248]. See the historical accounts in [82, 108] and the overviews by Bernardi and Welsh [41, 256] of the Tutte polynomial and its (surprisingly numerous) applications.

In the FK random cluster model, one begins with a finite graph $G = (V, E)$ and chooses a random $T \subset E$ (which need not necessarily be a spanning tree). The probability of a given T is proportional to a *weight* given by $W_G(T) := q^{k(T)} w^{|T|}$ where q and w are positive constants, and $k(T)$ represents the number of *clusters* (i.e., connected components) of the graph (V, T) . Note that $k(t)$ is maximized and equal to $|V|$ if T is empty. It is minimized and equal to 1 if (V, T) is connected.

If $q = 1$, then the FK random cluster model is just ordinary bond percolation with $p = w/(1 + w)$. Generally, a large w gives a bias toward T having more edges, and a small w gives a bias toward T having fewer edges. If q is very large, then the probability measure is biased in favor of (V, T) being highly *disconnected* (i.e., having lots of components). If w is close to zero, then the bias is in favor of T being more *connected* (i.e., having few components). The uniform spanning tree can be obtained as a limiting case of the FK random cluster model. (Taking $w \rightarrow 0$ quickly forces T to be a.s. connected; taking $p \rightarrow 0$ slowly makes $|T|$ as small as possible, and the connected T that minimize $|T|$ are spanning trees.) The *partition function* of the FK random cluster model is defined to be the sum of the weights, taken over all $T \subset E$,

$$Z_G(q, w) = \sum_{T \subset E} W_G(T),$$


which again is equivalent to the Tutte polynomial (up to a certain coordinate change). An important point to stress is that if $q = 1$, the partition function $W_G(T)$ depends only on the number of edges in G .

For certain parameters, the FK random cluster model is also closely related to the hugely influential *Ising model* which was presented by Ising in 1925 but actually introduced earlier by Lenz in 1920, see, e.g., the historical account at [63]. In the Ising model, not all of the possible face colorings are equally likely; rather one *weights* the probability of each coloring by a constant to the number of adjacent pairs on which the colors disagree.


The decorations we have discussed so far (bond percolation, site percolation, uniform spanning tree, FK random cluster model, Ising model) are some of the very simplest statistical physics models. But there are others. For example, one can decorate a graph G with a *bipolar orientation*, which is a way of *orienting* the edges in G so that only one vertex v (the sink) has all incident edges oriented toward v , and only one vertex w (the source) has all incident edges oriented away from w . The number of bipolar orientations is also encoded in the Tutte polynomial, see, e.g., Bernardi’s explanation in [41]. One can also decorate G with a so-called *Schnyder wood* or an instance of the Gaussian free field or a Brownian loop soup (all of which will be mentioned below). The list goes on. We will not properly survey the literature on decorated planar maps here, but a few examples of work in this area include [39, 42, 43, 52, 53, 55, 56, 78].

2.3. Decorated random planar maps

In this section we give some examples of how one can generate a pair of spanning trees if one starts with another type of decoration. The figures below were used in [236] to explain how to generate a pair of trees from an instance of the FK random cluster model. By way of setup, one starts with a planar map M in black, adds a vertex in the center of each face, and draws a green edge connecting it to each vertex on the boundary of the face, thereby producing a green quadrangulation Q . The planar map M^* (which is dual to M) is obtained by connecting red-to-red within each green quadrilateral of Q , and is shown with dotted lines. (Here M itself is recovered from Q by connecting blue to blue within each quadrilateral.)




Next, the figures below illustrate an edge subset T in solid black (which happens to form a spanning tree) with the dual T^* shown as dotted black lines. (The dual consists of a red-to-red dotted line contained within each quadrilateral that *does not* have a blue-to-blue solid line.) The interface between the two trees is represented by a red curve that passes through all of the green edges one at a time, and starts at a location between the root and dual root (shown as large blue and red vertices).



Finally, the figure below on the left illustrates an example of T (solid black lines) that is not a spanning tree (it contains cycles and there is one blue vertex that is not connected to any others). When one draws the red curve (which passes through green lines but never through black solid/dotted lines) one finds that it is not connected—in fact, it has five different components. In the figure on the right, four “swaps” have been made, where the edge in a quadrilateral was changed to a dual edge, or vice versa. The new edges/dual-edges produced by the swaps are colored in yellow. There is a simple algorithm (due to Bernardi in [41], see the author’s explanation in [236]) for choosing which edges to swap. (Very roughly speaking, one tries to form a single red loop that crosses all the green edges one at a time, as above—but whenever the path is about to “seal off” a region without traversing it, one replaces a

black blue-to-blue edge with a dotted dual red-to-red edge, or vice versa, to prevent this from happening.) Each swap reduces the number of red loops by one—and at the end, one is left with a single red loop, and a tree/spanning-tree pair. This is called the *hamburger–cheeseburger* construction because of its relationship to a simple inventory accumulation model in which two kinds of products (“burgers”) are created and consumed [236].



At the end of the day, [236] uses some algebra (which is too detailed to explain here) to show that the lattice path encoding this pair of trees converges to a two-dimensional Brownian motion where the correlation coefficient ρ depends on the FK-cluster parameter in a simple way. While the work in [236] was done for infinite graphs, Gwynne, Mao, and Sun have understood the scaling limits for the FK model on a finite random planar map, and have given much stronger forms of convergence, which also encode the structure of the loops themselves in a direct way [131, 141, 142]. See also [35, 77] for more results about critical FK random maps. Since then, work by Bernardi, Holden, and Sun has explained that a simpler, tailor-made bijection applies for site-percolation-decorated random triangulations, in which case $\rho = 0.5$.

With site or bond percolation, the partition functions depend only on the number of vertices or edges, and once one conditions on this the marginal law of M is the same as in the undecorated model [45]. This suggests that the peanosphere corresponding to $\rho = 0.5$ should be somehow equivalent to the Brownian map (together with extra randomness encoding the percolation structure on the Brownian map). And this is indeed the case—see Section 5.

Other bijections have been found for bipolar orientations and Schnyder woods, and so-called active spanning trees [42, 51, 130, 155, 178]. All of these models have different ρ values. For an extreme case, note that if we naively set $\rho = 1$, then we are effectively gluing two identical Brownian trees to each other, and we simply obtain the same Brownian tree back. On the other hand, there is a way of taking the $\rho \rightarrow 1$ limit that produces a nontrivial object, as shown in a remarkable paper by Aru, Holden, Powell, and Sun [22], which shows that there is a nontrivial construction that is far nicer than anyone could have expected, see also [24]. (The $\rho = -1$ limit is expected to correspond to an ordinary Euclidean sphere.)


As a short preview for the experts, let us now note that there is a general relationship between ρ and the κ parameter from SLE theory, a parameter we will say more about in Section 2.5. Using this relationship, one finds for example that the bipolar orientation model corresponds to SLE_{12} and the Schnyder wood model to SLE_{16} . This is because (see Section 5) there is a natural way to “conformally map” the peanosphere onto the Euclidean sphere, and when this is done the space-filling Peano curve becomes a space-filling form of SLE_κ (see Sections 2.5 and 2.6) with $\kappa > 4$, where $\rho = -\cos(4\pi/\kappa)$ [102]. Here κ ranges from 4 to ∞ as ρ ranges from 1 to -1 . Let us note that the peanosphere surface itself turns out to be equivalent to an LQG sphere (see Section 3) with relevant parameters satisfying $\gamma^2 = 16/\kappa$ and $Q = 2/\gamma + \gamma/2$ and $d = 25 - 6Q^2$. The special value $\rho = 0.5$ corresponds to $\kappa = 6$ (and $\gamma = \sqrt{8/3}$ and $d = 0$). We will say more about these quantities in subsequent sections.

2.4. Computing the scaling exponent

As mentioned above, it turns out that if we set $\rho = 0.5$ then the peanosphere is a random surface that is equivalent to the Brownian sphere (decorated by a percolation model that does not change the law of the surface itself) although this is far from obvious at first glance. In this case, the two continuum trees are somehow midway between perfectly uncorrelated (as in the Mullin limit) and perfectly correlated. Even before constructing that bijection, it is not hard to argue heuristically that 0.5 is the only correlation coefficient that is consistent with the $C\beta^n n^{-7/2+k}$ formula, just as 0 is the only correlation coefficient consistent with $C\beta^n n^{-4+k}$. In general, one can derive a relationship between ρ and the b in $C\beta^n n^{-b+k}$.

We will not give the full details here, but let us sketch the rough idea, just to show that the relationship between b and ρ is not something mysterious. What is the probability that a simple random walk on \mathbb{Z}^2 , started at the origin, remains in the wedge of angle θ until time n , and is back at the origin at time n ? By way of illustration, part of a wedge of angle $\pi/4$ is shown below on the left (below the generating Mathematica code) along with its conformal image under the map $z \rightarrow z^4$.


```
f[z_]=z^4;F[{x_,y_}]={Re[f[x+I y]],Im[f[x+I y]]};n=40;
{ParametricPlot[Table[{{t(1-j/n)+j/n,j/n},{j/n,t j/n}},{j,0,n}],{t,0,1}],
ParametricPlot[Table[{F[{t(1-j/n)+j/n,j/n}],F[{j/n,t j/n}]}},{j,0,n}],{t,0,1}]}
```



Here is the heuristic. (References to rigorous arguments along these lines are given in [155].) Suppose one does a simple random walk in a wedge of angle θ , with lattice size $1/\sqrt{n}$, and we want the probability of the walk starting at the origin and returning in n steps (without leaving the wedge). Using the fact that the random walk approximates Brownian motion, and Brownian motion is conformally invariant under the map shown above, it appears that the odds of escaping to a macroscopic distance within $n/3$ steps should be of order $(1/\sqrt{n})^{\pi/\theta} = n^{-\pi/(2\theta)}$. The same should hold if one selects the walk increments in reverse (starting at time n) for order $n/3$ steps; and then there is an order $1/n$ chance that the middle third lines up correctly, so overall one expects a probability of order $n^{-\pi/\theta-1}$.

The same should also hold if the simple random walk is replaced by another walk that has Brownian motion as a scaling limit. If $\rho \neq 0$ then it is necessary to *stretch* or *squash* the grid by some amount (in the $(1, -1)$ direction) in order to produce a walk with Brownian motion as a scaling limit, and one can easily work out the angle θ as a function of ρ .

The formula $n^{-\pi/\theta-1}$ corresponds to n^{-b+k} where $k = 1$ and $b = \pi/\theta + 2$ so that $\theta = \pi/(b - 2)$. If we have $b = 4$ then $\theta = \pi/2$, which makes sense since the Mullin bijection corresponds to a simple random walk in a quadrant. More generally, if we have $X_t = \tilde{B}_t + aB_t$ and $Y_t = \tilde{B}_t - aB_t$ (where B_t and \tilde{B}_t are independent standard one-dimensional Brownian motions), we have a correlation coefficient $\rho = (1 - a^2)/(1 + a^2)$ where a is the factor by which the space is scaled in the $(1, -1)$ direction. Then θ is the angle obtained when we squash the standard positive quadrant by a factor of a in that direction. Precisely, $1/a = \tan(\theta/2)$ so $a = \cot(\theta/2)$ and $\rho = (1 - \cot(\theta/2)^2)/(1 + \cot(\theta/2)^2)^2 = -\cos(\theta) = -\cos(\frac{\pi}{b-2})$. Plotting this relationship, we see that the b value increases from 3 to ∞ as ρ decreases from 1 to -1 .



2.5. Schramm–Loewner evolution

Schramm–Loewner evolution is by itself a large subject. It has been the topic of several other sets of ICM lecture notes, see, e.g., the ICM notes by Duplantier, Lawler, Schramm, Smirnov, and Werner [101,162,229,240,242,258]. Longer expository introductions to the subject include, e.g., [36,152,161,257].

Here we will briefly explain what SLE curves are. First, let us present the axioms that motivated the definition. In 1999 [228], Schramm set out to construct—for any simply connected domain D with boundary points a and b —a random non-self-crossing chordal

curve η connecting a and b . Schramm insisted that the definition of an SLE curve have two properties. First, the definition had to be *conformally invariant* meaning that if ψ is a conformal (i.e., analytic and one-to-one) map taking D to a domain $\psi(D)$ then the image of η under ψ should have the law of an SLE in $\psi(D)$ from $\psi(a)$ to $\psi(b)$ (up to a time change). Second, the path should be *Markovian* in the sense that *given* η up to a stopping time τ , the conditional law of *the rest* of η is (up to a time change) that of an SLE in $D \setminus \eta([0, \tau])$ from $\eta(\tau)$ to b . Schramm showed that there was only a one-parameter family of ways to define SLE if one insists on these properties. Schramm indexed this family by a parameter $\kappa \in [0, \infty)$.

For completeness, let us now give Schramm’s more explicit definition of SLE (though we will not say too much more about it here). By conformal invariance, it is enough to define the law of η for one domain and one pair of boundary points. It turns out to be convenient to work with the upper half plane $\mathbb{H} \subset \mathbb{C}$ with $a = 0$ and $b = \infty$. For any time t , we define the function g_t to be the unique conformal map from the unbounded component of $\mathbb{H} \setminus \eta([0, t])$ to \mathbb{H} that satisfies

$$\lim_{|z| \rightarrow \infty} g_t(z) - z = 0.$$

Schramm defined SLE in a rather indirect way: namely, he constructed the analytic functions g_t , and then used these functions to deduce what η must be. The g_t are defined by setting $g_0(z) = z$ and then requiring that for any fixed $z \in \mathbb{H}$, the ODE

$$\partial_t g_t(z) = \frac{2}{g_t - W_t}$$

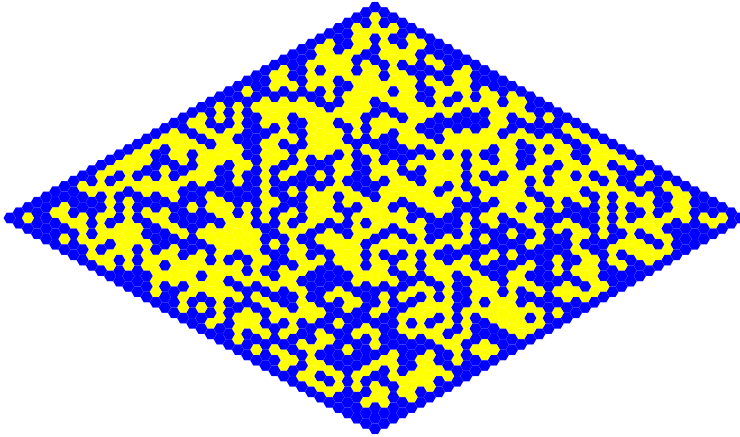
is satisfied up until the smallest time t at which z is hit (or cut off from infinity by) the curve $\eta([0, t])$, where $W_t := B_{\kappa t}$ is a standard Brownian motion sped up by a factor of κ . This requirement determines the functions g_t which in turn determine η .

In some sense, the larger the κ (and hence the faster the W_t is moving up and down), the “windier” the curve becomes. In fact, Rohde and Schramm showed in [224] that η is a.s. a simple curve when $\kappa \in [0, 4]$, that η a.s. hits (but does not cross) itself when $\eta \in (4, 8)$ and that η is a.s. a *space-filling* curve when $\kappa \geq 8$. Beffara showed that the Hausdorff dimension of the range of η is linear in κ , a.s. given by $\min\{1 + \kappa/8, 2\}$ [29]. (Beffara builds on a related dimension calculation in [224].) In particular, since the scaling limit of the percolation picture above corresponds to $\kappa = 6$, the Hausdorff dimension of the curve is almost surely $1 + 6/8 = 7/4$.

At least on the square lattice, scaling limits of critical Ising model interfaces and the FK–Ising interfaces (which correspond to the FK cluster model for a particular choice of parameters) have scaling limits given by SLE_3 and $SLE_{16/3}$, respectively, as shown in a remarkable series of papers by Smirnov and coauthors including Chelkak, Duminil-Copin, Hongler, Kemppainen [75, 76, 241]. Loop-erased random walks and uniform spanning tree boundaries have scaling limits given by forms of SLE_2 and SLE_8 , respectively, as shown in a remarkable paper by Lawler, Schramm, and Werner [163]. Level lines of the so-called Gaussian free field are given by SLE_4 as shown by Schramm and the author [230, 231].

2.6. Conformal loop ensembles

What happens if we consider a percolation model with all blue boundary conditions, and then consider the collection of *all* of the loops that form boundaries between blue and yellow regions?



Camia and Newman showed in [66] that in the fine-mesh scaling limit, these loops converge to a random collection of continuum loops, called a *conformal loop ensemble* (CLE) with parameter $\kappa = 6$. Conformal loop ensembles are defined for any $\kappa \in (8/3, 8]$, see the CLE construction by the author in [234] and/or the axiomatic characterization of simple CLE loops by the author and Werner in [238].

We will not give a formal definition of CLE here. But we wish to stress one point. As explained in [238], it turns out that there is a natural way to use an instance of CLE_κ for $\kappa \in (4, 8)$ to generate a *space-filling* version of SLE_κ (even though ordinary SLE_κ is not space-filling for κ in this range) which is somehow the continuum version of the procedure used in Section 2.3 to combine multiple red loops into one loop. One can also reverse the procedure and recover the loops from the space-filling curve. Very roughly speaking, one tries to follow the CLE interfaces, but any time the curve separates a region from the target point, one violates the rules and fills up that region before continuing. The space-filling curve divides space into a continuum tree-and-dual-tree pair. Although we are not giving details here, this implies that, at least in the continuum, if one wants to understand *loop-decorated* random surfaces one can equivalently try to understand *tree-decorated* random surfaces. We remark that these continuum trees also have an interpretation as coalescing rays within a so-called *imaginary geometry* [98, 190–193] where the dual tree corresponds to the coalescing tree of rays drawn in the opposite direction.

3. LIOUVILLE QUANTUM GRAVITY SPHERE: A RANDOM RIEMANNIAN GEOMETRY

The *Liouville quantum gravity sphere* is a random sphere-homeomorphic space whose law depends on the single parameter d (or Q or γ) as mentioned earlier. When $d = 0$ the LQG sphere is called the *pure LQG sphere*. The pure LQG sphere differs from the Brownian sphere in that the former *a priori* comes endowed with a conformal structure but no metric space structure, while the latter is *a priori* endowed with a metric space structure but no conformal structure.

A recent summary of LQG surfaces and the Gaussian free field can be found in the lecture notes by Berestycki and Powell [33, 37]. See also the Astérisque summary of LQG surfaces and the KPZ formula by Garban [113], the *Notices of the AMS* overview article by Gwynne [125], and the ICM proceedings article by Miller. A more recent set of ICM lecture notes by Ding, Dubédat and Gwynne gives an overview of recent works that have established a metric space structure for general values of $d < 25$ [91].

To summarize the latter point, recently various researchers (such as Ang, Basu, Bhatia, Ding, Dubédat, Dunlap, Falconet, Ganguly, Gwynne, Holden, Miller, Pfeffer, Remy, Sepúlveda, and Sun) have contributed to a spectacular international program to show that LQG spheres with $d \neq 0$ can also be given a canonical metric measure space structure, and to prove some basic properties about the resulting random metric spaces [13, 28, 90, 92, 93, 99, 100, 124, 127, 132, 140, 213]. The reader might start by looking at the metric constructions by Gwynne and Miller or by Ding and Gwynne [92, 133]. Viewed as random metric spaces, general LQG spheres can thus also be viewed as generalizations of the Brownian map. On the other hand, we stress that when $d \neq 0$ there is no reason to believe that the geodesic-tree/dual-tree pair is described by anything as simple as the Brownian snake. That is, as far as we know, the simple Brownian map construction given in Section 1 has no simple analog corresponding to $d \neq 0$.

3.1. The Gaussian free field

Brownian motion is a natural random function from \mathbb{R} to \mathbb{R} . A generalization of Brownian motion called the *Gaussian free field* (GFF) is a random (generalized) function from \mathbb{R}^d to \mathbb{R} for any d . Our use of the term GFF in this paper will be limited to $d = 2$.

There are many ways to define the Gaussian free field, see the author's survey [233]. It is a scaling limit of discrete random functions from \mathbb{Z}^2 to \mathbb{R} much as Brownian motion is the scaling limit of random functions from \mathbb{Z} to \mathbb{R} . One particularly concise definition is as follows.

Fix a bounded planar domain D . If f and g are functions on D whose gradients lie in L^2 then we can write $(f, g)_\nabla = (2\pi)^{-1} \int_D \nabla f(z) \cdot \nabla g(z) dz$ for the *Dirichlet inner product* of f and g . Let $H(D)$ be the Hilbert space closure of the space of compactly supported smooth functions with respect to this inner product. Then the *Gaussian free field* on D with zero boundary conditions is the sum $\sum \alpha_i f_i$ where the f_i are an orthonormal basis for $H(D)$ and the α_i are independent standard normal random variables (mean zero, vari-

ance one). The sum a.s. does not converge pointwise or in $H(D)$ but it a.s. does converge in the space of generalized functions (a.k.a. *distributions*) [233].

The GFF can also be defined as the Gaussian random distribution with covariance given by Green’s function $G(x, y)$. Here $G(x, \cdot)$ is given by $-\log|x - \cdot|$ minus the harmonic extension of $-\log|x - \cdot|$ from ∂D to D . The symbols h and ϕ are both commonly used to describe an instance of the Gaussian free field, depending on the context. We will use ϕ in this paper.

The GFF can also be defined on the whole plane, where one simply has $G(x, y) = -\log|x - y|$. In this setting the field ϕ is only defined up to additive constant. But one may nonetheless write $\text{Cov}((\phi, f), (\phi, g)) = \iint -\log|x - y|f(x)g(y)dx dy$ as long as both f and g have mean zero. The integration-by-parts identity

$$(f, g)_{\nabla} = \frac{1}{2\pi} \int \nabla f(z) \cdot \nabla g(z) dz = \frac{-1}{2\pi} \int f(z) \Delta g(z) dz = \frac{-1}{2\pi} (f, -\Delta g)$$

is frequently used.

3.2. Conformal parameterizations

The LQG-sphere has a long history. On the physics side, LQG surfaces come up in certain formulations of string theory and 2D quantum field theories based on the Einstein equations (which in two dimensions reduce to the very simple Liouville equation). This literature is rich and complex, with foundational contributions by Belavin, Brézin, David, Di Francesco, Distler, Dorn, Duplantier, Eynard, Fateev, Itzykson, Kawai, Kazakov, Knizhnik, Kostov, Migdal, Otto, Parisi, Polchinski, Polyakov, Segal, Seiberg, Teschner, Witten, the Zamolodchikov brothers, Zinn-Justin, Zuber, and many others. (This list is far from exhaustive.) We will not attempt to properly survey the physics literature in this paper, but we point the reader to the long list of references in [105] (or the articles cited in Section 4) as a place to start.

On the mathematics side, one might begin with Gauss [115] who explained in 1827 how curvature could be understood as an intrinsic property of a two-dimensional surface, independently of how the surface was “embedded” in a higher-dimensional space. The Riemann mapping theorem (formulated by Riemann in 1851, proved by Osgood in 1900) and the more general Riemann uniformization theorem (conjectured by Klein in 1893, Poincaré in 1892, proved by Poincaré in 1907, Koebe in 1907) also play a central role [119, 255].

It is standard in differentiable geometry to define a surface (or two-dimensional manifold) by covering the surface with a “chart” of open sets that can each be diffeomorphically mapped to a planar domain. Within one of these open sets, parameterized by pairs (x, y) , the “metric” can be written $A(x, y)dx^2 + B(x, y)dx dy + C(x, y)dy^2$. The parameterization is said to be *conformal* if $A = C$ and $B = 0$, and the Riemann uniformization theorem implies that one can always find a conformal parameterization. If we treat the parameterizing domain U as a subset of \mathbb{C} and write $z = x + iy$ then we can write the metric as $e^{\rho(z)}(dx^2 + dy^2)$ and the associated area measure as $e^{\rho(z)} dz$ where dz is Lebesgue measure on U . (In some conventions the definition of ρ may differ by a factor of two; the issue is whether $e^{\rho(z)}$ is interpreted as the length multiplier or the area multiplier.)

The Gaussian curvature is $-e^{-\rho(z)} \Delta \rho(z)$ so the integral of the Gaussian curvature over the set parameterized by a region R is equal to the integral of $\Delta \rho$ over R . In particular, the Gaussian curvature is zero if and only if $\Delta \rho = 0$ so that ρ is harmonic. A function ρ on a bounded domain D is harmonic if and only if it minimizes $\int (\nabla \rho(z) \cdot \nabla \rho(z)) dz = \int (-\rho(z) \Delta \rho(z)) dz$ given its boundary conditions. More generally, the Gaussian curvature is equal to the constant K if and only if $\Delta \rho = -K e^\rho$. The latter equation is called *Liouville's equation* and was formulated by Liouville in 1838 [179]. A function with constant curvature minimizes $\int (\nabla \rho(z) \cdot \nabla \rho(z)) + K e^\rho$, which is a linear combination of the Dirichlet energy $(\rho, \rho)_\nabla$ and the overall surface area. (Polyakov used the latter quantity to define the so-called *Liouville action*, see Section 3.5.2.)

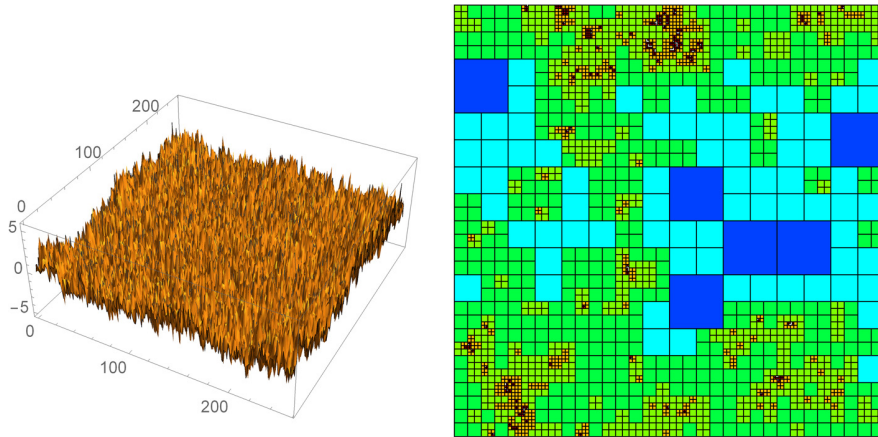
If we want to choose a random perturbation of a flat metric, we need to find a random function ρ that is in some sense a random perturbation of a harmonic function. Roughly speaking, one “randomly perturbs” a harmonic function by replacing the assertion that “the Dirichlet energy of ρ is minimal given its boundary values” with “the probability of a given ρ is proportional to the exponential of minus the Dirichlet energy of ρ .” Formally, this means taking ρ to be a constant $\gamma \in (0, 2)$ times the Gaussian free field. The induced area measure then takes the form $\nu_\phi := e^{\gamma \phi(z)} dz$. This cannot make literal sense (since ϕ is a random distribution) but there are various ways to make sense of this through regularization.

The first rigorous construction of a random measure with the law of ν_ϕ is due to Høegh-Krohn [158] who constructed the object for $\gamma \in [0, \sqrt{2}]$. Høegh-Krohn was motivated by earlier works in the quantum field theory literature that made sense of $V(\phi)$ where ϕ was the Gaussian free field and V was a polynomial. The exponential of the free field (viewed as a quantum field theory) was studied in several papers over the subsequent decade, and was cited by prominent quantum field theorists such as Glimm, Jaffe, and Simon.


In the 1980s Kahane derived a similar construction for all $\gamma \in [0, 2)$ and used the term *Gaussian multiplicative chaos* to describe the associated random measures [153]. Kahane was motivated by the multiplicative cascades popularized by Mandelbrot. (See, e.g., [103] for $\gamma = 2$.) Neither Høegh-Krohn nor Kahane *interpreted* these random measures as the pullback to a planar domain of the area measure on a random surface parameterized by that domain. The author's work with Duplantier [105] constructed the measure ν_ϕ as a weak limit as $\varepsilon \rightarrow 0$ of the random measures $\varepsilon^{\gamma/2} e^{\gamma \phi_\varepsilon(z)} dz$, where $\phi_\varepsilon(z)$ is the mean value of ϕ on the ball boundary $\partial B_\varepsilon(z)$, see also [237]. This construction showed that the measure ν_ϕ is a function of the Gaussian free field ϕ .

Computer visualization can help us generate some intuition about what the measures ν_ϕ look like. The following Mathematica code generates the two figures shown below.

```
K = 8; fieldmultiplier = 1.5; squarefraction = .001;
phi=Re[Fourier[Table[(InverseErf[2 Random[]-1]+I InverseErf[2 Random[]-1])*If[j+k == 2, 0,
1/Sqrt[(Sin[(j-1)*Pi/2^K]^2+Sin[(k-1)*Pi/2^K]^2)], {j, 2^K}, {k, 2^K}]]];
CO = squarefraction Sum[M[[i, j]], {i, 1, 2^K}, {j, 1, 2^K}]; M=Exp[fieldmultiplier phi];
{ListPlot3D[phi], Graphics[Table[Table[If[Sum[M[[2^k m+i, 2^k n+j]], {i, 1, 2^k}, {j, 1, 2^k}]<CO,
{Hue[k/8], EdgeForm[Thin], Rectangle[{2^k m, 2^k n}, {2^k m+2^k, 2^k n+2^k}]}],
{m, 0, 2^(K-k)-1}, {n, 0, 2^(K-k)-1}], {k, 0, K-1}]]}
```



The left figure is an instance ϕ of a discrete version of the Gaussian free field on the 256×256 torus (lines 2 and 3 encode the discrete GFF; see the explanation in [233]). The exponential of *fieldmultiplier* times ϕ (called *M* in the code) describes a random measure on the torus. To obtain the picture on the right one starts with the whole square, then divides it into four equal squares, then divides each of those into four equal squares and so on, except that one stops dividing whenever one reaches a square where the area in the square is less than some constant cutoff (taken in the code to be *squarefraction* times the total area). The squares are colored according to their Euclidean size. Each square S shown in the picture above has area less than the cutoff—but its dyadic parent S' must have area greater than the cutoff (as otherwise S' would have been drawn after S , and would have covered S). In this sense one expects that all of the squares shown are “about the same size” in the random geometry. The blue squares correspond to regions where ϕ is smaller on average (and the measure is less dense) while the orange and red squares correspond to high density areas. One can easily paste the code above into Mathematica and experiment with different variants: for example, shown below are the figures obtained by taking a *fieldmultiplier* of 0.5 (left) and 1 (right) instead of 1.5 above. The measure represented by the left figure is much closer to Euclidean measure.



3.3. LQG surfaces

We use the term *LQG surface* broadly to describe any surface whose area measure is $e^{\gamma\phi(z)} dz$ where ϕ has a law that is locally absolutely continuous with respect to the Gaussian free field. Formally, an LQG surface is an equivalence class of pairs (D, ϕ) , where (D, ϕ) and $(\tilde{D}, \tilde{\phi})$ are equivalent if they are related as in the following diagram with $Q = \gamma/2 + 2/\gamma$.



Using the definition $\nu_\phi := \lim_{\varepsilon \rightarrow 0} \varepsilon^{\gamma^2/2} e^{\gamma\phi_\varepsilon(z)} dz$, one can show that the $\nu_{\tilde{\phi}}$ area measure on \tilde{D} above must a.s. agree with the pullback of the ν_ϕ area measure on D . Roughly speaking, this is because $\exp(\gamma Q \log |\psi'|) = |\psi'|^2 |\psi'|^{\gamma^2/2}$. The $|\psi'|^2$ is part of the usual change of variables formula, and the $|\psi'|^{\gamma^2/2}$ accounts for the fact that ψ maps circles of diameter ε to (approximate) circles of size diameter $|\psi'|\varepsilon$, which affects the $\varepsilon^{\gamma^2/2}$ factor in the regularization.

In addition to the area measure ν_ϕ , it is possible to define a boundary length measure (in the case that ϕ is a GFF with free boundary [195]) or a fractal length associated to a line or an SLE curve or other fractal set, see [38]. All of these measures, as well as the LQG *distance function* discussed above, are preserved by the coordinate change—which makes sense since (D, ϕ) and $(\tilde{D}, \tilde{\phi})$ literally represent the same surface. The image under ψ of the so-called *Liouville Brownian motion* in D as defined in [34, 114] is a Liouville Brownian motion in \tilde{D} .

Furthermore, if x_1, \dots, x_k are points in D , and we write $\tilde{x}_i = \psi(x_i)$, then we say that $(D, \phi, x_1, \dots, x_k)$ and $(\tilde{D}, \tilde{\phi}, \tilde{x}_1, \dots, \tilde{x}_k)$ represent the same “LQG surface with k marked points.” Finally, let us stress that the above definition of LQG surface makes sense for any $Q \geq 0$, not only for $Q \geq 2$. (The values $Q \geq 2$ are those that have the form $Q = \gamma/2 + 2/\gamma$ for some $\gamma > 0$ or, equivalently, those for which $d = 25 - 6Q^2 < 1$.) One just has to accept that the *area measure* is not well defined when $Q < 2$. However, as we mentioned earlier, the metric space structure (along with the fractal measure of some other sets) continues to be well defined when $Q < 2$ (although the diameter of the surface becomes infinite when $Q < 2$, which corresponds to $d > 1$).

3.4. Constructing the LQG sphere

One simple way to define a unit-area LQG sphere is to consider a GFF ϕ on a simply connected bounded domain D with boundary conditions given by a constant C , and to condition on $\nu_\phi(D) = 1$. As $C \rightarrow -\infty$ this object converges in law to the unit area LQG sphere (with the boundary somehow shrinking—in the metric sense—to a single marked point in the limit) as explained in [235]. Another approach involves starting with an infinite cylinder and two marked points (the cylinder’s endpoints), as described in [102]. Yet another approach involves starting with the infinite-volume Polyakov measure (see the next subsection) and “pinning it down” at three points, in a manner described in [87]. The paper by Aru, Huang,

and Sun [23] established the equivalence of the approaches in [102] and [87]. (See also the alternate proof in [16] and the disk analog in [70].) A chart, presenting several equivalent definitions and the relationships between them, is included in [194].

All of these approaches have analogous constructions that produce *unrestricted area* LQG spheres (instead of unit area LQG spheres). In the $C \rightarrow -\infty$ construction above, for example, instead of conditioning on $v_\phi(D) = 1$, one can (for any C value) multiply the measure by a constant factor to ensure that the measure assigned to ϕ distributions with $v_\phi(D) \in [1, 2]$ is constant, and then take the vague $C \rightarrow -\infty$ limit.

3.5. Polyakov’s infinite measure on embedded LQG surfaces

The Polyakov measure is an *infinite* measure on the space of unmarked, unrestricted-area LQG spheres *embedded* in \mathbb{C} . As mentioned in the introduction, it corresponds to an unrestricted-area LQG sphere embedded in all possible ways (with the embedding chosen from Haar measure on the Möbius group). The measure is infinite for two reasons: first, we recall from the introduction that the law of the area for an unrestricted-area LQG sphere has the form $A^{-b} dA$ (or $A^{-b} e^{-\mu A} dA$ for some constant μ in the “off-critical” case) which is an infinite measure for the b values that we will encounter (namely $b > 3$, recall Section 3.7). Second, the embedding is chosen from the Haar measure on the Möbius group, which itself has infinite volume. Also, just to avoid confusion, let us clarify that it is the Polyakov measure on the space of surfaces—not the area measure on any individual surface—that is infinite. In the Polyakov measure, almost all surfaces have finite area (assuming $d \leq 1$; the total area is not defined if $d > 1$). Each embedded surface is described by a generalized function ϕ , so the Polyakov measure can be viewed as an infinite measure on the set of generalized functions.

The Polyakov measure is usually defined in a slightly different way. It is presented as a way to make sense of the expression “ $e^{-S(\phi)} d\phi$ ” where S is the so-called *Liouville action*, which we will discuss below, see the presentation by David, Kupiainen, Rhodes, and Vargas at [87] or the lecture notes by Kupianen at [159]. The simplest way to describe it (in the case $\mu = 0$) is to say that ϕ is an instance of the zero-mean GFF on the sphere plus an independent constant C chosen from the infinite measure $e^{-2Qc} dc$. To obtain the general- μ measure, one simply weights the zero- μ measure by $e^{-\mu A}$, where A is the area of the LQG surface.

The fact that the Liouville action produces a measure on embedded LQG surfaces of the form described above (in particular, a measure that is Möbius invariant) is counterintuitive at first glance, but it is carefully explained, e.g., by Ang, Holden, and Sun in [16], see also the conformal invariance discussion in [159]. There is a certain miracle (related to what we will call “semi-Gaussian” measures) that makes everything work out. Before we discuss the specifics, let us present a couple of simpler semi-Gaussian measures as a warmup.


3.5.1. Semi-Gaussian measures

A *semi-Gaussian* measure is a constant times $e^{-F(v)} dv$ where F is a quadratic that is strictly convex in all directions except one. For example, $\frac{1}{\sqrt{2\pi}} e^{y-x^2} dx dy$ is semi-

Gaussian. It is a product of a normal measure $\frac{1}{\sqrt{2\pi}}e^{-x^2/2}dx$ and an infinite measure $e^y dy$. If we restrict this measure to any nonvertical line, we obtain a finite measure, which is a constant times a normal probability measure. If we restrict to any vertical line, we obtain an infinite measure.

In the figures below, the orange lines are level sets of the function $y - x^2$ (and hence also the function $\frac{1}{\sqrt{2\pi}}e^{y-x^2} dx dy$). In each figure, one can imagine “sampling” (X, Y) from the infinite measure in two steps. First one decides which blue line (X, Y) belongs to. (The blue lines in the three figures are the lines of slope 0, -1 , and 1 , respectively. By integrating the density function along the blue lines, we find that the law of the y -intercept of the blue line is given by a constant multiple of $e^y dy$.) Then *given that* one chooses the location *on* the blue line. The *conditional law* of the location on the blue line is that of a Gaussian random variable centered at a point on the vertical green line (which is the point on the blue line where $y - x^2$ is largest).

```
n = 4; {ParametricPlot[{Table[{t, j/n}, {j, -2n, 2n}], Table[{t, t^2 + j/n}, {j, -2n, 2n}], {0, t}], {t, -2, 2}, PlotRange->{{-1, 1}, {-1, 1}}], ParametricPlot[{Table[{t, -t+j/n}, {j, -2n, 2n}], Table[{t, t^2+j/n}, {j, -2n, 2n}], {-0.5, t}], {t, -2, 2}, PlotRange->{{-1, 1}, {-1, 1}}], ParametricPlot[{Table[{t, t+j/n}, {j, -2n, 2n}], Table[{t, t^2+j/n}, {j, -2n, 2n}], {.5, t}], {t, -2, 2}, PlotRange->{{-1, 1}, {-1, 1}}]}
```



In the left figure, *no matter what* blue line we choose, the conditional expectation of X is 0. In the second and third figures, *no matter what* blue line we choose, the conditional expectation of X is (respectively) $-1/2$ or $1/2$. Changing the blue-line slope somehow has the effect of “shifting” the conditional law of X .

This is a bit counterintuitive. For a closely related example (essentially a rotation of the one above by 45 degrees), suppose

$$e^{-(a-b)^2/2t} e^{(a+b)/2} da db$$

is the (semi-Gaussian) density function for a pair (A, B) . Then we can formally compute $\mathbb{E}[B|A] = A + t/2$ and $\mathbb{E}[A|B] = B + t/2$. This is because if we restrict to a fixed value of a , we obtain a multiple of a Gaussian measure on b values centered at $a + t/2$ (and similarly with a and b reversed). Checking this fact is an elementary complete-the-square calculation: if b is fixed then $-(a^2 - 2ab + b^2)/2t + (a + b)/2$ is $-(a^2 - 2a(b + t/2) + (b + t/2)^2)/2t$ plus a term that does not depend on a .

At a glance, the above seems to suggest that A is $t/2$ units bigger than B on average, and B is $t/2$ units bigger than A on average, which in turn seems contradictory. This is somehow reminiscent of “envelope switching” paradoxes, where after one observes the amount of money in either *one* of two envelopes, one always expects the *other* to contain more, see, e.g., [62]. In fact, this is just the sort of paradox that one encounters when dealing with infinite measures and/or infinite expectations. We find that the following are equivalent:

- First “sample” A from the infinite measure $e^a da$. Then choose B as a normal with variance t and mean $A + t/2$.
- First “sample” B from the infinite measure $e^b db$. Then choose A as a normal with variance t and mean $B + t/2$.

One can imagine similar constructions in higher dimension. For example, we could replace the X in the figures above by a vector (X_1, X_2, \dots, X_n) and replace $y - x^2$ with $y - \sum_{i=1}^n x_i^2$, and replace the blue lines by blue hyperplanes. One could then argue in a similar way that changing the slope of the blue hyperplanes has the effect of shifting the location of the vertical green line.

By taking a limit of an increasing sequence of finite-dimensional subspaces, we can also make sense of an infinite dimensional semi-Gaussian—precisely the same way we make sense of the GFF as an infinite dimensional Gaussian, using the quadratic function $-\frac{1}{2}(\phi, \phi)_\nabla$, or the way we define Brownian motion, using the quadratic function $\frac{1}{2} \int_{-\infty}^{\infty} (\frac{\partial}{\partial t} B(t))^2 dt$.

Let us give an example. Define a process (whose law is an infinite measure) by $M(t) = B(t) + |t|/2 + Y$ where Y has density $e^y dy$ and $B(t)$ is an independent standard Brownian motion (defined for all $t \in \mathbb{R}$, with $B(0) = 0$). This is an infinite measure on paths that at first glance seems to be “mostly supported” on paths that have their minima near 0. If we tried to define an “action” for $M(t)$ it would have the form

$$S(M) = e^{M(0)} + \frac{1}{2} \int_{-\infty}^{\infty} \left(\frac{\partial}{\partial t} (M(t) - |t|/2)^2 \right) dt.$$

On the other hand, some thought reveals that, when t is fixed, the (infinite-measure) law of $B(0)$ and $B(t)$ is equivalent to the law of A and B in the example above, and in fact the law is symmetric with respect to swapping the roles 0 and t . One can also show that given $M(0)$ and $M(t)$ the conditional law of the rest of the process is given by a Brownian bridge on $(0, t)$ and Brownian motions with drift on (t, ∞) and (running time backward) on $(-\infty, 0)$, with the given values of $M(0)$ and $M(t)$ as endpoints. One can make a similar argument using $-t$ and 0 in place of 0 and t and use this to show that the the law of M is invariant w.r.t. to translation by t units (for any t).

This is a rather remarkable fact. The definition of $M(t)$ does not look translation invariant at all—it is clearly centered at 0. But somehow translating the definition by r units to the right does two things: it effectively *weights* the law of M by $e^{M(r)-M(0)} = e^{B(r)-B(0)}$ (this puts a bias in favor of functions that increase on the interval $[0, r]$) and it deterministi-

cally adds the function $f(t) = \frac{|t-r|}{2} - \frac{|t|}{2}$ (which decreases on the interval $[0, r]$ and is flat elsewhere). And these two changes magically cancel each other out.

We remark that another way to construct this measure is by considering a Brownian bridge measure on processes $B(t)$ indexed by $[-T, T]$, with boundary values $B(T) = B(-T) = T/2$, then multiplying the measure by a constant to ensure that the measure of paths with $B(0) < 0$ is some fixed constant, then taking the vague $T \rightarrow \infty$ limit. This construction might make the translation invariance a bit more intuitive.

To extend the story to two dimensions, note that the above construction is still translation invariant if we take the law of Y to be $e^{-2Qy} dy$ and write $M(t) = B(t) - Q|t| + Y$. This just corresponds to taking $\mathcal{B}(T) = \mathcal{B}(-T) = -QT$ in the limiting procedure mentioned above. On the cylinder $\mathcal{C} = \{t + \theta i : t \in \mathbb{R}, i \in [0, 2\pi)\}$ one can then define $\phi(t + \theta i)$ to be $M(t)$ plus the projection of the GFF (on the cylinder) onto the space of functions with mean zero on every fixed- t slice of the cylinder. If restrict to $M(t)$ that have their maximum at $t = 0$ (to eliminate the translation symmetry) we obtain the so-called “ $\alpha = 0$ quantum sphere” [102]. If we *don't* restrict in this way, we obtain the Polyakov measure (with $\mu = 0$), an infinite measure on surfaces that turns out to be invariant w.r.t. *all* Möbius coordinate changes, not only translations.

To see why, note that there are three “mostly flat” surfaces often used to parameterize an LQG sphere: the Riemann sphere $\mathbb{C} \cup \{0\}$, the cylinder $\mathcal{C} \cup \{\pm\infty\}$ and the gluing \mathcal{D} of two unit disks together their boundaries. There are obvious unit-circle-preserving conformal maps between these three spaces. If ψ is the obvious conformal map from \mathcal{C} to \mathcal{D} then $Q \log |\psi'(t + i\theta)| = -Q|t|$. This is the term that comes up in the change of coordinates from \mathcal{D} to \mathcal{C} . If we instead change coordinates from \mathcal{D} to \mathbb{C} , then this term is replaced by $-2Q \max\{\log |z|, 0\}$. Working in \mathbb{C} , the proof that the law of the construction is invariant under translations and rotations of \mathbb{C} (as well as inversions and dilations) is very similar to the one dimensional argument above, and since these operations generate the Möbius group, full Möbius invariance follows, see [16, 159].

3.5.2. Embedded Polyakov sphere

When building physical field theories it is often natural to define a measure on fields ϕ by writing formally $e^{-S(\phi)} d\phi$ where $d\phi$ represents some sort of uniform measure on the space of all functions. In some cases this is hard to make mathematically precise because it is not clear what the measure $d\phi$ corresponds to. One might attempt to approximate $d\phi$ by something well defined (like Lebesgue measure on a finite-dimensional space of piecewise-linear functions) but it may be unclear how to normalize the construction to obtain a non-trivial limit, or how to prove such a limit exists. In the special case of the Liouville action, however, these concerns can be overcome. Let us tell the story with a playful dialog.

INSTRUCTOR: Consider the measure $e^{-A(\phi)-B(\phi)} d\phi$ where $d\phi$ is the uniform measure on the space of all functions.

MATH POLICE: Sorry, there is no such thing as the uniform measure on all functions. Your object is not defined.

INSTRUCTOR: But my A is a quadratic function on the space of ϕ for which it is finite. If I restrict A to a codimension-one subspace then A is the norm for a Hilbert space. So $e^{-A(\phi)}d\phi$ is just a Gaussian Hilbert space cross an infinite measure that looks something like $e^y dy$. Check out Janson’s book on Gaussian Hilbert spaces [150]. These things are certainly well defined.

MATH POLICE: Okay fine, but you still need to weight by the $e^{-B(\phi)}$ factor.

INSTRUCTOR: My friends proved that (with appropriate normalizing) $B(\phi)$ is well-defined and finite for almost all ϕ taken from the measure $e^{-A(\phi)}d\phi$.

MATH POLICE: So $e^{-B(\phi)}$ is the Radon–Nikodym derivative with respect to the semi-Gaussian measure? And B is obviously measurable?

INSTRUCTOR: That’s right.

MATH POLICE: Okay, you’re free to go.

We can take the S in the above dialog to be the *Liouville action* defined as follows:

$$S(\phi) = \int \left(\frac{1}{2} \cdot \frac{1}{2\pi} |\nabla\phi(z)|^2 + \frac{1}{4\pi} QR(z)\phi(z) + \mu e^{\gamma\phi(z)} \right) dz,$$

where R is the Ricci curvature associated to the “reference metric,” see e.g. [87]. For example, we may assume that the reference metric is the ordinary sphere so that R is constant (with total integral 8π by Gauss-Bonnet—recall that the Ricci curvature is twice the Gaussian curvature) so that the second term is $2Q$ times the mean value of ϕ . Alternatively, if the reference metric is the glued pair of disks \mathcal{D} from the previous subsection, then the second term becomes $2Q$ times the mean value of ϕ on the unit circle.

The integral of the first two terms is the quadratic part $A(\phi)$ while the integral of the last term is the $B(\phi)$. The fact that the latter is defined (provided $Q > 2$) follows essentially from the ideas of Høegh-Krohn and Kahane—see also the author’s later work with Duplantier [105] which constructs $e^{\gamma\phi(z)}dz$ as a measure-valued function of the field ϕ and explains the LQG context. In light of the above dialog, the Polyakov measure $e^{-S(\phi)}d\phi$ is in fact rigorously defined; this is explained in more detail in [87] (see also the higher genus version in [88]).

Stories like the above, where S has a quadratic term (relatively simple to handle) and a nonquadratic term (requiring more thought) are relatively common in quantum field theory, and it is certainly not *always* the case that the nonquadratic part simply modifies the Gaussian part in an absolutely continuous way.

The $\mu = 0$ measure is already rather interesting; for example, it can be easily shown to be Möbius invariant (following the LQG coordinate change rules), despite appearing at first glance to be “centered” at a specific location within the Möbius group, a phenomenon analogous to the example at the end of Section 3.5.1. The miracle is that when one applies a Möbius transformation ψ , the $Q \log |\psi'|$ factor one has to add exactly compensates for the effect of the recentering (swapping the curvature measure $R(z)dz$ for its image under ψ), see

[16]. This implies that the Polyakov measure factors as the product of a measure on (unembedded) LQG surfaces and a measure on embeddings (given by the infinite-volume Haar measure on the Möbius group) see [16]. The same is true if one considers a nonzero μ . However, the Polyakov measure with a nonzero μ only makes sense as described above if $d \leq 1$ (since otherwise the natural volume measure is infinite), while the construction of the zero- μ measure makes sense for any $Q \geq 0$. Changing the choice of reference metric is analogous to changing the slope of the blue lines/hyperplanes in the previous section and (up to a constant multiplicative factor) it leaves the measure construction unchanged [87].

On a flat reference metric, the quadratic action is simply the Dirichlet energy of ϕ while the $\mu e^{\gamma\phi(z)}$ term corresponds to the Ke^{ρ} term associated to the Liouville equation with a nonzero K . Polyakov was not the first to work with variants of the free field action, and in [214] and [216] he attributes a closely related “quadratic action” to Douglas, who used it in his work on Plateau’s problem for minimal surfaces. (Douglas was awarded one of the first two Fields medals for this work in 1936 [97].) See also [60, 89] which introduce the action now called the *Polyakov action*, which Polyakov used later in [215].

3.6. More history

The Brownian map and peanosphere constructions were directly motivated by discrete objects—namely random planar map models. The construction of LQG surfaces, on the other hand, was motivated within physics as a “quantization” based on the Dirichlet energy or on the Liouville equation. Why would one expect these objects to be equivalent?

The discrete planar map models were well studied in physics due to their relationship with random matrices and the random particle systems corresponding to random matrix eigenvalues, see the seminal papers from the 1970s by ’t Hooft and by Breézin, Itzykson, Parisi, and Zuber [1, 59], along with more recent overviews by Eynard, Guionnet, and Maurel-Segala [107, 122]. However, it took some time for people to be persuaded that the Liouville theory corresponded to the scaling limit of these models. For example, Polyakov wrote in a memoir [216] (see also his previous memoir [214]) that he did not become convinced of the connection between the discrete models and Liouville quantum gravity until the late 1980s after jointly deriving, with Knizhnik and Zamolodchikov, the so-called *KPZ formula* for certain Liouville quantum gravity scaling dimensions [156] and comparing them with known combinatorial results for the discrete models (for a rigorous approach, see also [32, 103, 105, 113, 220]).

At this point the relationship between the planar map models and the continuum objects is on more solid mathematical ground. We have convergence results of various kinds and formal equivalence proofs between the different continuum objects, see Section 5. But these rely on a lot of machinery that was not yet available in the 1980s. The quantum gravity zipper perhaps gives the cleanest way to rigorously relate the peanosphere construction and the Liouville quantum gravity construction, as described in [102, 235], see also [26].

3.7. Computing the scaling exponent

On the other hand, just as in the peanosphere case, one can compute the scaling exponent quite easily. This is already enough to show that *if* for some γ the unrestricted-area LQG sphere is the scaling limit of the unrestricted-area discrete models, then we must have $\gamma = \sqrt{8/3}$ in the undecorated case and $\gamma = \sqrt{2}$ in the tree-decorated case.

Say we fix the Gaussian free field up to the mean value X on the unit circle. Then this mean value is sampled from the infinite measure $e^{-2Qx} dx$. The total area of $e^{\gamma\phi(z)} dz$ goes like $e^{\gamma x}$ times a constant. So we have to consider the image of the measure $e^{-2Qx} dx$ under the map $\psi(x) = e^{\gamma x}$. By standard change of variables the new measure is (up to a constant factor)

$$\frac{1}{(\psi^{-1})'(A)} e^{-2Q\phi^{-1}(A)} dA = \frac{1}{A} e^{-2Q \log(A)/\gamma} = A^{-1-2Q/\gamma} dA.$$

Here $1 + 2Q/\gamma = (4/\gamma^2 + 2)$ which comes to $7/2$ if $\gamma = \sqrt{8/3}$ and 4 if $\gamma = \sqrt{2}$.

3.8. Random surfaces embedded in d -dimensional space

Suppose we accept, based on the previous discussion, that the $\gamma = \sqrt{8/3}$ theory (which corresponds to $25 - 6Q^2 = 0$) describes the scaling limit of the undecorated $d = 0$ model. We can call this *pure Liouville quantum gravity* model. We would then like to argue that if we *weight* the law of the pure model by the d th power of the GFF partition function (or the corresponding and equivalent “loop soup” partition function) we obtain a new LQG model with a Q parameter that satisfies $d = 25 - 6Q^2$.

This has been explained heuristically in various ways over the years. The author with Ang, Park, and Pfeffer gave a rigorous version of this statement that applied to a certain way of regularizing the random surfaces [17]. We won't give details here but we mention below a few of the ingredients used to make the connection between loop soup weightings and changes to $d = 25 - 6Q^2$.

On a compact surface with boundary, the heat kernel trace can be written $Z = Z(t) = \text{sp } e^{t\Delta} = \sum e^{t\lambda_n}$ where λ_n are the eigenvalues of the Laplace–Beltrami operator Δ . By standard Tauberian theory, the asymptotics of Z (as $t \rightarrow 0$) are closely related to the asymptotics of λ_n (as $n \rightarrow \infty$). Weyl addressed the latter for bounded planar domains D in 1911 [260] (see discussion in [182]) by showing $-\lambda_n \sim \frac{2\pi n}{\text{area}(D)}$ as $n \rightarrow \infty$ which is equivalent to

$$Z \sim \frac{\text{area}(D)}{4\pi t}$$

as $t \rightarrow 0$. In 1966 Kac gave higher-order correction terms for Z on domains with piecewise linear boundaries (accounting for boundary length and corners) in his famously titled “Can you hear the shape of a drum?” which asks what features of the geometry of D can be deduced from the λ_n or, equivalently, from Z [151]. (Short answer is some but not all.)

McKean and Singer (among others) extended these asymptotics from planar domains to smooth manifolds with nonzero curvature [182] where the constant order correction term is a certain curvature integral. For two-dimensional surfaces with boundary, the integral $\int_\delta^\infty Z(t)/t dt$ turns out to describe the *Brownian loop soup measure* of the set of loops

longer than δ (as developed and explored by Lawler, Werner, Dubédat, and others; see [17] for further explanation and references). When the metric takes the form e^ρ times a flat metric, the small δ asymptotics have a constant order correction term that corresponds to the Dirichlet energy of ρ . This is the so-called Polyakov–Alvarez formula, also known as the Polyakov–Ray–Singer or Weyl anomaly formula, and it has played a major role in conformal field theory. See e.g. the early discussion of Laplacian determinants and this formula by Osgood, Phillips and Sarnak [212, 225].

4. CONFORMAL FIELD THEORY AND MULTIPOINT CORRELATIONS

Conformal field theory is a huge subject. For a broader overview of conformal field theory, beyond just Liouville theory, the reader might begin with the well-known (and very long) textbook by Di Francesco, Mathieu, and Sénéchal [95]. Alternatively, the overview at [221] begins with a list of several conformal field theory textbooks.

Liouville theory is one of many conformal field theories, but it is by itself a large and highly studied subject. Here one might begin with the 1990 survey by Seiberg [232], the 1993 lecture notes by Ginsparg and Moore [118], the 1995 survey by Di Francesco, Ginsparg, and Zinn-Justin [94], the 2004 retrospective by Nakayama [210], or the 2014 textbook by Ribault [221] (which among other things explains why Liouville conformal field theory is uniquely characterized by certain axioms).

What is a conformal field theory? An internet search for “A conformal field theory is” reveals several definitions, the first two from Wikipedia:

- a quantum field theory that is invariant under conformal transformations,
- a set of correlation functions that obey a number of axioms,
- a functor [204] between categories satisfying certain “sewing axioms,”
- a Virasoro module [227]

$$V = \bigoplus_{i \in B_1} W(c_i, \phi_i) \otimes W(\bar{c}_i, \bar{\phi}_i)$$

with unitary highest weight modules $W(c_i, \phi_i)$, $W(\bar{c}_i, \bar{\phi}_i)$ subject to [certain] axioms.

The first definition is standard but its precise meaning depends on how one defines a quantum field theory. The third and fourth definitions represent formalization efforts that would go somewhat beyond the scope of this note. So let us focus on the second definition, sometimes called the *conformal bootstrap* approach to CFT. Here a CFT is no more or less than a set of *correlation functions*, and the interpretation of these functions is that they represent (in some sense) expected products of random generalized functions called *fields* evaluated at different points.

The famous 1984 paper by Belavin, Polyakov, and Zamolodchikov (BPZ) [30] argued that conformal invariance symmetries should imply certain properties for these correlation functions, and that this should be sufficient to allow one to explicitly compute the correlation functions for *some* conformal field theories (the so called *minimal models*) including a theory that one would expect to describe the scaling limit of the Ising model. About a decade later Dorn, Otto, Zamolodchikov, and Zamolodchikov were able to compute certain three-point correlations for the Liouville theory [96, 262]. Other correlation functions could then be deduced from these using the BPZ theory [30] and further input proposed by Teschner [244–246]. These arguments involved mathematically nonrigorous steps, such as assuming without proof that formulas defined in one setting could be analytically continued and applied in other settings.

Building on [87], Guillarmou, Kupianen, Rhodes, and Vargas have produced a series of papers that define and derive the correlation functions for Liouville theory mathematically, building on earlier derivations from the physics literature that we mentioned above. The expressions describing the correlation functions are complicated (integrals, special functions, recursive definitions, etc.) but nonetheless explicit. The first work in this series is a proof of the DOZZ formula [160]. The second paper derives an analog of Plancharel’s theorem (which states that the Fourier transform preserves L^2 -norm) along with a certain “spectrum” relevant to this context [120]. The final paper completes the bootstrap program [120] with an extension to n -point functions and higher genus surfaces [121]. We recommend that the reader take a look at the introduction to [120], which summarizes this viewpoint and situates it within the larger enterprise of quantum field theory. Here we will give a much shorter overview of this viewpoint, aiming only to give a simple account of the relationship to the other perspectives in this paper.

The physics literature on conformal field theory can be challenging for mathematicians to follow. It comes with a large and very specialized jargon, and it does not always proceed in the order a mathematician would expect (where one first produces the measure space and the σ -algebra, then constructs the measure, then begins doing calculations). Partly this is because (when quantum wave functions are involved) not everything in quantum field theory *can* be described in a simple probabilistic way—sometimes “observables” are non-commuting operators that can only be defined in indirect ways. Fortunately, Liouville theory does have a simple probabilistic interpretation.

Let us make one more comment on nomenclature. In quantum field theory, a (generalized) function on \mathbb{R}^{n+1} can be interpreted as a *path* on the space of functions on \mathbb{R}^n , and the term *integral* is often a shorthand for *a measure with respect to which one integrates*. In this context, the object we call the Polyakov measure (on the space of surfaces) in Section 3.5 is also called the *Polyakov path integral* (with n taken to be 1). This language evokes the Feynman path integral, an integral over particle trajectories that appears in quantum mechanics. Similar interpretations arise in string theory, where an integral over the space of string trajectories is seen as an integral over the space of embedded surfaces (a.k.a. “worldsheets”). A very informal overview of this viewpoint (featuring Vargas) appears in an online video produced by Quantum Magazine (starting at time 7:20) https://youtu.be/9uASADiYe_8.

4.1. Gaussian case

It is generally instructive to do something easy before doing something hard. So let us start in the simple setting where ϕ is a whole plane Gaussian field with additive constant chosen so that the mean value on the unit circle is 0. (This is also the starting point in [120], for example.) In this case the Green's function is given by

$$G(x, y) = \ln \frac{1}{|x - y|} + \ln |x|_+ + \ln |y|_+,$$

where $|x|_+ = \max\{1, |x|\}$ and the construction of “random fields” of the form $V_\alpha(x) := e^{\alpha\phi(x)}$ is relatively straightforward. In some sense it can already be seen in the work of quantum field theorists of the 1970s, beginning with the work of Høegh-Krohn. As we have already discussed, one way to make sense of this is by writing $V_\alpha^\varepsilon(x) := e^{\alpha^2/2 e^{\alpha\phi_\varepsilon(z)}}$ and $V_\alpha(x) := \lim_{\varepsilon \rightarrow 0} V_\alpha^\varepsilon(x)$, where $\phi_\varepsilon(z)$ is the mean value of ϕ on $\partial B_\varepsilon(z)$ and the convergence holds locally a.s. in the space of random generalized functions (or in the space of random measures). Then for any sufficiently small ε (so that the balls $B_\varepsilon(x_i)$ do not overlap), we have

$$\left\langle \prod V_{\alpha_i}^\varepsilon(x_i) \right\rangle = e^{\sum \alpha_i \alpha_j \tilde{G}(x_i, x_j)}$$

where

$$\tilde{G}(x, y) = \begin{cases} G(x, y), & x \neq y, \\ \log |x|_+ + \log |y|_+, & x = y. \end{cases}$$

This comes up out to be

$$\prod_{i \neq j} |x_i - x_j|^{\alpha_i \alpha_j} \prod_{i, j} (\max\{|x_i|, 1\} \max\{|x_j|, 1\})^{\alpha_i \alpha_j},$$

and, in the case that all the x_i lie inside the unit disc, the expression is simply

$$\prod_{i \neq j} |x_i - x_j|^{\alpha_i \alpha_j}.$$

By starting with $V_{\alpha_i}^\varepsilon$ and taking $\varepsilon \rightarrow 0$, we can give meaning to the multipoint correlation function (also known as *Schwinger function*) written as

$$\left\langle \prod V_{\alpha_i}(x_i) \right\rangle = \prod_{i \neq j} |x_i - x_j|^{\alpha_i \alpha_j}.$$

This can be interpreted as the density function for a two-dimensional “Coulomb gas” of particles with charges α_i . The overall integral over all $x = (x_1, \dots, x_n)$ may be infinite (depending on the α_i values), but it is finite if one restricts to a subset of (x_1, \dots, x_n) values such that the $|x_i|$ are bounded above and the $|x_i - x_j|$ are bounded below. Just to clarify, for now we are making this Coulomb gas calculation only for the centered GFF ϕ , not for the full Polyakov measure from Section 3.5.

As a formal collection of correlation functions (obtained in the $\varepsilon \rightarrow 0$ limit) this expression makes perfect sense for any $\alpha_i \in \mathbb{C}$. Differentiation commutes with expectation, so one can use the same formula to compute expectations of products involving derivatives like $\frac{\partial}{\partial x_i} V_{\alpha_i}(x_i)$. Expectations of fields involving ϕ itself (or polynomials in ϕ) can also be defined mathematically.

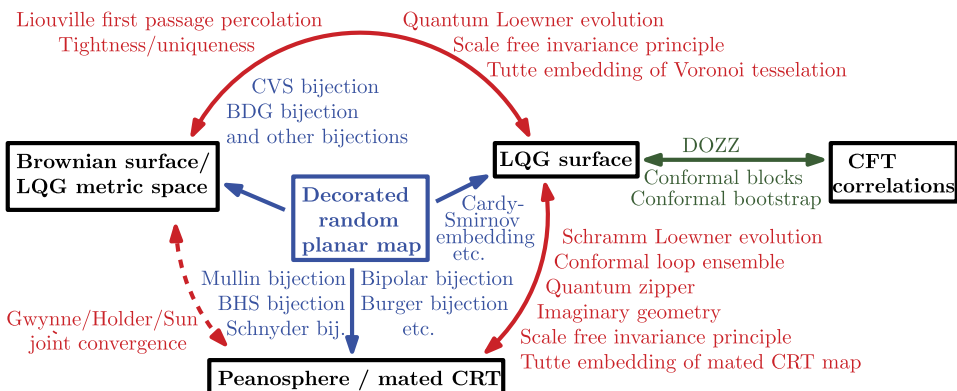
4.2. Incorporating the Liouville term or area conditioning

The ϕ used in Section 4.1 (plus a deterministic function) can be obtained by restricting the Polyakov measure to the set of ϕ whose mean value on the unit circle is zero. But what if we instead restrict the Polyakov measure to the set of surfaces of total quantum area 1 (or weight by $e^{-\mu A}$ where A is the surface area)? The answer is that after such a conditioning or such a weighting, ϕ is no longer Gaussian, and the n -point correlation computation transforms from *easy* to *doable but only barely*. On the other hand, in order to understand some fundamental things like the law of conformal modulus of four points (or n points) sampled independently from the measure on an LQG sphere, one has to address the harder question, and this is precisely what is done in [120, 121, 160]. Although these results are rather recent, they have already inspired a tremendous amount of activity, establishing *exact solvability* for many problems that could previously only be addressed more qualitatively.

Papers by Ang and Sun, including some with coauthors Holden and Remy, give applications of LCFT results to various domains (SLE, CLE, and a variance formula for the peanosphere) [15, 16, 18, 19]. These remarkable papers combine mating-of-trees and conformal-welding techniques with LCFT techniques, leading to rigorous proofs of physics results (such as the FZZ formula and the imaginary DOZZ formula), as well as entirely new results. They have to extend the welding/mating theory to finite-volume surfaces/trees, which is harder than the infinite-volume work, since the finite-volume surfaces lack scale invariance. Recent integrability achievements for the disk include works by Remy and Zhu [218, 219] and a proof of the Fyodorov–Bouchaud formula by Remy [217]. See also the work of Ghosal, Remy, Sun, and Sun on the torus [117].

5. RELATIONSHIPS

Although all four viewpoints are equivalent in some sense, the relationships are somewhat involved. The following diagram summarizes a few of the associated keywords, which we will discuss briefly below.



Miller and the current author recently showed that the LQG and Brownian spheres are *a posteriori* equivalent [195, 197, 198]. That is, we showed that there is a canonical way to endow each object with the other object's structure, and that once this is done the two objects agree in law. This is not all obvious. The papers are long and difficult, and build on hundreds of pages of prior work. The author's other joint works on the LQG construction include [38, 103–105, 154]. This establishes the upper arrow in the special case of the Brownian map. The construction of the metric space structure for general LQG surfaces (with $d \neq 0$) was discussed earlier.

In [235] the author showed that when two infinite half-plane-homeomorphic γ -LQG surfaces are “conformally welded” to one another along their boundaries—and the new surface is conformally mapped to the half-plane—the interface between these curves becomes an SLE_κ curve with $\kappa = \gamma^2$. The sphere version is established in a follow-up work [194], and a disk version appears in [14]. Establishing the correlation formula for $\kappa > 8$ was achieved in [126].

With Duplantier and Miller, the author showed the equivalence of the peanosphere approach and the SLE-decorated-LQG approach in [102], which draws from the imaginary geometry results in [190–193] and the quantum zipper construction in [235]. A remarkable series of papers by Holden and Sun has shown that if one embeds the uniformly random triangulation in the plane in a natural way (inspired by the conformal coordinates Smirnov developed for his proof of Cardy's formula) then the counting measure on the vertices converges (in an appropriate scaling limit) to the Liouville quantum gravity measure, see the overview at [147, 148].

Random planar maps converge to continuum objects in the limit, but there are also natural ways to generate random planar maps from the continuum constructions by “coarse graining” in some sense. The Poisson–Voronoi tessellation (constructed from the Brownian map) and mated-CRT map (constructed from the peanosphere) have been shown to converge to LQG when they are embedded according to the Tutte embedding. This was done in a series of four papers with Gwynne and Miller [136–139]. Effectively, this gives a way of putting a conformal structure on the Brownian map or the matings of trees that is more concrete than the one guaranteed by [195, 197, 198]. The heart of all of this is [138] which gives an invariance principle (i.e., Brownian motion convergence) for random walks in random environments that are “scale free” in the sense that there is no universally typical length scale. (All of the most natural discretizations of Liouville quantum gravity measures are scale free in this sense.)

Gwynne, Holden, and Sun established the *joint convergence* of random triangulations in the metric and peanosphere sense [129]. (Convergence in one topology coupled with convergence in another topology does not imply joint convergence in the product of the two topologies, but it was established in this particular case.) This result builds on earlier convergence work by Gwynne and Miller for percolation-interface-decorated surfaces [134, 135].

A series by the author plus Miller and Werner explores conformal loop ensembles on random surfaces, the surfaces obtained by cutting along the boundaries of these loops, and so forth [201, 202]. These papers extend the conformal welding stories described earlier

and show that much of the intuition one derives from discrete planar maps is correct in the continuum as well (e.g., the surfaces inside and outside of a CLE loop are independent given the interface length) and this leads to a number of interesting computations. These papers build on another recent paper by the same authors [199] which concerns continuum versions of the classical Edwards–Sokal couplings involving FK-models and Ising/Potts models and their variants, see also [200].

6. GAUGE THEORY

Random surfaces are related to many areas of math and physics, including random matrix theory, two-dimensional statistical physics, string theory, and so on. But we should also note that Polyakov’s influential 1981 paper began by mentioning an interest in gauge theory:

“In my opinion at the present time we have to develop an art of handling sums over random surfaces. These sums replace the old-fashioned (and extremely useful) sums over random paths. The replacement is necessary, because today gauge invariance plays the central role in physics. Elementary excitations in gauge theories are formed by the flux lines (closed in the absence of charges) and the time development of these lines forms the world surfaces. All transition amplitude[s] are given by the sums over all possible surfaces with fixed boundary.”
(A. M. Polyakov, Moscow, 1981) [215]

Over 40 years later, many fundamental gauge theory problems remain unsolved, including the famous Clay Millennial Prize Problem, and it is unclear how much random surface theory will help—see, e.g., the “skeptic vs. enthusiast” dialog in Section 2 of [79], written in 1995. Nonetheless, there has been some progress on so-called *gauge string duality*. For more on these efforts, the reader may consult the literature on the AGT conjecture and the AdS/CFT correspondence, or see the recent works of Chatterjee and Jafarov on lattice string trajectories and Yang–Mills theory [71,72,74,149]. Good entry points to the subject for probabilists include Chatterjee’s recent survey [73] and Thierry Lévy’s recent books about continuum Yang–Mills theory in two dimensions [164–166]. The 1986 article by Bridges, Giffen, Durhuus, and Fröhlich may also be read as a first step toward realizing Polyakov’s vision [64].

ACKNOWLEDGMENTS

Random surfaces and curves have featured prominently in most of my own work over the past decade, and I thank the many amazing coauthors and students who have collaborated with me on this subject, including Tom Alberts, Morris Ang, Nathanaël Berestycki, Manan Bhatia, Bertrand Duplantier, Ewain Gwynne, Nina Holden, Richard Kenyon, Sungwook Kim, Greg Lawler, Asad Lodhia, Oren Louidor, Jason Miller, Andrei Okounkov, Minjae Park, Yuval Peres, Joshua Pfeffer, Rémi Rhodes, Oded Schramm, Nike Sun, Xin Sun, Vincent Vargas, Sam Watson, Menglu Wang, Wendelin Werner, David Wilson, Catherine

Wolfram, Hao Wu and Pu Yu. We also thank Tom Alberts, Olivier Bernardi, Bertrand Duplantier, Morris Ang, Sky Cao, Grégory Miermont, Ewain Gwynne, Jason Miller, Minjae Park, Guillaume Remy, Rémi Rhodes, Steffen Rohde, Stanislav Smirnov, Yilin Wang, and Wendelin Werner for their help in reading and improving the manuscript. And finally we thank Thomas Budzinski for allowing us to use his wonderful illustration.

FUNDING

This work was partially supported by NSF Award: DMS 1712862.

REFERENCES

- [1] G. 't Hooft, A planar diagram theory for strong interactions. *Nuclear Phys. B* **72** (1974), 461–473.
- [2] C. Abraham, Rescaled bipartite planar maps converge to the Brownian map. *Ann. Inst. Henri Poincaré Probab. Stat.* **52** (2016), no. 2, 575–595.
- [3] C. Abraham, J. Bettinelli, G. Collet, and I. Kortchemski, Random maps. *ESAIM Proc. Surv.* **51** (2015), 133–149.
- [4] L. Addario Berry and M. Albenque, The scaling limit of random simple triangulations and random simple quadrangulations. *Ann. Probab.* **45** (2017), no. 5, 2767–2825.
- [5] L. Addario Berry and M. Albenque, Convergence of odd-angulations via symmetrization of labeled trees. 2019, arXiv:[1904.04786](https://arxiv.org/abs/1904.04786).
- [6] L. Addario Berry and M. Albenque, Convergence of non-bipartite maps via symmetrization of labeled trees. *Ann. Henri Lebesgue* **4** (2021), 653–683.
- [7] M. Albenque, N. Holden, and X. Sun, Scaling limit of triangulations of polygons. *Electron. J. Probab.* **25** (2020), 1–43.
- [8] M. Albenque and D. Poulalhon, A generic method for bijections between blossoming trees and planar maps. *Electron. J. Combin.* **22** (2015), no. 2, P2–38.
- [9] D. Aldous, The continuum random tree. I. *Ann. Probab.* (1991), 1–28.
- [10] D. Aldous, The continuum random tree. II. An overview. *Stoch. Anal.* **167** (1991), 23–70.
- [11] D. Aldous, The continuum random tree III. *Ann. Probab.* (1993), 248–289.
- [12] J. Ambjørn and T. G. Budd, Trees and spatial topology change in causal dynamical triangulations. *J. Phys. A* **46** (2013), no. 31, 315201, 33.
- [13] M. Ang, H. Falconet, and X. Sun, Volume of metric balls in Liouville quantum gravity. *Electron. J. Probab.* **25** (2020), 1–50.
- [14] M. Ang and E. Gwynne, Liouville quantum gravity surfaces with boundary as matings of trees. *Ann. Inst. Henri Poincaré Probab. Stat.* **57** (2021), no. 1, 1–53.
- [15] M. Ang, N. Holden, and X. Sun, Conformal welding of quantum disks. 2020, arXiv:[2009.08389](https://arxiv.org/abs/2009.08389).
- [16] M. Ang, N. Holden, and X. Sun, Integrability of SLE via conformal welding of random surfaces. 2021, arXiv:[2104.09477](https://arxiv.org/abs/2104.09477).

- [17] M. Ang, M. Park, J. Pfeffer, and S. Sheffield, Brownian loops and the central charge of a Liouville random surface. *Ann. Probab.* (2020, to appear).
- [18] M. Ang, G. Remy, and X. Sun, FZZ formula of boundary Liouville CFT via conformal welding. 2021, arXiv:2104.09478.
- [19] M. Ang and X. Sun, Integrability of the conformal loop ensemble. 2021, arXiv:2107.01788.
- [20] O. Angel, B. Kolesnik, and G. Miermont, Stability of geodesics in the Brownian map. *Ann. Probab.* **45** (2017), no. 5, 3451–3479.
- [21] O. Angel and O. Schramm, Uniform infinite planar triangulations. *Comm. Math. Phys.* **241** (2003), no. 2, 191–213.
- [22] J. Aru, N. Holden, E. Powell, and X. Sun, Mating of trees for critical Liouville quantum gravity. 2021, arXiv:2109.00275.
- [23] J. Aru, Y. Huang, and X. Sun, Two perspectives of the 2D unit area quantum sphere and their equivalence. *Comm. Math. Phys.* **356** (2017), no. 1, 261–283.
- [24] J. Aru, E. Powell, and A. Sepúlveda, Critical Liouville measure as a limit of subcritical measures. *Electron. Commun. Probab.* **24** (2019), 1–16.
- [25] M. Aspenberg and M. Yampolsky, Mating non-renormalizable quadratic polynomials. *Comm. Math. Phys.* **287** (2009), no. 1, 1–40.
- [26] K. Astala, A. Kupiainen, E. Saksman, and P. Jones, Random conformal weldings. *Acta Math.* **207** (2011), no. 2, 203–254.
- [27] J. C. Baez, The Brownian Map. *Not. Amer. Math. Soc.* **68** (2021), 801–803.
- [28] R. Basu, M. Bhatia, and S. Ganguly, Environment seen from infinite geodesics in Liouville Quantum Gravity. 2021, arXiv:2107.12363.
- [29] V. Beffara, The dimension of the SLE curves. *Ann. Probab.* **36** (2008), no. 4, 1421–1452.
- [30] A. A. Belavin, A. M. Polyakov, and A. B. Zamolodchikov, Infinite conformal symmetry in two-dimensional quantum field theory. *Nuclear Phys. B* **241** (1984), no. 2, 333–380.
- [31] J. Beltran and J.-F. Le Gall, Quadrangulations with no pendant vertices. *Bernoulli* **19** (2013), no. 4, 1150–1175.
- [32] I. Benjamini and O. Schramm, KPZ in one dimensional random geometry of multiplicative cascades. *Comm. Math. Phys.* **289** (2009), no. 2, 653–662.
- [33] N. Berestycki, Introduction to the Gaussian free field and Liouville quantum gravity. Lecture notes, 2015.
- [34] N. Berestycki, Diffusion in planar Liouville quantum gravity. *Ann. Inst. Henri Poincaré Probab. Stat.* **51** (2015), no. 3, 947–964.
- [35] N. Berestycki, B. Laslier, and G. Ray, Critical exponents on Fortuin–Kasteleyn weighted planar maps. *Comm. Math. Phys.* **355** (2017), no. 2, 427–462.
- [36] N. Berestycki and J. Norris, Lectures on Schramm–Loewner Evolution. Lecture notes, available on the webpages of the authors, 2014.
- [37] N. Berestycki and E. Powell, Gaussian free field, Liouville quantum gravity and Gaussian multiplicative chaos. *Lecture notes* (2021).

- [38] N. Berestycki, S. Sheffield, and X. Sun, Equivalence of Liouville measure and Gaussian free field. 2014, arXiv:[1410.5407](https://arxiv.org/abs/1410.5407).
- [39] O. Bernardi, Bijective counting of Kreweras walks and loopless triangulations. *J. Combin. Theory Ser. A* **114** (2007), no. 5, 931–956.
- [40] O. Bernardi, Bijective counting of tree-rooted maps and shuffles of parenthesis systems. *Electron. J. Combin.* **14** (2007), no. 1, Research Paper 9, 36 pp.
- [41] O. Bernardi, Tutte polynomial, subgraphs, orientations and sandpile model: new connections via embeddings. *Electron. J. Combin.* **15** (2008), no. 1, Research Paper 109, 53 pp.
- [42] O. Bernardi and N. Bonichon, Intervals in Catalan lattices and realizers of triangulations. *J. Combin. Theory Ser. A* **116** (2009), no. 1, 55–75.
- [43] O. Bernardi and M. Bousquet-Mélou, Counting colored planar maps: algebraicity results. *J. Combin. Theory Ser. B* **101** (2011), no. 5, 315–377.
- [44] O. Bernardi and E. Fusy, Unified bijections for maps with prescribed degrees and girth. *J. Combin. Theory Ser. A* **119** (2012), no. 6, 1351–1387.
- [45] O. Bernardi, N. Holden, and X. Sun, Percolation on triangulations: a bijective path to Liouville quantum gravity. 2018, arXiv:[1807.01684](https://arxiv.org/abs/1807.01684).
- [46] J. Bertoin, T. Budd, N. Curien, and I. Kortchemski, Martingales in self-similar growth-fragmentations and their connections with random planar maps. *Probab. Theory Related Fields* **172** (2018), no. 3, 663–724.
- [47] J. Bertoin, N. Curien, and I. Kortchemski, Random planar maps and growth-fragmentations. *Ann. Probab.* **46** (2018), no. 1, 207–260.
- [48] J. Bettinelli, Geodesics in Brownian surfaces (Brownian maps). *Ann. Inst. Henri Poincaré Probab. Stat.* **52** (2016), no. 2, 612–646.
- [49] J. Bettinelli, E. Jacob, and G. Miermont, The scaling limit of uniform random plane maps, via the Ambjørn–Budd bijection. *Electron. J. Probab.* **19** (2014), 1–16.
- [50] J. Bettinelli and G. Miermont, Compact Brownian surfaces I: Brownian disks. *Probab. Theory Related Fields* **167** (2017), no. 3–4, 555–614.
- [51] P. Biane, Mating of discrete trees and walks in the quarter-plane. *Electron. J. Combin.* **28** (2021), no. 3, Paper No. 3.56, 29 pp.
- [52] G. Borot, J. Bouttier, and E. Guitter, A recursive approach to the $O(n)$ model on random maps via nested loops. *J. Phys. A* **45** (2011), no. 4, 045002.
- [53] M. Bousquet-Mélou and G. Schaeffer, The degree distribution in bipartite planar maps: applications to the Ising model. In *Formal power series and algebraic combinatorics*, pp. 312–323, 2003.
- [54] J. Bouttier, P. Di Francesco, and E. Guitter, Geodesic distance in planar graphs. *Nuclear Phys. B* **663** (2003), no. 3, 535–567.
- [55] J. Bouttier, P. Di Francesco, and E. Guitter, Planar maps as labeled mobiles. *Electron. J. Combin.* **11** (2004), no. 1, R69.
- [56] J. Bouttier, P. Di Francesco, and E. Guitter, Combinatorics of bicubic maps with hard particles. *J. Phys. A: Math. Gen.* **38** (2005), no. 21, 4529.

- [57] J. Bouttier and E. Guitter, The three-point function of planar quadrangulations. *J. Stat. Mech. Theory Exp.* **2008** (2008), no. 07, P07020.
- [58] J. Bouttier and E. Guitter, Confluence of geodesic paths and separating loops in large planar quadrangulations. *J. Stat. Mech. Theory Exp.* **2009** (2009), no. 03, P03001.
- [59] E. Brézin, C. Itzykson, G. Parisi, and J. B. Zuber, Planar diagrams. *Comm. Math. Phys.* **59** (1978), no. 1, 35–51.
- [60] L. Brink, P. Di Vecchia, and P. Howe, A locally supersymmetric and reparametrization invariant action for the spinning string. *Phys. Lett. B* **65** (1976), no. 5, 471–474.
- [61] S. R. Broadbent and J. M. Hammersley, Percolation processes. I. Crystals and mazes. *Proc. Camb. Philos. Soc.* **53** (1957), 629–641.
- [62] J. Broome, The two-envelope paradox. *Analysis* **55** (1995), no. 1, 6–11.
- [63] S. G. Brush, History of the Lenz–Ising model. *Rev. Modern Phys.* **39** (1967), no. 4, 883.
- [64] D. Brydges, C. Giffen, B. Durhuus, and J. Fröhlich, Surface representations of Wilson loop expectations in lattice gauge theory. *Nuclear Phys. B* **275** (1986), no. 3, 459–487.
- [65] T. Budzinski and B. Louf, Local limits of uniform triangulations in high genus. *Invent. Math.* **223** (2021), no. 1, 1–47.
- [66] F. Camia and C. M. Newman, Two-dimensional critical percolation: the full scaling limit. *Comm. Math. Phys.* **268** (2006), no. 1, 1–38.
- [67] F. Camia and C. M. Newman, Critical percolation exploration path and SLE 6: a proof of convergence. *Probab. Theory Related Fields* **139** (2007), no. 3, 473–519.
- [68] A. Caraceni and A. Stauffer, Polynomial mixing time of edge flips on quadrangulations. *Probab. Theory Related Fields* **176** (2020), no. 1, 35–76.
- [69] A. Carrance, Convergence of Eulerian triangulations. *Electron. J. Probab.* **26** (2021), 1–48.
- [70] B. Cerclé, Unit boundary length quantum disk: a study of two different perspectives and their equivalence. *ESAIM Probab. Stat.* **25** (2021), 433–459.
- [71] S. Chatterjee, The leading term of the Yang–Mills free energy. *J. Funct. Anal.* **271** (2016), no. 10, 2944–3005.
- [72] S. Chatterjee, Rigorous solution of strongly coupled $SO(N)$ lattice gauge theory in the large N limit. *Comm. Math. Phys.* **366** (2019), no. 1, 203–268.
- [73] S. Chatterjee, Yang–Mills for probabilists. In *Probability and analysis in interacting physical systems*, pp. 1–16, Springer Proc. Math. Stat. 283, Springer, 2019.
- [74] S. Chatterjee and J. Jafarov, The $1/N$ expansion for $SO(N)$ lattice gauge theory at strong coupling. 2016, arXiv:1604.04777.
- [75] D. Chelkak, H. Duminil Copin, C. Hongler, A. Kemppainen, and S. Smirnov, Convergence of Ising interfaces to Schramm’s SLE curves. *C. R. Math.* **352** (2014), no. 2, 157–161.

- [76] D. Chelkak and S. Smirnov, Universality in the 2D Ising model and conformal invariance of fermionic observables. *Invent. Math.* **189** (2012), no. 3, 515–580.
- [77] L. Chen, Basic properties of the infinite critical-FK random map. *Ann. Inst. Henri Poincaré D* **4** (2017), no. 3, 245–271.
- [78] L. Chen, N. Curien, and P. Maillard, The perimeter cascade in critical Boltzmann quadrangulations decorated by an $O(n)$ loop model. *Ann. Inst. Henri Poincaré D* **7** (2020), no. 4, 535–584.
- [79] S. Cordes, G. Moore, and S. Ramgoolam, Lectures on 2D Yang–Mills theory, equivariant cohomology and topological field theories. *Nuclear Phys. B Proc. Suppl.* **41** (1995), no. 1–3, 184–244.
- [80] R. Cori and B. Vauquelin, Planar maps are well labeled trees. *Canad. J. Math.* **33** (1981), no. 5, 1023–1042.
- [81] H. Coxeter, The four-color map problem, 1840-1890. *Math. Teach.* **52** (1959), no. 4, 283–289.
- [82] H. H. Crapo, The Tutte polynomial. *Aequationes Math.* **3** (1969), 211–229.
- [83] N. Curien, Peeling random planar maps. *Saint-Flour lecture notes*, 2019.
- [84] N. Curien and J.-F. Le Gall, The Brownian plane. *J. Theor. Probab.* **27** (2014), no. 4, 1249–1291.
- [85] N. Curien and J.-F. Le Gall, The hull process of the Brownian plane. *Probab. Theory Related Fields* **166** (2016), no. 1, 187–231.
- [86] N. Curien, J.-F. Le Gall, and G. Miermont, The Brownian cactus I. Scaling limits of discrete cactuses. *Ann. Inst. Henri Poincaré Probab. Stat.* **49** (2013), no. 2, 340–373.
- [87] F. David, A. Kupiainen, R. Rhodes, and V. Vargas, Liouville quantum gravity on the Riemann sphere. *Comm. Math. Phys.* **342** (2016), no. 3, 869–907.
- [88] F. David, R. Rhodes, and V. Vargas, Liouville quantum gravity on complex tori. *J. Math. Phys.* **57** (2016), no. 2, 022302.
- [89] S. Deser and B. Zumino, A complete action for the spinning string. *Phys. Lett. B* **65** (1976), no. 4, 369–373.
- [90] J. Ding, J. Dubédat, A. Dunlap, and H. Falconet, Tightness of Liouville first passage percolation for $\gamma \in (0, 2)$. *Publ. Math. Inst. Hautes Études Sci.* **132** (2020), no. 1, 353–403.
- [91] J. Ding, J. Dubédat, and E. Gwynne, Introduction to the Liouville quantum gravity metric. 2021, arXiv:2109.01252.
- [92] J. Ding and E. Gwynne, Uniqueness of the critical and supercritical Liouville quantum gravity metrics. 2021, arXiv:2110.00177.
- [93] J. Ding, E. Gwynne, and A. Sepúlveda, The distance exponent for Liouville first passage percolation is positive. *Probab. Theory Related Fields* **181** (2021), no. 4, 1035–1051.
- [94] P. Di Francesco, P. Ginsparg, and J. Zinn Justin, 2D gravity and random matrices. *Phys. Rep.* **254** (1995), no. 1–2, 1–133.

- [95] P. Di Francesco, P. Mathieu, and D. Sénéchal, *Conformal field theory*. Springer, 2012.
- [96] H. Dorn and H.-J. Otto, Two- and three-point functions in Liouville theory. *Nuclear Phys. B* **429** (1994), no. 2, 375–388.
- [97] J. Douglas, Solution of the problem of Plateau. *Trans. Amer. Math. Soc.* **33** (1931), no. 1, 263–321.
- [98] J. Dubédat, SLE and the free field: partition functions and couplings. *J. Amer. Math. Soc.* **22** (2009), no. 4, 995–1054.
- [99] J. Dubédat and H. Falconet, Metric growth dynamics in Liouville quantum gravity. 2021, arXiv:2112.13933.
- [100] J. Dubédat, H. Falconet, E. Gwynne, J. Pfeffer, and X. Sun, Weak LQG metrics and Liouville first passage percolation. *Probab. Theory Related Fields* **178** (2020), no. 1, 369–436.
- [101] B. Duplantier, Liouville Quantum Gravity, KPZ and Schramm–Loewner Evolution. In *Proceedings of the ICM*, 2014.
- [102] B. Duplantier, J. Miller, and S. Sheffield, Liouville quantum gravity as a mating of trees. *Astérisque* (2021), no. 427, viii+257 pp.
- [103] B. Duplantier, R. Rhodes, S. Sheffield, and V. Vargas, Critical Gaussian multiplicative chaos: convergence of the derivative martingale. *Ann. Probab.* **42** (2014), no. 5, 1769–1808.
- [104] B. Duplantier, R. Rhodes, S. Sheffield, and V. Vargas, Renormalization of critical Gaussian multiplicative chaos and KPZ relation. *Comm. Math. Phys.* **330** (2014), no. 1, 283–330.
- [105] B. Duplantier and S. Sheffield, Liouville quantum gravity and KPZ. *Invent. Math.* **185** (2011), no. 2, 333–393.
- [106] E. Dynkin and S. Kuznetsov, Markov snakes and superprocesses. *Probab. Theory Related Fields* **103** (1995), no. 4, 433–473.
- [107] B. Eynard, Counting surfaces. *Prog. Math. Phys.* **70** (2016).
- [108] G. E. Farr, Tutte–Whitney polynomials: some history and generalizations. In *Combinatorics, complexity, and chance*, pp. 28–52, Oxford Lecture Ser. Math. Appl. 34, Oxford University Press, 2007.
- [109] C. M. Fortuin, On the random-cluster model II. The percolation model. *Physica* **58** (1972), no. 3, 393–418.
- [110] C. M. Fortuin, On the random-cluster model: III. The simple random-cluster model. *Physica* **59** (1972), no. 4, 545–570.
- [111] C. M. Fortuin and P. W. Kasteleyn, On the random-cluster model: I. Introduction and relation to other models. *Physica* **57** (1972), no. 4, 536–564.
- [112] R. Fritsch and G. Fritsch, *The four-color theorem*. Springer, New York, 1998.
- [113] C. Garban, Quantum gravity and the KPZ formula [after Duplantier–Sheffield]. *Astérisque* **352** (2013), no. 1052, 315–354.
- [114] C. Garban, R. Rhodes, and V. Vargas, Liouville Brownian motion. *Ann. Probab.* **44** (2016), no. 4, 3076–3110.

- [115] K. F. Gauss and P. Pesic, *General investigations of curved surfaces*. Courier Corporation, 2005.
- [116] I. M. Gessel and D. Zeilberger, Random walk in a Weyl chamber. *Proc. Amer. Math. Soc.* **115** (1992), no. 1, 27–31.
- [117] P. Ghosal, G. Remy, X. Sun, and Y. Sun, Probabilistic conformal blocks for Liouville CFT on the torus. 2020, arXiv:2003.03802.
- [118] P. Ginsparg and G. Moore, Lectures on 2D gravity and 2D string theory (TASI 1992). 1992, arXiv:hep-th/9304011.
- [119] J. Gray, On the history of the Riemann mapping theorem. *Supp. Rend.* **19** (1994).
- [120] C. Guillarmou, A. Kupiainen, R. Rhodes, and V. Vargas, Conformal bootstrap in Liouville Theory. 2020, arXiv:2005.11530.
- [121] C. Guillarmou, A. Kupiainen, R. Rhodes, and V. Vargas, Segal’s axioms and bootstrap for Liouville Theory. 2021, arXiv:2112.14859.
- [122] A. Guionnet and E. Maurel-Segala, Combinatorial aspects of matrix models. 2005, arXiv:math/0503064.
- [123] R. K. Guy, C. Krattenthaler, and B. E. Sagan, Lattice paths, reflections, & dimension-changing bijections. *Ars Combin.* **34** (1992), 3–15.
- [124] E. Gwynne, The dimension of the boundary of a Liouville quantum gravity metric ball. *Comm. Math. Phys.* **378** (2020), no. 1, 625–689.
- [125] E. Gwynne, Random surfaces and Liouville quantum gravity. *Not. Amer. Math. Soc.* **64** (2020), no. 4, 484–491.
- [126] E. Gwynne, N. Holden, J. Miller, and X. Sun, Brownian motion correlation in the peanosphere for $\kappa > 8$. *Ann. Inst. Henri Poincaré Probab. Stat.* **53** (2017), no. 4, 1866–1889.
- [127] E. Gwynne, N. Holden, J. Pfeffer, and G. Remy, Liouville quantum gravity with matter central charge in $(1, 25)$: a probabilistic approach. *Comm. Math. Phys.* (2020), 1–53.
- [128] E. Gwynne, N. Holden, and X. Sun, Mating of trees for random planar maps and Liouville quantum gravity: a survey. 2019, arXiv:1910.04713.
- [129] E. Gwynne, N. Holden, and X. Sun, Joint scaling limit of site percolation on random triangulations in the metric and peanosphere sense. *Electron. J. Probab.* **26** (2021), 1–58.
- [130] E. Gwynne, A. Kassel, J. Miller, and D. B. Wilson, Active Spanning Trees with Bending Energy on Planar Maps and SLE-Decorated Liouville Quantum Gravity for $\kappa > 8$. *Comm. Math. Phys.* **358** (2018), no. 3, 1065–1115.
- [131] E. Gwynne, C. Mao, and X. Sun, Scaling limits for the critical Fortuin–Kasteleyn model on a random planar map I: Cone times. *Ann. Inst. Henri Poincaré Probab. Stat.* **55** (2019), no. 1, 1–60.
- [132] E. Gwynne and J. Miller, Conformal covariance of the Liouville quantum gravity metric for $\gamma \in (0, 2)$. *Ann. Inst. Henri Poincaré Probab. Stat.* **57** (2021), no. 2, 1016–1031.

- [133] E. Gwynne and J. Miller, Existence and uniqueness of the Liouville quantum gravity metric for $\gamma \in (0, 2)$. *Invent. Math.* **223** (2021), no. 1, 213–333.
- [134] E. Gwynne and J. Miller, Characterizations of SLE_κ for $\kappa \in (4, 8)$ on Liouville quantum gravity. 2021, arXiv:1701.05174.
- [135] E. Gwynne and J. Miller, Convergence of percolation on uniform quadrangulations with boundary to SLE_6 on $\sqrt{8/3}$ -Liouville quantum gravity. 2021, arXiv:1701.05175.
- [136] E. Gwynne, J. Miller, and S. Sheffield, Harmonic functions on mated-CRT maps. *Electron. J. Probab.* **24** (2019), Paper No. 58, 55 pp.
- [137] 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.* **374** (2020), no. 2, 735–784.
- [138] E. Gwynne, J. Miller, and S. Sheffield, An invariance principle for ergodic scale-free random environments. *Acta Math.* (2021, to appear).
- [139] E. Gwynne, J. Miller, and S. Sheffield, The Tutte embedding of the mated-CRT map converges to Liouville quantum gravity. *Ann. Probab.* **49** (2021), no. 4, 1677–1717.
- [140] E. Gwynne and J. Pfeffer, KPZ formulas for the Liouville quantum gravity metric. 2019, arXiv:1905.11790.
- [141] E. Gwynne and X. Sun, Scaling limits for the critical Fortuin–Kastelyn model on a random planar map III: finite volume case. 2015, arXiv:1510.06346.
- [142] E. Gwynne and X. Sun, Scaling limits for the critical Fortuin–Kasteleyn model on a random planar map II: local estimates and empty reduced word exponent. *Electron. J. Probab.* **22** (2017), 1–56.
- [143] K. Hartnett, A unified theory of randomness, 2016, https://www.quantamagazine.org/20160802-unified_theory_of_randomness.
- [144] K. Hartnett, Random surfaces hide an intricate order, 2019, <https://www.quantamagazine.org/random-surfaces-hide-an-intricate-order-20190702>.
- [145] K. Hartnett, Mathematicians prove 2D version of quantum gravity really works, 2021, <https://www.quantamagazine.org/mathematicians-prove-2d-version-of-quantum-gravity-really-works-20210617>.
- [146] P. J. Heawood, Map color theorems. *Q. J. Math.* **24** (1890), 332–338.
- [147] N. Holden, X. Li, and X. Sun, Natural parametrization of percolation interface and pivotal points. *Ann. Inst. Henri Poincaré Probab. Stat.* **58** (2022), no. 1, 7–25.
- [148] N. Holden and X. Sun, Convergence of uniform triangulations under the Cardy embedding. 2019, arXiv:1905.13207.
- [149] J. Jafarov, Wilson loop expectations in $SU(N)$ lattice gauge theory. 2016, arXiv:1610.03821.
- [150] S. Janson, *Gaussian Hilbert spaces*. 129. Cambridge University Press, Cambridge, 1997.

- [151] M. Kac, Can one hear the shape of a drum? *Amer. Math. Monthly* **73** (1966), 1–23.
- [152] W. Kager and B. Nienhuis, A guide to stochastic Löwner evolution and its applications. *J. Stat. Phys.* **115** (2004), no. 5, 1149–1229.
- [153] J.-P. Kahane, Sur le chaos multiplicatif. *Ann. Sci. Math. Québec* **9** (1985), no. 2, 105–150.
- [154] R. Kenyon, J. Miller, S. Sheffield, and D. B. Wilson, Six-vertex model and Schramm–Loewner evolution. *Phys. Rev. E* **95** (2017), no. 5, 052146.
- [155] R. Kenyon, J. Miller, S. Sheffield, and D. B. Wilson, Bipolar orientations on planar maps and $SLE_1/2$. *Ann. Probab.* **47** (2019), no. 3, 1240–1269.
- [156] V. G. Knizhnik, A. M. Polyakov, and A. B. Zamolodchikov, Fractal structure of 2d-quantum gravity. *Modern Phys. Lett. A* **3** (1988), no. 08, 819–826.
- [157] M. A. Krikun, Uniform infinite planar triangulation and related time-reversed critical branching process. *J. Math. Sci.* **131** (2005), no. 2, 5520–5537.
- [158] R. H. Krohn, A general class of quantum fields without cut-offs in two space-time dimensions. *Comm. Math. Phys.* **21** (1971), no. 3, 244–255.
- [159] A. Kupiainen, Constructive Liouville conformal field theory. 2016, arXiv:1611.05243.
- [160] A. Kupiainen, R. Rhodes, and V. Vargas, Integrability of Liouville theory: proof of the DOZZ Formula. *Ann. of Math.* **191** (2020), no. 1, 81–166.
- [161] G. F. Lawler, An introduction to the stochastic Loewner evolution. In *Random walks and geometry*, pp. 261–293, De Gruyter, 2004.
- [162] G. F. Lawler, Conformally invariant loop measures. In *Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. I. Plenary lectures*, pp. 669–703, 2018.
- [163] G. F. Lawler, O. Schramm, and W. Werner, Conformal invariance of planar loop-erased random walks and uniform spanning trees. In *Selected works of Oded Schramm. Volume 1, 2*, pp. 931–987, Sel. Works Probab. Stat., Springer, 2011.
- [164] T. Lévy, Yang–Mills measure on compact surfaces. *Mem. Amer. Math. Soc.* **166** (2003), no. 790, xiv+122 pp.
- [165] T. Lévy, The master field on the plane. *Astérisque* **388** (2017).
- [166] T. Lévy, Two-dimensional quantum Yang–Mills theory and the Makeenko–Migdal equations. In *Frontiers in analysis and probability—in the spirit of the Strasbourg–Zürich meetings*, pp. 275–325, 2020.
- [167] J.-F. Le Gall, The Brownian snake and solutions of $\Delta u = u^2$ in a domain. *Probab. Theory Related Fields* **102** (1995), no. 3, 393–432.
- [168] J.-F. Le Gall, The topological structure of scaling limits of large planar maps. *Invent. Math.* **169** (2007), no. 3, 621–670.
- [169] J.-F. Le Gall, Geodesics in large planar maps and in the Brownian map. *Acta Math.* **205** (2010), no. 2, 287–360.
- [170] J.-F. Le Gall, Uniqueness and universality of the Brownian map. *Ann. Probab.* **41** (2013), no. 4, 2880–2960.

- [171] J.-F. Le Gall, Random geometry on the sphere. In *Proceedings of the ICM*, 2014.
- [172] J.-F. Le Gall, The Brownian cactus II: Upcrossings and local times of super-Brownian motion. *Probab. Theory Related Fields* **162** (2015), no. 1, 199–231.
- [173] J.-F. Le Gall, Subordination of trees and the Brownian map. *Probab. Theory Related Fields* **171** (2018), no. 3, 819–864.
- [174] J.-F. Le Gall, Brownian disks and the Brownian snake. *Ann. Inst. Henri Poincaré Probab. Stat.* **55** (2019), no. 1, 237–313.
- [175] J.-F. Le Gall, Brownian geometry. *Jpn. J. Math.* **14** (2019), no. 2, 135–174.
- [176] J.-F. Le Gall and G. Miermont, Scaling limits of random trees and planar maps. In *Probability and statistical physics in two and more dimensions*, p. 155, 2010.
- [177] J.-F. Le Gall and F. Paulin, Scaling limits of bipartite planar maps are homeomorphic to the 2-sphere. *Geom. Funct. Anal.* **18** (2008), no. 3, 893–918.
- [178] Y. Li, X. Sun, and S. S. Watson, Schnyder woods, SLE(16), and Liouville quantum gravity. 2017, arXiv:1705.03573.
- [179] J. Liouville, Note sur la Théorie de la Variation des constantes arbitraires. *J. Math. Pures Appl.* (1838), 342–349.
- [180] J.-F. Marckert and A. Mokkadem, Limit of normalized quadrangulations: The Brownian map. *Ann. Probab.* **34** (2006), no. 6, 2144–2202.
- [181] C. Marzouk, Scaling limits of random bipartite planar maps with a prescribed degree sequence. *Random Structures Algorithms* **53** (2018), no. 3, 448–503.
- [182] H. P. Jr. McKean and I. M. Singer, Curvature and the eigenvalues of the Laplacian. *J. Differential Geom.* **1** (1967), 43–69.
- [183] G. Miermont, On the sphericity of scaling limits of random planarquadrangulations. *Electron. Commun. Probab.* **13** (2008), 248–257.
- [184] G. Miermont, Random maps and their scaling limits. In *Fractal geometry and stochastics IV*, pp. 197–224, Progr. Probab. 61, Springer, 2009.
- [185] G. Miermont, The Brownian map is the scaling limit of uniform random plane quadrangulations. *Acta Math.* **210** (2013), no. 2, 319–401.
- [186] G. Miermont, Aspects of random maps. *Saint-Flour lecture notes*, 2014.
- [187] G. Miermont, Compact Brownian surfaces. In *Stochastic Analysis, Random Fields and Integrable Probability—Fukuoka 2019*, vol. 87, pp. 173–199, Mathematical Society of Japan, 2021.
- [188] J. Miller, Liouville quantum gravity as a metric space and a scaling limit. In *Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. IV. Invited lectures*, pp. 2945–2971, 2018.
- [189] J. Miller and W. Qian, Geodesics in the Brownian map: Strong confluence and geometric structure. 2020, arXiv:2008.02242.
- [190] J. Miller and S. Sheffield, Imaginary geometry I: Interacting SLEs. *Probab. Theory Related Fields* **164** (2016), no. 3–4, 553–705.
- [191] J. Miller and S. Sheffield, Imaginary geometry II: Reversibility of $SLE_\kappa(\rho_1; \rho_2)$ for $\kappa \in (0, 4)$. *Ann. Probab.* **44** (2016), no. 3, 1647–1722.

- [192] J. Miller and S. Sheffield, Imaginary geometry III: Reversibility of SLE_κ for $\kappa \in (4, 8)$. *Ann. of Math. (2)* **184** (2016), no. 2, 455–486.
- [193] J. Miller and S. Sheffield, Imaginary geometry IV: Interior rays, whole-plane reversibility, and space-filling trees. *Probab. Theory Related Fields* **169** (2017), no. 3–4, 729–869.
- [194] J. Miller and S. Sheffield, Liouville quantum gravity spheres as matings of finite-diameter trees. *Ann. Inst. Henri Poincaré Probab. Stat.* **55** (2019), no. 3, 1712–1750.
- [195] J. Miller and S. Sheffield, Liouville quantum gravity and the Brownian map I: The $QLE(8/3, 0)$ metric. *Invent. Math.* **219** (2020), no. 1, 75–152.
- [196] J. Miller and S. Sheffield, An axiomatic characterization of the Brownian map. *J. Éc. Polytech. Math.* **8** (2021), 609–731.
- [197] J. Miller and S. Sheffield, Liouville quantum gravity and the Brownian map II: Geodesics and continuity of the embedding. *Ann. Probab.* **49** (2021), no. 6, 2732–2829.
- [198] J. Miller and S. Sheffield, Liouville quantum gravity and the Brownian map III: The conformal structure is determined. *Probab. Theory Related Fields* **179** (2021), no. 3–4, 1183–1211.
- [199] J. Miller, S. Sheffield, and W. Werner, CLE percolations. *Forum Math. Pi* **5** (2017), e4, 102 pp.
- [200] J. Miller, S. Sheffield, and W. Werner, Non-simple SLE curves are not determined by their range. *J. Eur. Math. Soc. (JEMS)* **22** (2020), no. 3, 669–716.
- [201] J. Miller, S. Sheffield, and W. Werner, Simple conformal loop ensembles on Liouville quantum gravity. 2020, arXiv:2002.05698.
- [202] J. Miller, S. Sheffield, and W. Werner, Non-simple conformal loop ensembles on Liouville quantum gravity and the law of CLE percolation interfaces. *Probab. Theory Related Fields* **181** (2021), no. 1–3, 669–710.
- [203] J. Milnor, Pasting together Julia sets: a worked out example of mating. *Exp. Math.* **13** (2004), no. 1, 55–92.
- [204] G. Moore and N. Seiberg, Classical and quantum conformal field theory. *Comm. Math. Phys.* **123** (1989), no. 2, 177–254.
- [205] R. L. Moore, Concerning upper semi-continuous collections of continua. *Trans. Amer. Math. Soc.* **27** (1925), no. 4, 416–428.
- [206] R. Mullin, Enumeration of rooted triangular maps. *Amer. Math. Monthly* **71** (1964), no. 9, 1007–1010.
- [207] R. Mullin, The enumeration of Hamiltonian polygons in triangular maps. *Pacific J. Math.* **16** (1966), no. 1, 139–145.
- [208] R. C. Mullin, On counting rooted triangular maps. *Canad. J. Math.* **17** (1965), 373–382.
- [209] R. C. Mullin, On the enumeration of tree-rooted maps. *Canad. J. Math.* **19** (1967), 174–183.

- [210] Y. Nakayama, Liouville field theory: A decade after the revolution. *Internat. J. Modern Phys. A* **19** (2004), no. 17–18, 2771–2930.
- [211] M. Noy, C. Requilé, and J. Rué, On the expected number of perfect matchings in cubic planar graphs. *Publ. Mat.* **66** (2022), no. 1, 325–353.
- [212] B. Osgood, R. Phillips, and P. Sarnak, Extremals of determinants of Laplacians. *J. Funct. Anal.* **80** (1988), no. 1, 148–211.
- [213] J. Pfeffer, Weak Liouville quantum gravity metrics with matter central charge $c \in (-\infty, 25)$. 2021, arXiv:2104.04020.
- [214] A. Polyakov, Confinement and liberation. In *50 Years of Yang–Mills Theory*, pp. 311–329, World Scientific, 2005.
- [215] A. M. Polyakov, Quantum geometry of bosonic strings. *Phys. Lett. B* **103** (1981), no. 3, 207–210.
- [216] A. M. Polyakov, From quarks to strings. 2008, arXiv:0812.0183.
- [217] G. Remy, The Fyodorov–Bouchaud formula and Liouville conformal field theory. *Duke Math. J.* **169** (2020), no. 1, 177–211.
- [218] G. Remy and T. Zhu, The distribution of Gaussian multiplicative chaos on the unit interval. *Ann. Probab.* **48** (2020), no. 2, 872–915.
- [219] G. Remy and T. Zhu, Integrability of boundary Liouville conformal field theory. 2020, arXiv:2002.05625.
- [220] R. Rhodes and V. Vargas, KPZ formula for log-infinitely divisible multifractal random measures. *ESAIM Probab. Stat.* **15** (2011), 358–371.
- [221] S. Ribault, Conformal field theory on the plane. 2014, arXiv:1406.4290.
- [222] A. Riera, *Brownian Geometry*. PhD thesis, Université Paris-Saclay, 2021.
- [223] L. Rogers, The four colour theorem. *nrich*, September, 2008.
- [224] S. Rohde and O. Schramm, Basic properties of SLE. *Ann. of Math.* (2005), 883–924.
- [225] P. Sarnak, Determinants of Laplacians; heights and finiteness. In *Analysis, et cetera*, pp. 601–622, Elsevier, 1990.
- [226] G. Schaeffer, *Conjugaison d’arbres et cartes combinatoires aléatoires*. PhD thesis, Bordeaux I, 1998.
- [227] M. Schottenloher, *A mathematical introduction to conformal field theory* 759, Springer, Berlin, 2008.
- [228] O. Schramm, Scaling limits of loop-erased random walks and uniform spanning trees. *Israel J. Math.* **118** (2000), no. 1, 221–288.
- [229] O. Schramm, Conformally invariant scaling limits: an overview and a collection of problems. In *Selected works of Oded Schramm*, pp. 1161–1191, Springer, 2011.
- [230] O. Schramm and S. Sheffield, Contour lines of the two-dimensional discrete Gaussian free field. *Acta Math.* **202** (2009), no. 1, 21–137.
- [231] O. Schramm and S. Sheffield, A contour line of the continuum Gaussian free field. *Probab. Theory Related Fields* **157** (2013), no. 1–2, 47–80.
- [232] N. Seiberg, Notes on quantum Liouville theory and quantum gravity. *Progr. Theoret. Phys. Suppl.* **102** (1990), 319–349.

- [233] S. Sheffield, Gaussian free fields for mathematicians. *Probab. Theory Related Fields* **139** (2007), no. 3–4, 521–541.
- [234] S. Sheffield, Exploration trees and conformal loop ensembles. *Duke Math. J.* **147** (2009), no. 1, 79–129.
- [235] S. Sheffield, Conformal weldings of random surfaces: SLE and the quantum gravity zipper. *Ann. Probab.* **44** (2016), no. 5, 3474–3545.
- [236] S. Sheffield, Quantum gravity and inventory accumulation. *Ann. Probab.* **44** (2016), no. 6, 3804–3848.
- [237] S. Sheffield and M. Wang, Field-measure correspondence in Liouville quantum gravity almost surely commutes with all conformal maps simultaneously. 2016, arXiv:1605.06171.
- [238] S. Sheffield and W. Werner, Conformal loop ensembles: the Markovian characterization and the loop-soup construction. *Ann. of Math. (2)* **176** (2012), no. 3, 1827–1917.
- [239] S. Smirnov, Critical percolation in the plane: conformal invariance, Cardy’s formula, scaling limits. *C. R. Acad. Sci., Sér. I Math.* **333** (2001), no. 3, 239–244.
- [240] S. Smirnov, Towards conformal invariance of 2D lattice models. In *International Congress of Mathematicians. Vol. II*, pp. 1421–1451, 2006.
- [241] S. Smirnov, Conformal invariance in random cluster models. I. Holomorphic fermions in the Ising model. *Ann. of Math.* (2010), 1435–1467.
- [242] S. Smirnov, Discrete complex analysis and probability. In *Proceedings of the International Congress of Mathematicians. Volume I*, pp. 595–621, 2010.
- [243] N. Sun, Conformally invariant scaling limits in planar critical percolation. *Probab. Surv.* **8** (2011), 155–209.
- [244] J. Teschner, On the Liouville three-point function. *Phys. Lett. B* **363** (1995), no. 1–2, 65–70.
- [245] J. Teschner, Liouville theory revisited. *Classical Quantum Gravity* **18** (2001), no. 23, R153.
- [246] J. Teschner, A lecture on the Liouville vertex operators. *Internat. J. Modern Phys. A* **19** (2004), 436–458.
- [247] V. Timorin, Topological regluing of rational functions. *Invent. Math.* **179** (2010), no. 3, 461–506.
- [248] W. T. Tutte, A contribution to the theory of chromatic polynomials. *Canad. J. Math.* **6** (1954), 80–91.
- [249] W. T. Tutte, A census of Hamiltonian polygons. *Canad. J. Math.* **14** (1962), 402–417.
- [250] W. T. Tutte, A census of planar triangulations. *Canad. J. Math.* **14** (1962), 21–38.
- [251] W. T. Tutte, A census of slicings. *Canad. J. Math.* **14** (1962), 708–722.
- [252] W. T. Tutte, A census of planar maps. *Canad. J. Math.* **15** (1963), 249–271.
- [253] W. T. Tutte, *Graph theory as I have known it. Vol. 11*. Oxford University Press, Oxford, 2012.

- [254] V. Vargas, Lecture notes on Liouville theory and the DOZZ formula. 2017, arXiv:1712.00829.
- [255] J. Walsh, History of the Riemann mapping theorem. *Amer. Math. Monthly* **80** (1973), no. 3, 270–276.
- [256] D. J. Welsh and C. Merino, The Potts model and the Tutte polynomial. *J. Math. Phys.* **41** (2000), no. 3, 1127–1152.
- [257] W. Werner, Random planar curves and Schramm–Loewner evolutions. In *Lectures on probability theory and statistics*, pp. 107–195, Lecture Notes in Math. 1840, Springer, 2004.
- [258] W. Werner, Conformal restriction properties. In *International Congress of Mathematicians. Vol. III*, pp. 741–762, 2006.
- [259] W. Werner and E. Powell, Lecture notes on the Gaussian free field. *Cours Spéc.* **28** (2022).
- [260] H. Weyl, Über die asymptotische Verteilung der Eigenwerte. *Nachr. Ges. Wiss. Gött., Math.-Phys. Kl.* **1911** (1911), 110–117.
- [261] M. Yampolsky and S. Zakeri, Mating Siegel quadratic polynomials. *J. Amer. Math. Soc.* **14** (2001), no. 1, 25–78.
- [262] A. Zamolodchikov and A. Zamolodchikov, Conformal bootstrap in Liouville field theory. *Nuclear Phys. B* **477** (1996), no. 2, 577–605.

SCOTT SHEFFIELD

77 Mass. Ave., Cambridge, MA, USA, sheffield@math.mit.edu

