

Geometry of Brownian surfaces

Jean-François LE GALL
 Laboratoire de Mathématiques d'Orsay
 Université Paris-Saclay, Centre d'Orsay
 91405 Orsay Cedex, France*

Résumé.

We discuss the random metric spaces called Brownian surfaces, which arise as scaling limits of large graphs embedded in the sphere or in a surface of higher genus. In particular, we give a detailed construction of the Brownian sphere, and we briefly present recent results of Bettinelli and Miermont concerning Brownian surfaces with holes in higher genus. We also state a few open problems and conjectures.

Table des matières

1	Introduction	2
2	Quadrangulations and the CVS bijection	4
2.1	Planar maps	4
2.2	The Cori-Vauquelin-Schaeffer bijection	5
3	The Brownian sphere as the limit of rescaled planar maps	8
3.1	The Gromov-Hausdorff distance	8
3.2	Convergence to the Brownian sphere	8
3.3	The construction of the Brownian sphere	10
3.3.1	The Brownian tree	11
3.3.2	Brownian labels on the Brownian tree	12
3.3.3	Defining the Brownian sphere	12
3.4	k -point functions and Voronoï cells	14
3.5	The free Brownian sphere	15
4	Geodesics in the Brownian sphere	16
5	Planar maps with a boundary and Brownian disks	18
5.1	Quadrangulations with a boundary	18
5.2	The construction of the Brownian disk	19
6	Non-compact models	21
6.1	The Brownian plane	21
6.2	The Brownian plane as the scaling limit of the UIPQ	22
6.3	The Brownian half-plane	23

*Laboratoire de Mathématiques d'Orsay, Université Paris-Saclay, 91405 Orsay Cedex, France

7 Spatial Markov property in Brownian surfaces 24

8 Brownian surfaces in higher genus 26

1 Introduction

The study of the random surfaces that appear as scaling limits of various models of random graphs embedded in surfaces has aroused much interest in the probabilistic community in the recent years. The purpose of the present article is to describe some of these contributions, focussing on the Gromov-Hausdorff convergence of random graphs viewed as discrete metric spaces to the continuous models we call Brownian surfaces, and on the relations existing between these continuous models.

The idea of constructing continuous random geometries as limits of random graphs embedded in surfaces appeared in the physics literature in the years 1980 in the setting of two-dimensional quantum gravity [32, 44, 70]. However, it took a few more years before mathematicians got interested in this question. A pioneering work of Chassaing and Schaeffer [25] exploited the bijection between quadrangulations and labeled trees now known as the CVS bijection to derive certain asymptotic distributions for the metric properties of the vertex set of a uniformly distributed random planar quadrangulation with n faces (equipped with the graph distance) when n tends to ∞ . In particular, this paper showed that the diameter of a random planar quadrangulation with n faces grows like $n^{1/4}$. As already mentioned in the introduction of [25], these results strongly suggest the existence of a “Continuum Random Map that would describe the continuum limit of scaled random quadrangulations” (by the previously mentioned result, the graph distance has to be rescaled by $n^{-1/4}$). For planar triangulations instead of quadrangulations, this problem was stated in a more precise form by Schramm in his plenary paper [71] at the 2006 ICM, in terms of the Gromov-Hausdorff convergence of compact metric spaces. At about the same time, Marckert and Mokkadem [57] proposed a candidate for the scaling limit of planar quadrangulations, which they called the Brownian map (this is now more commonly called the Brownian sphere). However, Marckert and Mokkadem could only prove a weak form of the convergence of (rescaled) random planar quadrangulations to the Brownian sphere, and the question of proving a strong form of this convergence in the Gromov-Hausdorff sense remained open for several years. This problem was solved independently by Miermont [64] and the author of the present work [49] in 2011. The paper [49] also derived a universality property of the Brownian sphere by showing that it appears as the scaling limit of other classes of random planar maps, including triangulations (the case discussed by Schramm), $2p$ -angulations for any integer $p \geq 2$, and bipartite planar maps under Boltzmann distributions. Since then, the universality class of the Brownian sphere has been extended considerably by a number of different authors, see in particular [1–3, 13, 58, 59]. Roughly speaking, large random planar maps where the degree of faces is “not too large” converge to the Brownian sphere. On the other hand, considering random planar maps where the degree of faces may be very large leads to different scaling limits [54].

As the name suggests, the Brownian sphere has the topology of the two-dimensional sphere [55, 62], and an obvious question is to construct Brownian surfaces with other topologies. Bettinelli [11] considered planar quadrangulations with a boundary

(meaning that a distinguished face called the boundary face has an arbitrary even degree, but all other faces are quadrangles) and was able to obtain the existence of sequential Gromov-Hausdorff limits when the boundary size grows like the square root of the number of faces. The uniqueness of the limit, which is called the Brownian disk and is homeomorphic to the closed unit disk, was then established by Bettinelli and Miermont [14]. Relations between the Brownian sphere and the Brownian disk were discussed in [50], which proves in particular that connected components of the complement of a ball in the Brownian sphere are Brownian disks.

Another question was to obtain analogs of the Brownian sphere in higher genus. Here also, the first step is due to Bettinelli [12], who proved the existence of sequential limits (in the Gromov-Hausdorff sense) for uniformly distributed random bipartite quadrangulations on a surface of genus g . The uniqueness of the limit was solved only recently in the important paper of Bettinelli and Miermont [15], which constructs more general Brownian surfaces in arbitrary genus g , with a finite number of boundary components (holes in the surface). This includes of course the Brownian sphere (genus 0, no hole) the Brownian disk (genus 0, one hole), but also the Brownian annulus (genus 0, two holes) discussed recently in [6], and surfaces of genus g without holes.

Finally, non-compact versions of Brownian surfaces have also been studied in genus 0. The Brownian plane, which may be viewed as a tangent cone in distribution to the Brownian sphere at its distinguished point and is homeomorphic to the plane, was introduced in the paper [28]. The Brownian plane is interpreted as the scaling limit of the infinite random lattices known as the uniform infinite planar triangulation (UIPT) [9] and the uniform infinite planar quadrangulation (UIPQ) [26, 45], which have both been studied extensively. The Brownian half-plane, which corresponds to the tangent cone to the Brownian disk at a boundary point and is homeomorphic to the half-plane, appeared in [41] and [10] (the latter paper also introduces the infinite-volume Brownian disk). It is worth noting that both the Brownian plane and the Brownian half-plane play an important role in the paper [15] discussing general (compact) Brownian surfaces.

In the present article, we concentrate on the Gromov-Hausdorff convergence of rescaled discrete random maps and on the construction and properties of the limiting Brownian surfaces viewed as random metric spaces. We however mention that other approaches to the same objects have been investigated. In particular, Miller and Sheffield [66–69] have developed a completely different approach to the Brownian sphere, using a growth model called Quantum Lowner Evolution to define a random metric on the plane. The latter metric can also be obtained as a special case of the Liouville Quantum Gravity (LQG) metrics defined as first-passage percolation distances associated with the exponential of the Gaussian free field in two dimensions. The construction of these LQG metrics was achieved in the paper [35], and their uniqueness was proved in [43]. Finally, we mention the important recent works dealing with the mathematical approach to Liouville conformal field theory (see in particular [33, 38, 46]). In these works, the integrability properties of planar Gaussian multiplicative chaoses lead to exact distributions for quantities related to the LQG metrics. In particular the recent work [6] has computed the law of the conformal modulus of a Brownian annulus.

The paper is organized as follows. In Section 2, we present the CVS bijection between quadrangulations and labeled trees, which has played a major role in the de-

velopment of the subject. The CVS bijection also helps to understand the construction of the Brownian sphere which is presented in Section 3 together with the basic limit theorem for rescaled random planar maps (Theorem 2). We emphasize the role of the Brownian tree in the construction, and we state some basic properties of the Brownian sphere. We also briefly present a conjecture of Chapuy about Voronoï cells in the Brownian sphere, which is one of the most challenging open questions in this area. Section 4 is devoted to geodesics in the Brownian sphere. The construction from the Brownian tree yields detailed information about geodesics from a typical point, but we also discuss geodesics from the exceptional points called geodesic stars, and we state a couple of open problems. In Section 5, we introduce the Brownian disks as scaling limits of quadrangulations with a boundary. We provide a detailed construction of Brownian disks in terms of a Poisson collection of Brownian trees. In Section 6, we present the two basic non-compact models, namely the Brownian plane and the Brownian half-plane. Section 7 discusses several forms of the spatial Markov property for these compact and non-compact models. Finally, Section 8 is a brief presentation of the recent work of Bettinelli and Miermont dealing with Brownian surfaces in higher genus.

2 Quadrangulations and the CVS bijection

2.1 Planar maps

Recall that a planar map is a proper embedding of a finite connected multigraph in the two-dimensional sphere \mathbb{S}^2 , and that two planar maps are identified if they correspond via an orientation-preserving homeomorphism of the sphere. The word “proper” means that edges are not allowed to cross. Also note that we use “multigraph” instead of graph, meaning that we allow self-loops and multiple edges. The example in Fig. 1 has a self-loop and a double edge.

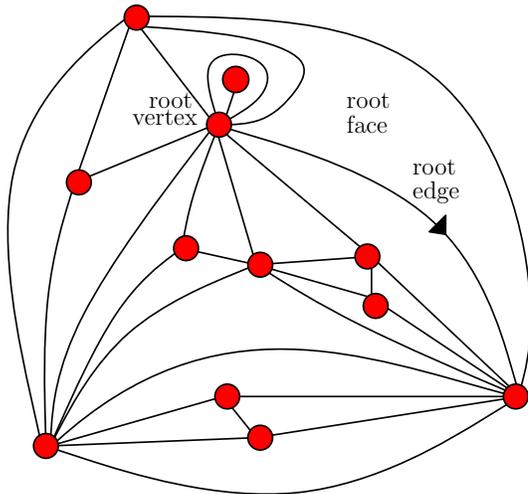


FIGURE 1 – A rooted triangulation with 20 faces.

Faces are the connected components of the complement of edges, or equivalently the regions bounded by the edges. The degree of a face is the number of edges in its boundary, with the convention that, if both sides of an edge are incident to the

same face, this edge is counted twice in the degree of the face (for instance the face inside the self-loop in Fig. 1 has degree 3 though there are only two edges in its boundary). If $p \geq 3$ is an integer, a planar map is called a p -angulation if all its faces have degree p . It is a *triangulation* when $p = 3$, a *quadrangulation* when $p = 4$. The planar map in Fig. 1 is a triangulation with 20 faces.

As a last technical point, we will consider rooted planar maps, meaning that we have distinguished an edge and that this edge is oriented. This oriented edge is called the rooted edge, its origin is the root vertex, and the face to the left of the root edge is the root face (see Fig. 1). Enumeration questions, or bijections with simpler objects such as trees are much more tractable for rooted objects because rooting avoids problems coming from the existence of symmetries.

For every integers $p \geq 3$ and $n \geq 1$, we write \mathbb{M}_n^p for the set of all rooted p -angulations with n faces. When p is odd, \mathbb{M}_n^p is empty for all odd integers n , so in that case we will implicitly assume that n is even. Thanks to the identification modulo homeomorphisms of \mathbb{S}^2 , the set \mathbb{M}_n^p is finite (there are only finitely many “shapes”) and it thus make sense to consider a random p -angulation with n faces which is uniformly distributed over \mathbb{M}_n^p . Let us write M_n for such a uniformly distributed p -angulation with n faces. Our aim is to understand the metric and geometric properties of M_n when n is large.

To this end, we will view planar maps as metric space. If M is a planar map, we denote the vertex set of M by $V(M)$, and we equip $V(M)$ with the usual graph distance d_{gr}^M . The starting point of the work on scaling limits of random planar maps was the question of the convergence of the space $V(M_n)$, equipped with the suitably rescaled metric $d_{\text{gr}}^{M_n}$, when $n \rightarrow \infty$. This convergence makes sense thanks to the Gromov-Hausdorff distance between compact metric spaces that we will introduce in Section 3.1 below.

2.2 The Cori-Vauquelin-Schaeffer bijection

The Cori-Vauquelin-Schaeffer bijection [25,27] (in short the CVS bijection) provides a way of coding rooted quadrangulations by discrete trees equipped with integer labels. The reason why the CVS bijection and its generalizations have been so successful in the study of scaling limits of random planar maps comes from the fact that scaling limits of discrete random trees (possibly equipped with labels) had been studied and well understood by the probabilistic community. Let us immediately mention that the CVS bijection can be extended to much more general planar maps than quadrangulations : See in particular the paper [16] by Bouttier, Di Francesco and Guitter, which was used in [49] to deal with triangulations and $2p$ -angulations. Here however, we will concentrate on the case of quadrangulations, because the bijection with trees is easier to describe in that case.

First recall that a *plane tree* τ is a rooted and ordered discrete tree. Each vertex of a plane tree can be represented as a finite word made of positive integers, in such a way that the empty word \emptyset corresponds to the root, and, for instance, the word 21 corresponds to the first child of the second child of the root. This should be clear from the left side of Fig. 2. In view of the connection with planar maps, we may and will assume that plane trees are drawn in the plane (or rather on the sphere) in the way illustrated in the left side of Fig. 2, so that in particular the edges connecting a vertex to its parent, its first child, its second child, etc., appear in clockwise order

around that vertex.

A *labeled tree* is a plane tree τ , with vertex set $V(\tau)$, whose vertices are assigned integer labels $(\ell_v)_{v \in V(\tau)}$ in such a way that the following two properties hold :

- (i) $\ell_\emptyset = 0$;
- (ii) $|\ell_v - \ell_{v'}| \leq 1$ whenever $v, v' \in V(\tau)$ are adjacent.

The circled figures in the left side of Fig. 2 show a possible assignment of labels. For every $n \geq 2$, let \mathbb{T}_n stand for the set of all labeled trees with n edges.

We also need to introduce rooted and pointed quadrangulations. A rooted and pointed quadrangulation is a rooted quadrangulation given with a distinguished vertex (which can be any vertex, including the root vertex). For every $n \geq 2$, let $\mathbb{M}_n^{4,\bullet}$ stand for the set of all rooted and pointed quadrangulations.

The CVS bijection is a one-to-one correspondence between the sets $\mathbb{M}_n^{4,\bullet}$ and $\mathbb{T}_n \times \{-1, 1\}$. To explain this correspondence, let us start from a labeled tree $(\tau, (\ell_v)_{v \in V(\tau)})$ in \mathbb{T}_n and a sign $\varepsilon \in \{-1, +1\}$. We need to consider corners of the tree τ : A corner incident to a vertex v of τ is an angular sector between two successive edges incident to v . By convention, the root corner c_0 is the corner “below” the root vertex. The set of all corners is given a cyclic ordering by moving clockwise around the tree : starting from the root corner c_0 , the $2n$ corners can be listed as $c_0, c_1, \dots, c_{2n-1}$ in cyclic ordering, as shown in the middle part of Fig. 2. We finally agree that every corner inherits the label of the vertex to which it is incident.

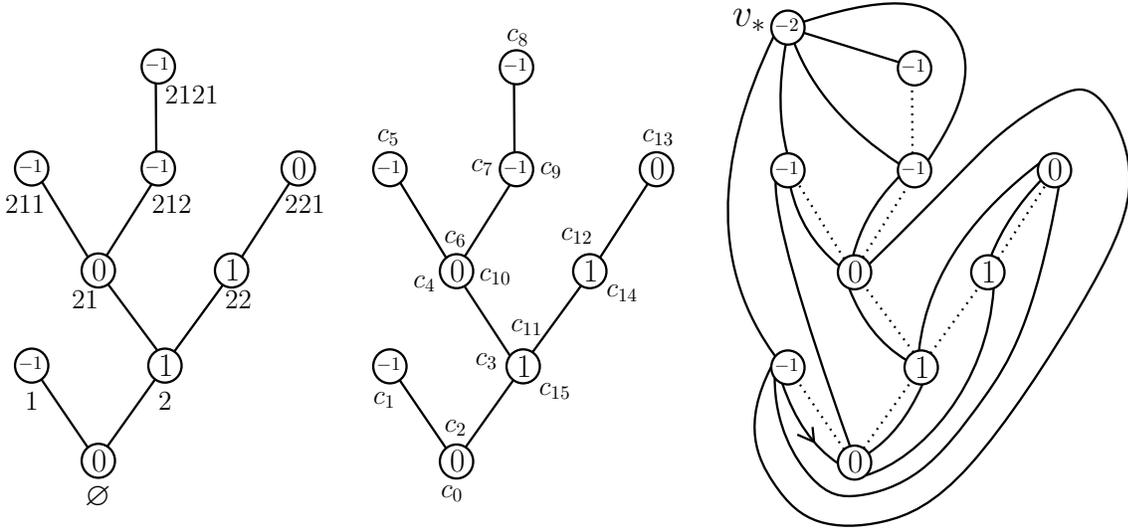


FIGURE 2 – The CVS bijection. Left : a labeled tree with 8 edges. Middle : the sequence c_0, c_1, \dots, c_{15} of corners enumerated in cyclic order. Right : the edges of the associated rooted and pointed quadrangulation with 8 faces (case $\varepsilon = -1$).

From the labeled tree $(\tau, (\ell_v)_{v \in V(\tau)})$, we can construct a rooted and pointed quadrangulation M as follows. First, the vertex set of M is the union of the vertex set of τ and an extra vertex v_* , which by convention is assigned the label

$$\ell_{v_*} = \min_{v \in V(\tau)} \ell_v - 1.$$

Then, in order to obtain the edges of the quadrangulation M , we proceed in the following way. For every corner c of τ , with label ℓ_c , we draw an edge starting from

this corner and ending at the next corner of τ (next refers to the cyclic ordering) with label $\ell_c - 1$ — this corner will be called the successor of c . This makes sense unless ℓ_c is equal to the minimal label on the tree τ , in which case we draw an edge starting from c and ending at v_* . All these edges can be drawn, in a unique manner (up to homeomorphisms of the sphere), in such a way that they do not cross and do not cross the edges of τ , and the resulting planar map is a quadrangulation (see Fig. 2 for an example).

We still have to define the root of the quadrangulation M and its distinguished vertex. The root edge is the edge starting from c_0 and ending at the successor of c_0 , and its orientation is determined by the sign ε : The root vertex is \emptyset if and only if $\varepsilon = +1$. Finally the distinguished vertex of M is v_* , and we have indeed obtained a rooted and pointed quadrangulation.

Proposition 1 *The preceding construction yields a bijection from $\mathbb{T}_n \times \{-1, 1\}$ onto $\mathbb{M}_4^{n, \bullet}$. Moreover, if the rooted and pointed quadrangulation M is the image of the pair $((\tau, (\ell_v)_{v \in V(\tau)}), \varepsilon)$ under this bijection, the vertex set $V(M)$ is canonically identified with $V(\tau) \cup \{v_*\}$ where v_* is the distinguished vertex of M , and with this identification we have, for every $v \in V(\tau)$,*

$$d_{\text{gr}}^M(v_*, v) = \ell_v - \min_{u \in V(\tau)} \ell_u + 1. \quad (1)$$

Let us explain why property (1) holds. Let v be a vertex of M distinct from v_* , so that v is identified to a vertex of τ . Choose any corner c incident to v in the tree τ . The construction of edges in the CVS bijection shows that there is an edge connecting c to a corner c' of a vertex v' with label $\ell_v - 1$. But similarly, there is an edge of M connecting the corner c' of v' to a corner of a vertex with label $\ell_v - 2$. We can continue inductively, and we get a path in M of length $\ell_v - \min_{u \in V(\tau)} \ell_u$ connecting v to a vertex with minimal label, which itself (by the rules of the CVS bijection) is adjacent to v_* in M . In this way we get the upper bound

$$d_{\text{gr}}^M(v_*, v) \leq \ell_v - \min_{u \in V(\tau)} \ell_u + 1.$$

The corresponding lower bound is also very easy, using the fact that $|\ell_v - \ell_{v'}| = 1$ whenever v and v' are adjacent in M , again by the construction of the CVS bijection.

Note that (1) only gives information about distances from the distinguished vertex v_* , which is far from sufficient if one is interested in the Gromov-Hausdorff convergence we will discuss. If v and v' are two arbitrary vertices of M , there is however a useful upper bound for the graph distance $d_{\text{gr}}^M(v, v')$. To state this bound, recall that $c_0, c_1, \dots, c_{2n-1}$ is the sequence of corners of the tree τ associated with M listed in the cyclic ordering, which was already used in the definition of the CVS bijection. For every $i \in \{0, 1, \dots, 2n-1\}$, let v_i be the vertex corresponding to the corner v_i . Then, if $0 \leq i < j \leq 2n-1$, we have

$$d_{\text{gr}}^M(v_i, v_j) \leq \ell_{v_i} + \ell_{v_j} - 2 \max \left(\min_{k \in [i, j]} \ell_{v_k}, \min_{k \in [j, 2n-1] \cup [0, i]} \ell_{v_k} \right) + 2. \quad (2)$$

The proof is easy. Consider the path γ from the corner c_i to v_* constructed as in the proof of (1), and the similar path from the corner c_j . A simple argument shows that these two geodesic paths coalesce at a vertex whose label is the maximum appearing in (2) minus 1. The concatenation of these two paths up to their coalescence time thus gives a path from v_i to v_j whose length is the right-hand side of (2).

3 The Brownian sphere as the limit of rescaled planar maps

3.1 The Gromov-Hausdorff distance

To make sense of the convergence of rescaled discrete planar maps to a continuous object, we will use the Gromov-Hausdorff distance between compact metric spaces, which we now introduce (we refer to [21] for a more detailed presentation). We first recall that, if K_1, K_2 are two compact subsets of a metric space (E, d) , the Hausdorff distance between K_1 and K_2 is defined by

$$d_{\text{Haus}}^E(K_1, K_2) = \inf\{\varepsilon > 0 : K_1 \subset U_\varepsilon(K_2) \text{ and } K_2 \subset U_\varepsilon(K_1)\}$$

where $U_\varepsilon(K_1) = \{x \in E : d(x, K_1) \leq \varepsilon\}$ is the ε -enlargement of K_1 .

Definition 1 (Gromov-Hausdorff distance) *Let (E_1, d_1) and (E_2, d_2) be two compact metric spaces. The Gromov-Hausdorff distance between E_1 and E_2 is*

$$d_{\text{GH}}(E_1, E_2) = \inf\{d_{\text{Haus}}^E(\psi_1(E_1), \psi_2(E_2))\}$$

where the infimum is over all isometric embeddings $\psi_1 : E_1 \rightarrow E$ and $\psi_2 : E_2 \rightarrow E$ of E_1 and E_2 into the same metric space (E, d) .

Let \mathbb{K} stand for the set of all compact metric spaces, where as usual two compact metric spaces are identified if they are isometric. Then the Gromov-Hausdorff distance d_{GH} is a metric on \mathbb{K} , and furthermore $(\mathbb{K}, d_{\text{GH}})$ is complete and separable. In other words, $(\mathbb{K}, d_{\text{GH}})$ is a Polish space, which makes it especially suitable to study the convergence in distribution of random variables with values in \mathbb{K} .

One can prove that a sequence (E_n) of compact metric spaces converges to a limiting space E_∞ in \mathbb{K} if and only if all spaces E_n and the limit E_∞ can be embedded isometrically in the same metric space E in such a way that the convergence holds in the sense of the Hausdorff distance.

There are many variants of the Gromov-Hausdorff distance involving metric spaces marked with one or several distinguished points (or more generally one or several closed subsets) and equipped with one or several Borel measures. These variants can be used to give stronger forms of the convergence theorems that will be stated below, especially in Section 8. For the sake of simplicity, we will content ourselves with the Gromov-Hausdorff distance defined as above.

3.2 Convergence to the Brownian sphere

Let $p \geq 3$ be an integer, and let M_n be a uniformly distributed p -angulation with n faces (recall that n is even when p is odd). We will now state the convergence in distribution of $(M_n, n^{-1/4} d_{\text{gr}}^{M_n})$. Note that we rescale the graph distance $d_{\text{gr}}^{M_n}$ by the factor $n^{-1/4}$. The reason for this rescaling is easily understood (at least for the case $p = 4$ of quadrangulations) using the CVS bijection. In this bijection, distances from the distinguished point are of the same order as labels on the associated labeled tree (cf. formula (1)), but since the height of a plane tree with n edges is of order \sqrt{n} and since labels evolve like random walk along the tree, it is intuitively clear that labels are typically of order $\sqrt{\sqrt{n}} = n^{1/4}$.

The following theorem was obtained in the case $p = 4$ by Miermont [64] and independently in [49] for the cases $p = 3$ and p even, and finally by Addario-Berry and Albenque [3] in the case of odd $p \geq 5$.

Theorem 2 (The scaling limit of p -angulations) *There exists a constant $c_p > 0$ such that we have the following convergence in distribution,*

$$(V(M_n), c_p n^{-1/4} d_{\text{gr}}^{M_n}) \xrightarrow[n \rightarrow \infty]{(d)} (\mathbf{m}_\infty, D)$$

in the Gromov-Hausdorff sense. The limit (\mathbf{m}_∞, D) is a random compact metric space (that is, a random variable with values in \mathbb{K}) that does not depend on p and is called the Brownian sphere.

We note that the role of the scaling constants c_p in the theorem is only to ensure that the limit does not depend on p . With the normalization of \mathbf{m}_∞ that will be given below, one has the explicit expressions :

$$c_3 = 6^{1/4} \quad , \quad c_p = \left(\frac{9}{p(p-2)} \right)^{1/4} \quad \text{if } p \text{ is even.}$$

There is no such formula when $p \geq 5$ is odd.

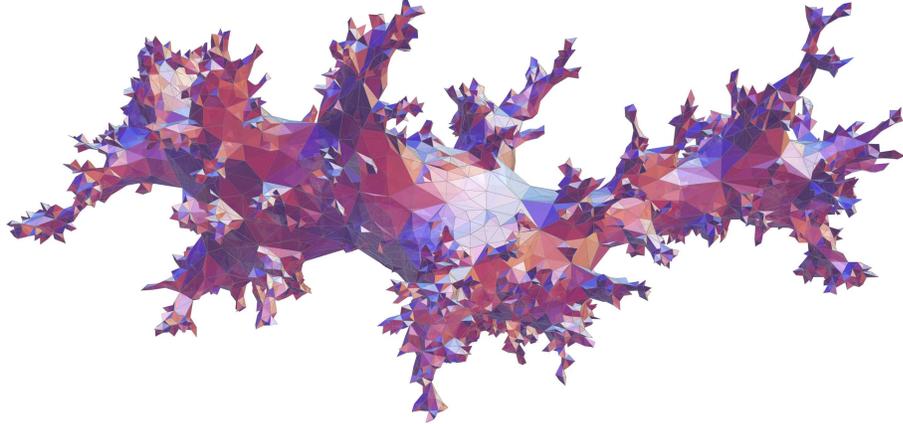


FIGURE 3 – A large planar triangulation represented in space (simulation : N. Curien).

To get a better understanding of the convergence in distribution stated in the theorem, one may use a classical representation theorem of Skorokhod. According to this theorem, one can couple all planar maps M_n in such a way that the convergence $(V(M_n), c_p n^{-1/4} d_{\text{gr}}^{M_n}) \rightarrow (\mathbf{m}_\infty, D)$ now holds almost surely for the Gromov-Hausdorff distance : By a preceding observation, this implies that we can embed isometrically all spaces $(V(M_n), c_p n^{-1/4} d_{\text{gr}}^{M_n})$, as well as the limit (\mathbf{m}_∞, D) , into the same metric space (E, Δ) , in such a way that $V(M_n)$ converges to \mathbf{m}_∞ in the sense of the Hausdorff distance between compact subsets of (E, Δ) .

In much of the previous work in this area, the Brownian sphere is called the Brownian map after Marckert and Mokkadem [57] who obtained a weak form of the theorem in the case of quadrangulations. The name “Brownian sphere” however seems more appropriate in view of Theorem 4 below and of the related objects called the Brownian disk and the Brownian plane that we shall discuss later.

The fact that the limit does not depend on p is of course an important feature of Theorem 2. Roughly speaking, it means that at large scales the metric properties of a typical (large) planar map are the same if this planar map is a triangulation, or a

quadrangulation, or a p -angulation. This is the universality property of the Brownian sphere, which has been confirmed in many subsequent works : In particular, analogs of Theorem 2, always with the same limit (\mathbf{m}_∞, D) hold for general planar maps with a fixed number of edges [13], for bipartite planar maps with a fixed number of edges [1], for simple triangulations or quadrangulations (where self-loops and multiple edges are not allowed) [2], for planar maps with a prescribed degree sequence [58], etc. We also mention that results similar to Theorem 2 hold if the graph distance is replaced by a “local modification” : The paper [29] considers the so-called first-passage percolation distance on random triangulations (independent random weights are assigned to the edges and the distance between two vertices is the minimal total weight of a path between them). Perhaps surprisingly, this local modification does not change the scaling limit, which is still the Brownian sphere up to a deterministic scale factor for the distance.

It is implicit in Theorem 2 that the limit (\mathbf{m}_∞, D) is not the degenerate space with a single point. We make this more explicit in the following two theorems that give some useful information about the Brownian sphere.

Theorem 3 ([47]) *The Hausdorff dimension of (\mathbf{m}_∞, D) is almost surely equal to 4.*

Theorem 4 ([55, 62]) *The compact metric space (\mathbf{m}_∞, D) is almost surely homeomorphic to the 2-sphere \mathbb{S}^2 .*

Both these theorems can be deduced from the construction of the Brownian sphere from Brownian motion indexed by the Brownian tree that will be given below. The proof of Theorem 3 is in fact relatively easy, but that of Theorem 4 is more intricate (in [55], the proof relies in part on an old theorem of Moore giving conditions for a quotient space of the sphere to be homeomorphic to the sphere).

Since planar maps are defined as graphs embedded in the sphere, and since we take a limit where the number of vertices tends to infinity, it is maybe not surprising that the limiting metric space has the topology of the sphere. Still, Theorem 4 implies a non-trivial combinatorial fact about the non-existence of small “bottlenecks” in a large planar map : Informally, for a random triangulation with n faces, the probability that there exists a cycle with length $o(n^{1/4})$ such that both sides of the cycle (meaning both components of its complement) have a diameter greater than $\delta n^{1/4}$, for some fixed $\delta > 0$, will tend to 0 as $n \rightarrow \infty$. The question of the existence of small separating cycles in random planar maps has been investigated recently in connection with isoperimetric inequalities [53].

3.3 The construction of the Brownian sphere

We will now explain the construction of the Brownian sphere (\mathbf{m}_∞, D) appearing in Theorem 2. The approach that we will describe, which is essentially due to Markert and Mokkadem [57], was the first construction of the Brownian sphere, but we emphasize that other approaches, which we do not discuss here, have been developed since : see [67–69] for a construction via the Quantum Loewner Evolution, and [35] and [43] for the construction and uniqueness of the Liouville quantum gravity metric (a special case of which is related to the Brownian sphere metric).

Our construction of the Brownian sphere is a kind of continuous analog of the CVS bijection described in Section 2.2. The role of the (discrete) labeled trees will

be played by a continuous random tree (called the Brownian tree) equipped with real labels, which evolve according to Brownian motion along the line segments of the tree — this is similar to the discrete setting where labels evolves like a random walk along the tree. The distance on the Brownian sphere will satisfy properties analogous to (1) and (2). In contrast with the discrete setting, the set \mathbf{m}_∞ will not exactly coincide with the set of points on the (continuous) tree, because we will need to make certain identifications.

3.3.1 The Brownian tree

Let us start by introducing the Brownian tree. Recall that an \mathbb{R} -tree is a metric space (\mathcal{T}, d) such that, for every $a, b \in \mathcal{T}$ there is, up to reparameterization, a unique continuous injective path γ from a to b , and the range of γ (called the line segment between a and b) is isometric to the line segment $[0, d(a, b)]$. An \mathbb{R} -tree \mathcal{T} is rooted if there is a distinguished point $\rho \in \mathcal{T}$, which is called the root.

We will consider only compact \mathbb{R} -trees, and we will use the fact that such trees can be coded by continuous functions. Let $\sigma > 0$, and let $h : [0, \sigma] \rightarrow \mathbb{R}_+$ be a nonnegative continuous function on $[0, \sigma]$ such that $h(0) = h(\sigma) = 0$. For every $s, t \in [0, \sigma]$, we set

$$d_h(s, t) := h(s) + h(t) - 2 \min_{s \wedge t \leq r \leq s \vee t} h(r).$$

We note that d_h is a pseudo-metric on \mathbb{R}_+ (this means that d_h satisfies all properties of a distance, except that we may have $d_h(s, t) = 0$ for some $s \neq t$). We then introduce the associated equivalence relation on $[0, \sigma]$, defined by setting $s \sim_h t$ if and only if $d_h(s, t) = 0$, or equivalently

$$h(s) = h(t) = \min_{s \wedge t \leq r \leq s \vee t} h(r).$$

Then, d_h induces a distance on the quotient space $[0, \sigma] / \sim_h$. Moreover, one can verify [36] that the quotient space $\mathcal{T}_h := [0, \sigma] / \sim_h$ equipped with the distance d_h is a compact \mathbb{R} -tree.

The \mathbb{R} -tree (\mathcal{T}_h, d_h) is called the tree coded by h . The canonical projection from \mathbb{R}_+ onto \mathcal{T}_h is denoted by p_h . By definition, \mathcal{T}_h is rooted at $\rho_h = p_h(0)$. In addition, we equip \mathcal{T}_h with a *volume measure*, which is defined as the pushforward of Lebesgue measure on $[0, \sigma]$ under p_h .

Remark. It is not hard to verify that any compact \mathbb{R} -tree can be represented as \mathcal{T}_h for some (not unique) function h .

The coding by a function makes it possible to define “lexicographical” intervals on the tree. Let us explain this. If $s, t \in [0, \sigma]$ and $s > t$, we make the convention that $[s, t] = [s, \sigma] \cup [0, t]$ (of course, if $s \leq t$, $[s, t]$ is the usual interval). If $a, b \in \mathcal{T}_h$, there is a smallest “interval” $[s, t]$ with $s, t \in [0, \sigma]$ (but not necessarily $s \leq t$) such that $p_h(s) = a$ and $p_h(t) = b$, and we then set $[a, b] = p_h([s, t])$. Note that $[a, b]$ is typically different from $[b, a]$. Intuitively, $[a, b]$ is the set of all points of \mathcal{T}_h that are visited when going from a to b around the tree in “clockwise order”.

Let us now randomize h . We let $\mathbf{e} = (\mathbf{e}_t)_{0 \leq t \leq 1}$ be a normalized Brownian excursion. Informally, this means that \mathbf{e} is distributed as a linear Brownian motion started at 0 at time 0, conditioned to be back at 0 at time 1 and to stay positive over the interval $(0, 1)$.

Definition 2 *The (standard) Brownian tree is the tree $\mathcal{T}_{\mathbf{e}}$ coded by $\mathbf{e} = (\mathbf{e}_t)_{0 \leq t \leq 1}$.*

Up to an unimportant scaling factor 2 for the distance, the Brownian tree coincides with the Continuum Random Tree introduced by Aldous [4, 5], which has been shown to be the scaling limit of many different classes of discrete random trees, and in particular of the uniformly distributed plane trees considered in Section 2.2.

3.3.2 Brownian labels on the Brownian tree

In a way analogous to the labeled trees of Section 2.2, we now need to assign (real) labels to the points of the Brownian tree $\mathcal{T}_{\mathbf{e}}$. Recall our notation $\rho_{\mathbf{e}} = p_{\mathbf{e}}(0)$ for the root of $\mathcal{T}_{\mathbf{e}}$. Informally, if we condition on the tree $\mathcal{T}_{\mathbf{e}}$, or equivalently on \mathbf{e} , we can introduce a collection $(Z_a)_{a \in \mathcal{T}_{\mathbf{e}}}$, which is a centered Gaussian process vanishing at the root ($Z_{\rho_{\mathbf{e}}} = 0$), whose (conditional) covariance structure is determined by

$$\mathbb{E}[(Z_a - Z_b)^2 \mid \mathbf{e}] = d_{\mathbf{e}}(a, b), \quad a, b \in \mathcal{T}_{\mathbf{e}}.$$

In a mathematically more precise way, we start from a random process $(Y_t)_{t \in [0,1]}$ which, conditionally given \mathbf{e} , is centered Gaussian with (conditional) covariance

$$\mathbb{E}[Y_s Y_t \mid \mathbf{e}] = \min_{s \wedge t \leq r \leq s \vee t} \mathbf{e}_r, \quad s, t \in [0, 1].$$

Then a.s. $d_{\mathbf{e}}(s, t) = 0$ implies $Y_s = Y_t$, and thus, for every $a \in \mathcal{T}_{\mathbf{e}}$, we can define $Z_a = Y_s$ whenever a is the equivalence class of s in the quotient space $\mathcal{T}_{\mathbf{e}} = [0, 1] / \sim_{\mathbf{e}}$. We view Z_a as a label assigned to the point a of $\mathcal{T}_{\mathbf{e}}$. Labels evolve like linear Brownian motion when moving along a line segment of the tree $\mathcal{T}_{\mathbf{e}}$ — this is of course similar to the assignment of labels in discrete labeled trees.

3.3.3 Defining the Brownian sphere

Recall the Brownian tree $\mathcal{T}_{\mathbf{e}}$ and the collection of labels $(Z_a)_{a \in \mathcal{T}_{\mathbf{e}}}$. Also recall that, for every $a, b \in \mathcal{T}_{\mathbf{e}}$, we can define an “interval” $[a, b]$, which we interpret as the set of all points of $\mathcal{T}_{\mathbf{e}}$ that are visited when going from a to b around the tree in “clockwise order”.

We then set, for every $a, b \in \mathcal{T}_{\mathbf{e}}$,

$$D^\circ(a, b) = Z_a + Z_b - 2 \max \left(\min_{c \in [a, b]} Z_c, \min_{c \in [b, a]} Z_c \right). \quad (3)$$

The right-hand side of (3) is a continuous analog of the right-hand side of the bound (2).

We note that $D^\circ(a, b) = 0$ if and only if

$$Z_a = Z_b = \max \left(\min_{c \in [a, b]} Z_c, \min_{c \in [b, a]} Z_c \right), \quad (4)$$

which informally means that a and b have the same label and that we can go from a to b around the tree (clockwise or counterclockwise) visiting only points whose label is at least as large as the label of a and b . We then let $D(a, b)$ be the largest symmetric function of the pair (a, b) that is bounded above by $D^\circ(a, b)$ and satisfies the triangle inequality : For every $a, b \in \mathcal{T}_{\zeta}$,

$$D(a, b) = \inf \left\{ \sum_{i=1}^k D^\circ(a_{i-1}, a_i) \right\}, \quad (5)$$

where the infimum is over all choices of the integer $k \geq 1$ and of the elements a_0, a_1, \dots, a_k of \mathcal{T}_ζ such that $a_0 = a$ and $a_k = b$. Then, D is a pseudo-metric on \mathcal{T}_e , and therefore we can consider the associated equivalence relation : for $a, b \in \mathcal{T}_e$,

$$a \approx b \quad \text{if and only if} \quad D(a, b) = 0.$$

One can also prove [47] that $D(a, b) = 0$ holds if and only if $D^\circ(a, b) = 0$ (the “if” part is trivial), and thus $a \approx b$ is equivalent to the property (4).

Definition 3 *The (standard) Brownian sphere is the quotient space $\mathbf{m}_\infty := \mathcal{T}_e / \approx$ equipped with the distance induced by D , for which we keep the same notation D .*

In contrast with the discrete CVS bijection, we need to make certain identifications in the tree \mathcal{T}_e to get the Brownian sphere. However, in a sense, we do not make many identifications : it is not hard to check that the equivalence class of a typical point of \mathcal{T}_e is a singleton, and moreover equivalence classes can contain at most 3 points. Still these identifications drastically change the topology, since \mathbf{m}_∞ has the topology of the sphere (Theorem 4).

We will write Π for the canonical projection from \mathcal{T}_e onto \mathbf{m}_∞ . The Brownian sphere comes with two distinguished points. The first one is $x_* = \Pi(a_*)$, where a_* is the (unique) element of \mathcal{T}_e with minimal label,

$$Z_{a_*} = \min_{a \in \mathcal{T}_e} Z_a.$$

The second one is $x_0 = \Pi(\rho_e)$, where we recall the notation ρ_e for the root of \mathcal{T}_e .

Let us briefly explain why the CVS bijection makes this definition of the Brownian sphere as the scaling limit of random quadrangulations at least plausible. Since vertices of quadrangulations correspond to vertices of the associated tree in the CVS bijection, and it is well known that the scaling limit of discrete plane trees is the Brownian tree, it is not surprising that points of the Brownian sphere can be represented by points of the Brownian tree. However, we need to make certain identifications in the Brownian tree, which can be explained as follows. A simple argument shows that, in a large quadrangulation, there exist pairs of vertices of the associated tree which are at a macroscopic distance in the tree and still linked by an edge of the quadrangulation. Since distances are rescaled by a factor tending to 0, such a pair of vertices has to be identified in the scaling limit.

We can also provide a heuristic explanation of the formula for the distance D . We first note that

$$D(a, b) \geq |Z_a - Z_b| \tag{6}$$

as an immediate consequence of the similar bound for D° . Then, it follows from our definitions and the choice of a_* , that

$$D(a_*, a) \leq D^\circ(a_*, a) = Z_a - Z_{a_*}.$$

Comparing with (6), we get that, for every $a \in \mathcal{T}_e$

$$D(a_*, a) = Z_a - Z_{a_*}. \tag{7}$$

This is clearly the analog of formula (1) in the discrete setting of the CVS bijection.

Then, let $x, y \in \mathbf{m}_\infty$ and $a, b \in \mathcal{T}_e$ such that $x = \Pi(a)$ and $y = \Pi(b)$. In the same way as we interpreted the right-hand side of (2) as the length of the concatenation of

two geodesic paths to v_* , we can interpret the quantity $D^\circ(a, b)$ as the length of the path from x to y in \mathbf{m}_∞ obtained by concatenating two geodesics, respectively from x to x_* and from y to x_* , up to the time when they coalesce (a complete description of geodesics to x_* and their coalescing properties is given in Section 4 below). This justifies the fact that the Brownian sphere distance is bounded above by $D^\circ(a, b)$, hence by $D(a, b)$ since it has to satisfy the triangle inequality. Of course, proving that the Brownian sphere distance is indeed equal to $D^\infty(a, b)$ requires much more work : Roughly speaking, one needs to verify that a geodesic from x to y can be approximated by the concatenation of paths that are portions of geodesics to x_* (see [64] or [49]).

The volume measure on \mathbf{m}_∞ is the pushforward of the volume measure on \mathcal{T}_e under the canonical projection. One can prove that the two distinguished points x_* and x_0 are independently uniformly distributed according to the volume measure. In a precise form, this means that the law of the two-pointed compact metric space $(\mathbf{m}_\infty, D, x_*, x_0)$ remains the same if one replaces x_* and x_0 by two independent points distributed according to the volume measure.

3.4 k -point functions and Voronoï cells

Following ideas of Gromov, a convenient way to characterize the distribution of the Brownian sphere as a random metric space equipped with a volume measure would be to compute explicitly the k -point functions. Here, for every integer $k \geq 2$, the k -point function is the distribution of the collection

$$(D(y_i, y_j))_{1 \leq i < j \leq k}$$

where the points y_1, \dots, y_k are chosen independently according to the volume measure on \mathbf{m}_∞ .

Let us first consider the case $k = 2$. Then, as was mentioned above, we can take $y_1 = x_*$ and $y_2 = x_0$, and we have

$$D(y_1, y_2) = D(x_*, x_0) = -Z_{a_*}$$

where the last equality follows from (7). So the two-point function is just the distribution of

$$|Z_{a_*}| = |\inf\{Z_a : a \in \mathcal{T}_e\}|.$$

Using connections between Brownian labels on the Brownian tree and certain semilinear partial differential equations, Delmas [34] found the explicit formula : for every $\beta > 0$ and $\lambda > 0$,

$$\int_0^\infty \frac{dr}{r^{3/2}} e^{-\lambda r} \mathbb{P}(D(x_*, x_0) > \beta r^{-1/4}) = \frac{6\sqrt{\pi}\sqrt{\lambda}}{(\sinh((2\lambda)^{1/4}\beta))^2},$$

and this formula characterizes the two-point function.

In a remarkable paper [17], Bouttier and Guitter were able to analyse the three-point function of large random quadrangulations, thus characterizing in the continuous limit the three-point function of the Brownian sphere (see also [37] for the analogous result for general planar maps). Unfortunately, it seems hopeless to extend these calculations to other values of k .

Voronoi cells in the Brownian sphere. We conclude this section with the presentation of an intriguing open problem concerning Voronoi tessellations in the Brownian sphere. Let $k \geq 2$, and suppose again that y_1, \dots, y_k are chosen independently according to the volume measure on \mathbf{m}_∞ . Then, the i -th Voronoi cell is defined by

$$C_i := \{y \in \mathbf{m}_\infty : D(y, y_i) < D(y, y_j) \text{ for every } j \neq i\}.$$

Conjecture (Chapuy [24]). The distribution of the k -tuple

$$(\text{Vol}(C_1), \dots, \text{Vol}(C_k))$$

is uniform over the simplex $\{(r_1, \dots, r_k) \in [0, 1]^k : r_1 + \dots + r_k = 1\}$.

This conjecture was motivated by certain moment calculations [24] which are consistent with the conjectured distribution. It is also worth noting that the analogous result for the Brownian tree has been proved.

At the cost of heavy calculations, the special case $k = 2$ of Chapuy's conjecture has been proved by Guitter [39]. Still it seems that we are very far from the solution for general k . The main reason why a direct approach in the continuous setting is difficult is the fact that our construction of the Brownian sphere gives insight in distances from one typical point (cf. formula (7)) but cannot handle simultaneously distances from several typical points. In relation with this remark, it might be possible to use Miermont's work [63], which gives analogs of the CVS bijection involving distances from several points of the quadrangulation.

3.5 The free Brownian sphere

By construction, the Brownian sphere has a total volume equal to 1. It turns out that it is also useful to consider a variant of the Brownian sphere where the total volume is random. To this end, it suffices to imitate the preceding construction, just replacing the normalized Brownian excursion \mathbf{e} (with duration 1) by an excursion distributed according to the so-called Itô excursion measure (with random duration).

To define the Itô excursion measure, first introduce the Brownian scaling operator. For every $\lambda > 0$ and every continuous function $g : [0, \sigma] \rightarrow \mathbb{R}$, the function $\varphi_\lambda(g) : [0, \lambda\sigma] \rightarrow \mathbb{R}$ is defined by

$$\varphi_\lambda(g)(t) = \sqrt{\lambda} g\left(\frac{t}{\lambda}\right).$$

Also, let \mathbf{E} stand for the set of all excursions, that is, of all continuous functions $e : [0, \sigma] \rightarrow \mathbb{R}_+$ such that $e(0) = e(\sigma) = 0$ and $e(t) > 0$ for every $t \in (0, \sigma)$. Then, the (infinite) Itô measure $\mathbf{n}(de)$ is the measure on \mathbf{E} defined by

$$\int F(e) \mathbf{n}(de) = \int_0^\infty \mathbb{E}[F(\varphi_\ell(\mathbf{e}))] \frac{d\ell}{2\sqrt{2\pi\ell^3}}.$$

We can then construct the free Brownian sphere $(\mathbf{m}_\infty^{\text{free}}, D^{\text{free}})$ from an excursion e distributed according to $\mathbf{n}(de)$, by exactly the same method, starting from the tree \mathcal{T}_e coded by e , assigning Brownian labels to the points of this tree, and defining the analogs of the functions D° and D , and then the quotient space $\mathbf{m}_\infty^{\text{free}} := \mathcal{T}_e / \approx$. It turns out that the free Brownian sphere has nicer spatial Markov properties than

the standard Brownian sphere, but the price to pay is to work under an infinite measure instead of the usual probability space setting.

As a final remark, one can state an analog of Theorem 2 with convergence to the free Brownian sphere. For instance in the case of quadrangulations, the idea is to consider quadrangulations with a random number of faces (instead of a fixed number n in Theorem 2) such that the probability of a given quadrangulation with k faces is proportional to 12^{-k} (here 12 is the critical value for which we can define such a probability measure).

4 Geodesics in the Brownian sphere

As a limit of (rescaled) planar maps equipped with the graph distance, the Brownian sphere is a geodesic space, meaning that any two points are connected by a geodesic path. Geodesics in the Brownian sphere exhibit a number of remarkable properties. We start by describing geodesics to the special point x_* , noting that similar properties will hold if x_* is replaced by a point chosen at random according to the volume measure. Recall our notation Π for the canonical projection from the tree \mathcal{T}_e onto the Brownian sphere \mathbf{m}_∞ .

Let $x \in \mathbf{m}_\infty$, and write $x = \Pi(a)$ for $a \in \mathcal{T}_e$. The basic idea to construct a geodesic from x to x_* is to explore the tree \mathcal{T}_e from the point a in counterclockwise order, and to record, for every $r \in [Z_{a_*}, Z_a]$, the first visited point with label r . The collection of (the images under Π of) these points forms a geodesic from x to x_* — recall from formula (7) that $D(x_*, x) = Z_a - Z_{a_*}$. To make this more precise, let $s \in [0, 1)$ such that $p_e(s) = a$, and recall the process $(Y_t)_{t \in [0, 1]}$ introduced in Section 3.3.2 to define labels on \mathcal{T}_e (in particular $Y_s = Z_a$). For every $r \in [0, Z_a - Z_{a_*}]$, set

$$\psi_s(r) = \begin{cases} \sup\{t \in [0, s] : Y_t = Z_{a_*} + r\} & \text{if } \{t \in [0, s] : Y_t = Z_{a_*} + r\} \neq \emptyset, \\ \sup\{t \in [s, 1] : Y_t = Z_{a_*} + r\} & \text{otherwise.} \end{cases}$$

Then, it is not hard to verify that the curve

$$[0, Z_a - Z_{a_*}] \ni r \mapsto \Gamma_s(r) := \Pi \circ p_e(\psi_s(r))$$

is a geodesic from x_* to x . This is indeed a continuous analog of the construction of geodesics of quadrangulations via the CVS bijection that was outline after Proposition 1. We call such a geodesic a *simple geodesic*. In fact, we obtain in this way all geodesics to the distinguished point x_* .

Theorem 5 [48] *All geodesics that end at x_* are simple geodesics.*

This result was a key ingredient of the proof of the uniqueness of the Brownian sphere in [49]. We observe that, if x and a such that $\Pi(a) = x$ are given, there may be several choices for $s \in [0, 1)$ such that $p_e(s) = a$ (the choice of s is unique only if a is a leaf of \mathcal{T}_e). When a is not a leaf, there are typically two possible choices for s (three when a is a branching point of \mathcal{T}_e) and these choices lead to different simple geodesics. On the other hand, the choice of a such that $\Pi(a) = x$ may also not be unique, but the identifications in the Brownian sphere show that changing a does not lead to a different geodesic.

To summarize the preceding discussion, let the skeleton of \mathcal{T}_e be the set of all points of \mathcal{T}_e that are not leaves. Then the *cut-locus* of the Brownian sphere relative

to the distinguished point x_* (that is, the set of all points that are connected to x_* by at least two different geodesics) is exactly the image under Π of the skeleton of \mathcal{T}_e . There is a striking analogy between this result and classical results in the geometry of Riemannian surfaces, where one proves that the cut-locus is always a tree. We also observe that the dimension of the skeleton of \mathcal{T}_e is 1 (whereas the dimension of \mathcal{T}_e is 2) and one can infer that the dimension of the cut-locus is 2, in contrast with the dimension 4 of the Brownian sphere.

Furthermore, it is easy to verify from the preceding considerations that the cut-locus (relative to x_*) has zero volume, which means that if x is a typical point of \mathbf{m}_∞ , there is a unique geodesic from x to x_* , and more generally there is a unique geodesic between two typical points of \mathbf{m}_∞ (here “typical” means chosen according to the volume measure). The latter result had been obtained earlier by Miermont [63], in fact in a more general setting.

Theorem 5 has another important consequence, namely the so-called “confluence property of geodesics”.

Corollary 6 [48] *Two geodesic paths to x_* coalesce before hitting x_* .*

In other words, if γ and γ' are two (non-trivial) geodesic paths starting from x_* , there exists $\varepsilon > 0$ such that $\gamma(t) = \gamma'(t)$ for every $t \in [0, \varepsilon]$. The proof of the corollary is easy from Theorem 5. With the previous notation, we can assume that $\gamma = \Gamma_s$ and $\gamma' = \Gamma_{s'}$ for some $s, s' \in [0, 1]$. Recall our special convention for the “intervals” $[s, s']$ and $[s', s]$ that was explained in Section 3.3.1 to define lexicographical intervals on \mathcal{T}_e . The construction of simple geodesics then ensures that $\Gamma_s(t) = \Gamma_{s'}(t)$ whenever

$$t \leq \max \left(\min_{u \in [s, s']} Y_u, \min_{u \in [s', s]} Y_u \right) - Z_{a_*}$$

and the max in the right-hand side is greater than Z_{a_*} because one of the two “intervals” $[s, s']$ and $[s', s]$ does not contain the (unique) minimizing time t_* such that $Y_{t_*} = Z_{a_*}$.

Corollary 6 remains valid if x_* is replaced by a typical point of the Brownian sphere. There are however exceptional points where the property of Corollary 6 fails — it obviously fails if x_* is replaced by an interior point of a geodesic. These exceptional points are called geodesic stars.

Definition 4 *Let $n \geq 2$ be an integer. A point x of the Brownian sphere is called a geodesic star if there exist $\varepsilon > 0$ and n geodesic paths $\gamma_1, \dots, \gamma_n$ parametrized by the interval $[0, \varepsilon]$ such that $\gamma_1(0) = \dots = \gamma_n(0) = x$, but the sets $\{\gamma_j(t) : t \in (0, \varepsilon]\}$, $j = 1, \dots, n$, are disjoint.*

Geodesic stars were introduced in [64] where they play an important role in the proof of the uniqueness of the Brownian sphere.

Theorem 7 *Let $n \in \{2, 3, 4\}$. The dimension of the set of all geodesic stars with n arms is equal to $5 - n$.*

The upper bound for the dimension was obtained in [65], and the lower bound was proved in [52]. The paper [65] also proves that geodesic stars with n arms do not exist when $n \geq 6$. This leaves the following open problem.

Open problem. Is the set of geodesic stars with 5 arms nonempty? According to [65], if this set is nonempty, its Hausdorff dimension is 0.

Thanks in particular to the papers [8] and [65], we now know a lot about geodesics in the Brownian sphere (see also Bouttier and Guitter [18, 19] for a discussion of geodesics in large quadrangulations). However, there remain some intriguing problems, and in particular the following one, which was communicated to the author by Omer Angel.

Open problem. Can one find $\varepsilon > 0$ and two geodesic paths γ and γ' parameterized by the interval $[0, \varepsilon]$ such that the following holds : there exists $r \in (0, \varepsilon)$ such that $\gamma(r) = \gamma'(r)$ and the two sets $\{\gamma(t) : t \in [0, \varepsilon]\}$ and $\{\gamma'(t) : t \in [0, \varepsilon]\}$ intersect only at the point $\gamma(r)$?

In other words, can two geodesic paths in the Brownian sphere have a “proper” crossing? If they do, the crossing point will be a geodesic star with 4 arms, but of a very special type so that it is very likely that such points do not exist. A precise argument is however still missing.

5 Planar maps with a boundary and Brownian disks

5.1 Quadrangulations with a boundary

For simplicity, we restrict our attention to quadrangulations in this section, although similar results hold for more general planar maps, in particular for triangulations.

A quadrangulation with a (general) boundary is a rooted planar map Q such that all faces but the root face (lying to the left of the root edge) have degree 4. The root face is also called the outer face and the other faces are the inner faces. The degree of the outer face, which is an even integer, is then called the boundary size or the perimeter of Q .

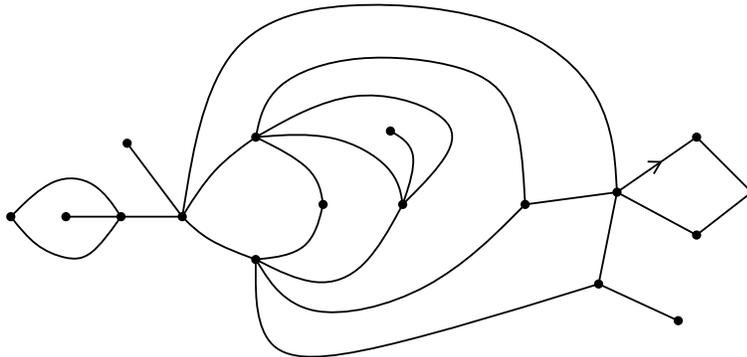


FIGURE 4 – A quadrangulation with a boundary of size 16.

For every integer $k \geq 1$, we denote the set of all quadrangulations with boundary size $2k$ by $\mathbb{Q}^{\partial, k}$. For every integer $n \geq 0$, the subset of $\mathbb{Q}^{\partial, k}$ consisting of those quadrangulations q that have n inner faces is denoted by $\mathbb{Q}_n^{\partial, k}$. Then, for every $k \geq 1$, there is a constant $b_k > 0$ such that

$$\#\mathbb{Q}_n^{\partial, k} \underset{n \rightarrow \infty}{\sim} b_k 12^n n^{-3/2}.$$

See formula (4) in [31].

A random variable $B_{(k)}$ with values in $\mathbb{Q}^{\partial,k}$ is called a Boltzmann quadrangulation with a boundary of size $2k$ if, for every integer $n \geq 0$ and every $Q \in \mathbb{Q}_n^{\partial,k}$,

$$\mathbb{P}(B_{(k)} = Q) = \tilde{b}_k 12^{-n},$$

where $\tilde{b}_k > 0$ is the appropriate normalizing constant so that the sum of the quantities $\tilde{b}_k 12^{-n}$ over all choices of n and Q is equal to 1.

The following theorem is proved in [14].

Theorem 8 *Let $B_{(k)}$ be a Boltzmann quadrangulation with a boundary of size $2k$, for every $k \geq 1$. Let $d_{\text{gr}}^{B_{(k)}}$ denote the graph distance on the vertex set $V(B_{(k)})$. Then*

$$\left(V(B_{(k)}), \sqrt{\frac{3}{2}} k^{-1/2} d_{\text{gr}}^{B_{(k)}} \right) \xrightarrow[k \rightarrow \infty]{(d)} (\mathbb{D}_1, D^\partial)$$

where the convergence holds in distribution in the Gromov-Hausdorff space $(\mathbb{K}, d_{\text{GH}})$. The limit $(\mathbb{D}_1, D^\partial)$ is a random compact metric space called the free Brownian disk with perimeter 1.

For every $r > 0$, the free Brownian disk with perimeter r , denoted by $(\mathbb{D}_r, D^\partial)$ can then be defined by scaling :

$$(\mathbb{D}_r, D^\partial) \stackrel{(d)}{=} (\mathbb{D}_1, \sqrt{r} D^\partial).$$

One proves [11] that the (free) Brownian disk is homeomorphic to the closed unit disk of the plane, which makes it possible to define the boundary $\partial\mathbb{D}_1$.

5.2 The construction of the Brownian disk

In this section, we recall the construction of (free) Brownian disks. Our presentation is slightly different from the one in [14].

The general idea is the same as for the construction of the Brownian sphere, but instead of considering a single Brownian tree $\mathcal{T}_{\mathbf{e}}$, we need to introduce a Poisson collection of such trees, which will be rooted on a “floor” represented by the interval $[0, 1]$ ($[0, r]$ for the Brownian disk with perimeter r). This floor will indeed correspond to the boundary of the Brownian disk.

To make this precise, we consider a Poisson point measure $\mathcal{N} = \sum_{i \in I} \delta_{(t_i, e_i)}$ on $[0, 1] \times \mathbf{E}$ with intensity

$$2 \mathbf{1}_{[0,1]}(t) dt \mathbf{n}(de),$$

where we recall the notation $\mathbf{n}(de)$ for the Itô measure of Brownian excursions. Recall the notation $(\mathcal{T}_{e_i}, d_{e_i})$ for the tree coded by e_i , and also write $\sigma(e_i)$ for the duration of the excursion e_i . We introduce the compact metric space \mathfrak{H} , which is obtained from the disjoint union

$$[0, 1] \cup \left(\bigcup_{i \in I} \mathcal{T}_{e_i} \right) \tag{8}$$

by identifying the root ρ_{e_i} of \mathcal{T}_{e_i} with the point t_i of $[0, 1]$, for every $i \in I$. The metric $d_{\mathfrak{H}}$ on \mathfrak{H} is defined in the obvious manner. In particular, the restriction of $d_{\mathfrak{H}}$

to each tree \mathcal{T}_{e_i} is the metric d_{e_i} , the restriction of $d_{\mathfrak{H}}$ to $[0, 1]$ is the usual metric and, if $u \in \mathcal{T}_{e_i}$ and $v \in \mathcal{T}_{e_j}$, with $j \neq i$,

$$d_{\mathfrak{H}}(u, v) = d_{e_i}(u, \rho_{e_i}) + |\rho_{e_i} - \rho_{e_j}| + d_{e_j}(\rho_{e_j}, v).$$

The volume measure on \mathfrak{H} is just the sum of the volume measures on the trees \mathcal{T}_{e_i} , $i \in I$.

If $\Sigma := \sum_{i \in I} \sigma(e_i)$ is the total mass of the volume measure, we can define a clockwise exploration $(\mathcal{E}_t)_{0 \leq t \leq \Sigma}$ of \mathfrak{H} by concatenating the mappings $p_{e_i} : [0, \sigma(e_i)] \rightarrow \mathcal{T}_{e_i}$ in the order prescribed by the t_i 's.

The clockwise exploration allows us to define ‘‘intervals’’ in \mathfrak{H} . Similarly as in the construction of the Brownian sphere, we make the convention that, if $s, t \in [0, \Sigma]$ and $s > t$, the interval $[s, t]$ is defined by $[s, t] := [s, \Sigma] \cup [0, t]$ (of course, if $s \leq t$, $[s, t]$ is the usual interval). Then, for every $u, v \in \mathfrak{H}$, such that $u \neq v$, there is a smallest interval $[s, t]$, with $s, t \in [0, \Sigma]$, such that $\mathcal{E}_s = u$ and $\mathcal{E}_t = v$, and we define

$$[[u, v]] := \{\mathcal{E}_r : r \in [s, t]\}.$$

Note that we use the notation $[[u, v]]$ rather than $[u, v]$ to avoid confusion with intervals of the real line.

We next assign real labels to the points of \mathfrak{H} . We first assign (independently) Brownian labels $(Z_a^i)_{a \in \mathcal{T}_{e_i}}$ to each tree \mathcal{T}_{e_i} , in the way explained in Section 3.3.2. Then, we let $(\mathbf{e}_t)_{0 \leq t \leq 1}$ be a normalized Brownian excursion, which is independent of \mathcal{N} and the labels $(Z_a^i)_{a \in \mathcal{T}_{e_i}}$. For $t \in [0, 1]$, we set $\Lambda_t := \sqrt{3} \mathbf{e}_t$, and for $u \in \mathcal{T}_{e_i}$, $i \in I$,

$$\Lambda_u := \sqrt{3} \mathbf{e}_{t_i} + Z_u^i.$$

By [11, Lemma 11], $\min\{\Lambda_u : u \in \mathfrak{H}\}$ is attained at a unique point v_* of \mathfrak{H} .

Just as in the construction of the Brownian sphere, labels allow us to define the pseudo-metric D^∂ on \mathfrak{H} . For every $u, v \in \mathfrak{H}$, we first set

$$D^\circ(u, v) := \Lambda_u + \Lambda_v - 2 \max \left(\inf_{w \in [[u, v]]} \Lambda_w, \inf_{w \in [[v, u]]} \Lambda_w \right), \quad (9)$$

and then

$$D^\partial(u, v) := \inf_{u_0=u, u_1, \dots, u_p=v} \sum_{i=1}^p D^\circ(u_{i-1}, u_i), \quad (10)$$

where the infimum is over all choices of the integer $p \geq 1$ and of the finite sequence u_0, u_1, \dots, u_p in \mathfrak{H} such that $u_0 = u$ and $u_p = v$. One immediately verifies that $D^\partial(u, v) \geq |\Lambda_u - \Lambda_v|$ for every $u, v \in \mathfrak{H}$. It easily follows that, for every $u \in \mathfrak{H}$, $D^\partial(u, v_*) = D^\circ(u, v_*) = \Lambda_u - \Lambda_*$.

We abuse notation by writing $\mathfrak{H}/\{D^\partial = 0\}$ for the quotient space of \mathfrak{H} with respect to the equivalence relation defined by setting $u \approx v$ if and only if $D^\partial(u, v) = 0$.

Definition 5 *The free pointed Brownian disk with perimeter 1 is the quotient space $\mathbb{D}_1^\bullet := \mathfrak{H}/\{D^\partial = 0\}$, which is equipped with the distance induced by D^∂ and with a distinguished point x_* which is the equivalence class of v_* .*

The space \mathbb{D}_1^\bullet is homeomorphic to the closed unit disk, and the boundary $\partial \mathbb{D}_1^\bullet$ is identified with the image of $[0, 1]$ under the canonical projection. The volume

measure Vol on \mathbb{D}_1^\bullet is obtained as the pushforward of the volume measure on \mathfrak{H} under the canonical projection. From the preceding construction, one verifies that the variable $\text{Vol}(\mathbb{D}_1^\bullet)$ has a density equal to

$$\frac{1}{\sqrt{2\pi t^3}} \exp\left(-\frac{1}{2t}\right).$$

One can also define a boundary measure, which is in a sense uniformly distributed over $\partial\mathbb{D}_1^\bullet$. This boundary measure μ_∂ may be defined by the approximation

$$\langle \mu_\partial, \varphi \rangle = \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon^2} \int \varphi(x) \mathbf{1}_{\{D^\partial(x, \partial\mathbb{D}_1^\bullet) < \varepsilon\}} \text{Vol}(dx), \quad (11)$$

for any continuous function φ on \mathbb{D}_1^\bullet .

The free (unpointed) Brownian disk with perimeter 1 is obtained from the free pointed Brownian disk by “forgetting” the distinguished point and putting a weight proportional to the inverse of the volume : If F is a function defined on the Gromov-Hausdorff space, the formula

$$\mathbb{E}[F(\mathbb{D}_1)] = \mathbb{E}\left[\frac{1}{\text{Vol}(\mathbb{D}_1^\bullet)} F(\mathbb{D}_1^\bullet)\right]$$

defines the distribution of the space \mathbb{D}_1 in Theorem 8. Of course the approximation (11) still makes sense to define the boundary measure μ_∂ on $\partial\mathbb{D}_1$.

As a final remark, one can also define the Brownian disk with fixed perimeter (say, equal to 1) and volume v by conditioning \mathbb{D}_1 on $\text{Vol}(\mathbb{D}_1) = v$.

6 Non-compact models

In this section, we briefly describe the two basic non-compact models called the Brownian plane and the Brownian half-plane.

6.1 The Brownian plane

Recall that, in Theorem 2, the graph distances were rescaled by the factor $n^{-1/4}$, which was the proper factor to get a (non-trivial) compact limit (as we explained in the case of quadrangulations, the typical diameter of a planar map with n faces is of order $n^{1/4}$). The Brownian plane will arise in a similar scaling limit for random planar maps if we rescale the graph distance by a factor tending to 0 at a slower rate than $n^{1/4}$. In order to explain this convergence, we need to introduce the local Gromov-Hausdorff convergence. Let (E_n, d_n, x_n) be a sequence of pointed metric spaces, where the word pointed means that there is a distinguished point $x_n \in E_n$, for every n . Assume that, for every n , E_n is a geodesic space (the distance between two points is the length of a shortest path between these points) and closed balls of E_n are compact. Then, we say that (E_n, d_n, x_n) converges to a limiting pointed metric space (E, d, x) for the local Gromov-Hausdorff topology if, for every $r > 0$, the closed ball of radius r centered at x_n in E_n converges to the same ball centered at x in E , for the Gromov-Hausdorff distance — implicitly we assume that closed balls in E are also compact.

Recall our notation \mathbb{M}_n^4 for the set of all rooted planar quadrangulations with n faces. For $Q \in \mathbb{M}_n^4$, we view the metric space $(V(Q), d_{\text{gr}}^Q)$ as a pointed metric space, whose distinguished point is the root vertex of Q .

Theorem 9 [28] *Suppose that, for every $n \geq 1$, Q_n is uniformly distributed over \mathbb{M}_n^4 , and let (a_n) be a sequence of positive reals such that $a_n \rightarrow 0$ and $n^{1/4}a_n \rightarrow \infty$ as $n \rightarrow \infty$. Then*

$$(V(Q_n), a_n d_{\text{gr}}^{Q_n}) \xrightarrow[n \rightarrow \infty]{(d)} (\mathcal{P}, D^\infty),$$

where the convergence holds in the local Gromov-Hausdorff sense. The limiting random pointed metric space (\mathcal{P}, D^∞) is called the Brownian plane.

One expects that a similar result holds for more general random planar maps.

The Brownian plane is homeomorphic to the usual plane \mathbb{R}^2 and has Hausdorff dimension 4 like the Brownian sphere. The Brownian plane shares many properties of the Brownian sphere, and can be viewed as a tangent cone in distribution to the Brownian sphere (\mathbf{m}_∞, D) , meaning that

$$(\mathbf{m}_\infty, \lambda D) \xrightarrow[\lambda \rightarrow \infty]{(d)} (\mathcal{P}, D^\infty)$$

again in the local Gromov-Hausdorff sense. In the latter convergence, we view the Brownian sphere \mathbf{m}_∞ as pointed at the distinguished point x_* (but x_* could be replaced by a point uniformly distributed on \mathbf{m}_∞). In fact, a much stronger result holds : For every $\delta > 0$, one can find $\varepsilon > 0$ and couple the Brownian sphere (\mathbf{m}_∞, D) with the Brownian plane (\mathcal{P}, D^∞) so that, with probability at least $1 - \delta$, the balls of radius ε centered at the distinguished point in \mathbf{m}_∞ and in \mathcal{P} are isometric. In other words, one can construct simultaneously \mathbf{m}_∞ and \mathcal{P} so that, with high probability, the small balls centered at the distinguished point are the same in \mathcal{P} and in \mathbf{m}_∞ . As a final remark, the Brownian plane is scale invariant in distribution, meaning that, for every $\lambda > 0$, $(\mathcal{P}, \lambda D^\infty)$ has the same distribution as (\mathcal{P}, D^∞) .

6.2 The Brownian plane as the scaling limit of the UIPQ

Both in Theorem 2 and in Theorem 9, the graph distance is rescaled by a factor tending to 0. If we do not rescale the distance, we can still obtain a limiting object, in the sense of local convergence of maps. To make this precise, consider a sequence \mathbf{m}_n of random rooted planar maps. For every n and for every integer $k \geq 0$, write $B_k(\mathbf{m}_n)$ for the (random) rooted planar map obtained by keeping only those faces of \mathbf{m}_n that contain at least one vertex at distance less than or equal to k from the root vertex of \mathbf{m}_n . We say that \mathbf{m}_n converges locally to a (possibly infinite but locally finite) rooted planar map \mathbf{m} if, for every $k \geq 0$, for every finite rooted planar map m_0 ,

$$\mathbb{P}(B_k(\mathbf{m}_n) = m_0) \xrightarrow[n \rightarrow \infty]{} \mathbb{P}(B_k(\mathbf{m}) = m_0).$$

For every (even) integer n , let T_n denote a uniformly distributed rooted planar triangulation with n faces. In a pioneering paper, Angel and Schramm [9] proved that T_n converges locally to a random infinite planar map T_∞ called the UIPQ (for Uniform Infinite Planar Triangulation). The proofs of [9] rely much on exact enumeration formulas for triangulations. Using similar formulas, Krikun [45] proved the analogous result for quadrangulations : If M_n denotes a uniformly distributed rooted planar quadrangulation with n faces, the sequence M_n converges locally to the random infinite planar map Q_∞ called the UIPQ. Both the UIPQ T_∞ and the UIPQ Q_∞ have been studied extensively, see in particular [26, 30, 40, 60].

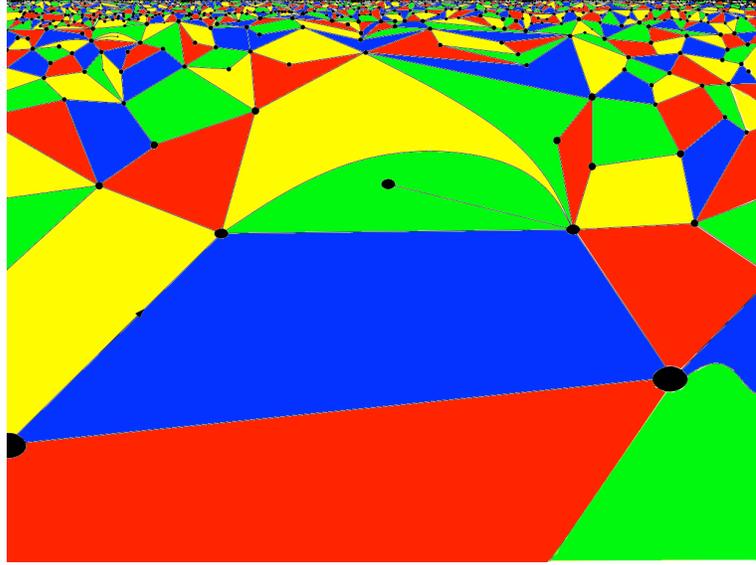


FIGURE 5 – An artistic view of the UIPQ (artist : N. Curien)

The Brownian plane appears as the scaling limit of both the UIPQ and the UIPT. Write $V(T_\infty)$ and $V(Q_\infty)$ for the respective vertex sets of T_∞ and Q_∞ .

Theorem 10 *We have*

$$(V(T_\infty), \lambda d_{\text{gr}}^{T_\infty}) \xrightarrow{\lambda \rightarrow 0} (\mathcal{P}, D^\infty)$$

and

$$(V(Q_\infty), \lambda d_{\text{gr}}^{Q_\infty}) \xrightarrow{\lambda \rightarrow 0} (\mathcal{P}, D^\infty)$$

where in both cases convergence holds in the local Gromov-Hausdorff sense.

We refer to [28] for the case of quadrangulations and to [20] for the case of triangulations.

6.3 The Brownian half-plane

The Brownian half-plane is another important non-compact model, which was introduced via two different constructions by Caraceni and Curien [23] and Gwynne and Miller [41] (the equivalence between these two constructions was later obtained in [51]). Roughly speaking, the Brownian half-plane appears as the scaling limit of random planar maps with a boundary of size $2k$, provided that the scaling factor tends to 0 but is large in comparison with $1/\sqrt{k}$ (recall Theorem 8 where the scaling factor was $\sqrt{3/2k}$).

To state this in a more precise form, let $B_{(k)}$ be a Boltzmann quadrangulation with a boundary of size $2k$ (see Section 5.1). We view $(V(B_{(k)}), d_{\text{gr}}^{B_{(k)}})$ as a pointed metric space, where the distinguished point is the root vertex of $B_{(k)}$ (note that this root vertex lies on the boundary of $B_{(k)}$).

Theorem 11 *Let (a_k) be a sequence such that $a_k \rightarrow 0$ but $a_k \sqrt{k} \rightarrow \infty$ as $k \rightarrow \infty$. We have*

$$(V(B_{(k)}), a_k d_{\text{gr}}^{B_{(k)}}) \xrightarrow[k \rightarrow \infty]{(d)} (\mathcal{H}, D^{\mathcal{H}}),$$

in the local Gromov-Hausdorff sense. The limiting space $(\mathcal{H}, D^{\mathcal{H}})$ is a random non-compact pointed metric space called the Brownian half-plane.

This theorem is proved in [10], which contains a thorough discussion of all possible scaling limits of quadrangulations with a boundary. The Brownian half-plane is homeomorphic to the usual half-plane, and so one can define its boundary $\partial\mathcal{H}$, which is homeomorphic to the line, and the distinguished point of \mathcal{H} lies on $\partial\mathcal{H}$.

In the same way as the Brownian plane appears as the scaling limit of both the UIPT and the UIPQ, the Brownian half-plane was proved to be the scaling limit of the infinite planar map called the uniform infinite half-planar triangulation or UIHPQ, see Gwynne and Miller [41].

7 Spatial Markov property in Brownian surfaces

The Brownian surfaces that we have introduced satisfy certain remarkable spatial Markov properties which we now describe. Roughly speaking, for certain special subsets H of the Brownian sphere (or the Brownian disk, or the Brownian half-plane), the complement of H viewed as a metric space for an intrinsic metric is independent of H (also viewed as a metric space for its intrinsic metric) conditionally on the “boundary size” of H .

Consider first the free Brownian sphere $(\mathbf{m}_{\infty}^{\text{free}}, D^{\text{free}})$ of Section 3.5. This metric space comes with two distinguished points that we denoted by x_* and x_0 , which should be viewed as uniformly distributed. Let us fix $r > 0$, and consider the closed ball of radius r centered at x_* in $\mathbf{m}_{\infty}^{\text{free}}$, which we denote by $B_r(x_*)$. We will condition on the event that $x_0 \notin B_r(x_*)$, or equivalently $D^{\text{free}}(x_*, x_0) > r$. Although the free Brownian sphere is defined under an infinite measure, the latter event has finite measure, and so the conditioning takes us back to a probability measure. The complement of $B_r(x_*)$ has infinitely many connected components, but we define the “hull” $B_r^{\bullet}(x_*)$ by saying that $\mathbf{m}_{\infty}^{\text{free}} \setminus B_r^{\bullet}(x_*)$ is the connected component that contains x_0 (informally, we obtain the hull by adding to the ball of radius r all connected components of its complement but the one containing x_0). One can then define the boundary size of $B_r^{\bullet}(x_*)$ by the formula

$$\mathcal{Z}_r := \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon^2} \text{Vol}(\{x \in \mathbf{m}_{\infty}^{\text{free}} \setminus B_r^{\bullet}(x_*) : D^{\text{free}}(x, B_r^{\bullet}(x_*)) < \varepsilon\}), \quad (12)$$

which is analogous to (11). Write $B_r^{\circ}(x_*)$ for the interior of the hull $B_r^{\bullet}(x_*)$. We then consider the intrinsic distance on $B_r^{\circ}(x_*)$, which is defined by declaring that the distance between two points is the infimum of the lengths of paths staying in $B_r^{\circ}(x_*)$ that connect these two points. It turns out that this intrinsic distance can be extended continuously to $B_r^{\bullet}(x_*)$. Similarly, we can define an intrinsic distance on the open set $\mathbf{m}_{\infty}^{\text{free}} \setminus B_r^{\bullet}(x_*)$ and verify that it can be extended continuously to its closure $\mathbf{m}_{\infty}^{\text{free}} \setminus B_r^{\circ}(x_*)$.

Theorem 12 [52] *Conditionally on the boundary size \mathcal{Z}_r , the hull $B_r^{\bullet}(x_*)$ and the hull complement $\mathbf{m}_{\infty}^{\text{free}} \setminus B_r^{\circ}(x_*)$, both equipped with their respective (extended) intrinsic distances, are independent random metric spaces. Moreover, the hull complement $\mathbf{m}_{\infty}^{\text{free}} \setminus B_r^{\circ}(x_*)$ is a free Brownian disk with perimeter \mathcal{Z}_r .*

A version of this theorem for the standard Brownian sphere can be found in [50], where it is proved that the connected components of the complement of a closed ball centered at x_* in \mathbf{m}_∞ are independent Brownian disks, conditionally on their boundary sizes and volumes (in contrast with Theorem 12, it is here necessary to condition also on the volumes, and as a result one does not get free Brownian disks but Brownian disks with given perimeter and volume).

Let us turn to an analog of Theorem 12 for the free Brownian disk $(\mathbb{D}_1, D^\partial)$. Here, we consider two points y_0 and y_1 which are independently distributed according to the boundary measure μ_∂ . Fix $r > 0$ and argue conditionally on the event $D^\partial(y_0, y_1) > r$. We can then define the hull $B_r^\bullet(y_0)$ by saying that $\mathbb{D}_1 \setminus B_r^\bullet(y_0)$ is the connected component containing y_1 of the complement of the ball of radius r centered at y_0 . To simplify notation, also write $\widehat{B}_r^\bullet(y_0)$ for the closure of $\mathbb{D}_1 \setminus B_r^\bullet(y_0)$. The boundary size \mathcal{Y}_r is then defined by the analog of formula (12) above, replacing x_* by y_0 , $\mathbf{m}_\infty^{\text{free}} \setminus B_r^\bullet(x_*)$ by $\widehat{B}_r^\bullet(y_0)$, and D^{free} by D^∂ . Furthermore, we also define \mathcal{X}_r as the size (measured with respect to the boundary measure μ_∂) of $\partial\mathbb{D}_1 \cap \widehat{B}_r^\bullet(y_0)$. See Fig. 6.

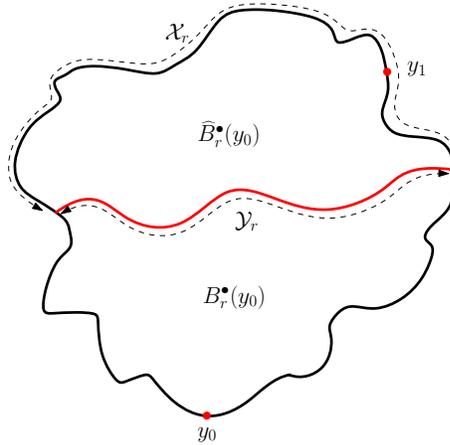


FIGURE 6 – Spatial Markov property in a Brownian disk.

Theorem 13 [56] *The random metric spaces $B_r^\bullet(y_0)$ and $\widehat{B}_r^\bullet(y_0)$ both equipped with their (extended) intrinsic measures are independent conditionally on the pair $(\mathcal{X}_r, \mathcal{Y}_r)$. Furthermore, $\widehat{B}_r^\bullet(y_0)$ is a free Brownian disk with perimeter $\mathcal{X}_r + \mathcal{Y}_r$.*

Instead of hulls centered at a boundary point of the Brownian disk, we could also consider hulls centered at an interior point. To this end, it is convenient to deal with the free pointed Brownian disk \mathbb{D}_1^\bullet constructed in Section 5.2, which has a distinguished (interior) point x_* . For every $0 < r < D^\partial(x_*, \partial\mathbb{D}_1^\bullet)$, we define the hull $B_r^\bullet(x_*)$ by declaring that $\mathbb{D}_1^\bullet \setminus B_r^\bullet(x_*)$ is the connected component containing $\partial\mathbb{D}_1^\bullet$ of the complement of the ball of radius r centered at x_* . For a given $\alpha > 0$, let r_α be the smallest r such that the boundary size of $B_r^\bullet(x_*)$ (defined via an approximation similar to (12)) is equal to α — we need here to condition on the event of positive probability where there exists such a value of r . Then [61] $\mathbb{D}_1^\bullet \setminus B_{r_\alpha}^\bullet(x_*)$ (equipped with its intrinsic metric) is a free Brownian annulus with boundary sizes 1 and α , which moreover is independent of the hull $B_{r_\alpha}^\bullet(x_*)$. The Brownian annulus appears

as the scaling limit of planar quadrangulations with two boundaries : this is a very special case of the convergence results proved in [15] and briefly discussed in Section 8 below.

We finally give a version of the spatial Markov property in the Brownian half-plane. In a sense, this result is nicer as it does not involve any conditioning (compare with Theorems 12 and 13. Recall that the Brownian half-plane $(\mathcal{H}, D^{\mathcal{H}})$ is homeomorphic to the usual half-plane, and that it has a distinguished point, which lies on the boundary $\partial\mathcal{H}$ and which we denote here by y_* . In a sense that can be made precise, the point y_* is a “typical point” of the boundary $\partial\mathcal{H}$.

For every $r > 0$, we define the hull $B_r^{\bullet, \mathcal{H}}$ by saying that $\mathcal{H} \setminus B_r^{\bullet, \mathcal{H}}$ is the (unique) unbounded connected component of the complement of the closed ball of radius r centered at y_* (informally, the hull $B_r^{\bullet, \mathcal{H}}$ is obtained by filling in the bounded holes in the ball of radius r centered at y_*).

We then let \mathcal{H}_r be the closure of $\mathcal{H} \setminus B_r^{\bullet, \mathcal{H}}$. Again, the intrinsic metric on $\mathcal{H} \setminus B_r^{\bullet, \mathcal{H}}$ has a continuous extension to \mathcal{H}_r . The following theorem is proved in a joint work in preparation with Armand Riera.

Theorem 14 *The space \mathcal{H}_r equipped with the extended intrinsic metric is again a Brownian half-plane, and is independent of the hull $B_r^{\bullet, \mathcal{H}}$ also equipped with its intrinsic metric.*

The formulation of Theorem 14 is slightly imprecise as we should specify the distinguished point of \mathcal{H}_r (recall that the Brownian half-plane is a random *pointed* metric space). This distinguished point may be chosen on the boundary of the hull $B_r^{\bullet, \mathcal{H}}$ in a way that is measurable with respect to the hull, but we omit the details.

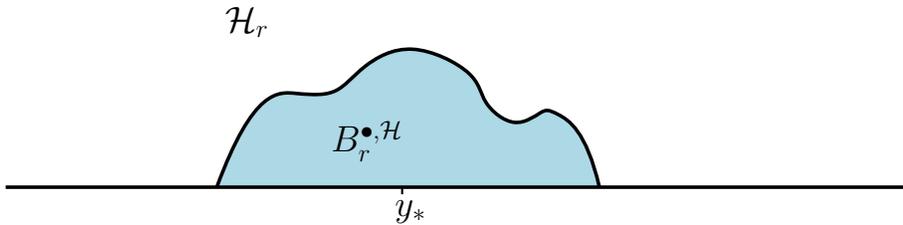


FIGURE 7 – Illustration of Theorem 14 : The complement of the hull $B_r^{\bullet, \mathcal{H}}$ is again a Brownian half-plane

Roughly speaking, Theorem 14 states that, if we remove a hull centered at a typical point of the boundary of the Brownian half-plane \mathcal{H} , what remains is again a Brownian half-plane (provided of course that we use the intrinsic metric). We can iterate this operation and find many subsets of \mathcal{H} such that the space obtained after their removal has the same distribution as \mathcal{H} . An intriguing question is whether the distribution of the Brownian half-plane could be characterized via such spatial Markov properties (and some other properties to specify).

8 Brownian surfaces in higher genus

In this section, we briefly present recent results of Bettinelli and Miermont [15] concerning the construction of Brownian surfaces in arbitrary genus, with a finite number of “boundaries”. This construction relies on an approximation by random

quadrangulations, but it is expected that quadrangulations can be replaced by more general planar maps.

Let $g \in \mathbb{Z}_+$. A map in genus g is a proper cellular embedding of a finite multi-graph in a compact orientable surface of genus g . Here the word cellular means that the connected components of the complement of edges (the faces of the map) are homeomorphic to the open unit disk of the plane. As previously, maps in genus g are considered up to orientation-preserving homeomorphisms of the surface in which they are embedded. Rooted maps are defined as in the planar case. We say that the map is bipartite if the vertex set can be partitioned in two subsets such that no edge links two vertices of the same subset.

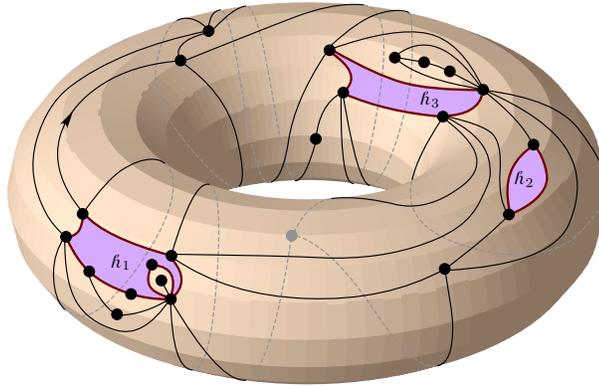


FIGURE 8 – A quadrangulation with three holes marked h_1, h_2, h_3 in genus 1 (figure from [15]).

We consider special cases of maps which we call quadrangulations with holes. For every integer $k \geq 0$, a quadrangulation with k holes (in genus g) is a bipartite rooted map in genus g having k distinguished faces f_1, \dots, f_k (called the holes) with arbitrary (even) degree, and such that all the other faces, which are called internal faces, have degree 4. The degree of each hole is called the perimeter of this hole.

If $\ell = (\ell^1, \dots, \ell^k)$ is a k -tuple of positive integers, and $n \geq 1$, we write $\mathbf{Q}_{n,\ell}^{[g]}$ for the set of all rooted quadrangulations in genus g with n internal faces and k holes of respective perimeters $2\ell^1, \dots, 2\ell^k$. Then, if $Q \in \mathbf{Q}_{n,\ell}^{[g]}$, we can equip the vertex set $V(Q)$ with the graph distance d_{gr}^Q .

Theorem 15 [15] *Let $g, k \geq 0$, and let $L = (L^1, \dots, L^k)$ be a k -tuple of positive real numbers. For every integer $n \geq 1$, let $\ell_n = (\ell_n^1, \dots, \ell_n^k) \in \mathbb{N}^k$, and assume that $\ell_n^i / \sqrt{2n} \rightarrow L^i$ as $n \rightarrow \infty$, for every $i \in \{1, \dots, k\}$. Finally, let Q_n be uniformly distributed over $\mathbf{Q}_{n,\ell_n}^{[g]}$. Then,*

$$\left(V(Q_n), \left(\frac{9}{8n} \right)^{1/4} d_{\text{gr}}^{Q_n} \right) \xrightarrow[n \rightarrow \infty]{(d)} (S_L^{[g]}, D_L^{[g]}),$$

where the limit is a random compact metric space, and the convergence holds in distribution in the Gromov-Hausdorff sense.

Bettinelli and Miermont [15, Theorem 1] give in fact a much more precise result, involving also the convergence of the uniform measure on $V(Q_n)$ and of the uniform measure on the boundary of each hole — the statement of such a result requires a

generalization of the Gromov-Hausdorff convergence which we chose not to present here.

The random metric space $(S_L^{[g]}, D_L^{[g]})$ may be interpreted as the (standard) Brownian surface in genus g with k holes of respective perimeters L_1, \dots, L_k . The proof of (a strong form of) Theorem 15 relies on a decomposition into elementary pieces called slices and quadrilaterals (with geodesic sides), which holds both for the discrete quadrangulations and for the limiting continuous objects. A key point is then to obtain the convergence of the (properly rescaled) discrete slices and quadrilaterals towards their continuous analogs. To get this convergence, it turns out to be convenient to view the discrete quadrilaterals, resp. the discrete slices, as subsets of the UIPQ, resp. of the UIHPQ, and similarly to view the continuous analogs as subsets of the Brownian plane and the Brownian half-plane. Although Theorem 15 deals with compact Brownian surfaces, the proof thus makes a heavy use of the non-compact models discussed in Section 6.

As a final remark, the standard Brownian surface $(S_L^{[g]}, D_L^{[g]})$ is normalized in the sense that its volume is equal to 1 (we start from quadrangulations with a fixed number of faces). Bettinelli and Miermont [15, Theorem 5] also discuss *free* Brownian surfaces in genus g that are scaling limits of quadrangulations distributed according to Boltzmann weights (in the same way as the free Brownian disk is obtained in Theorem 8). Interestingly, these free Brownian surfaces are defined under probability measures only in two special cases in genus 0, namely the Brownian disk studied in Section 5 and the Brownian annulus. Note in particular that the free Brownian sphere of Section 3.5 is defined under an infinite measure.

Acknowledgements. I thank Nicolas Curien for providing Figures 3 and 5, and Jérémie Bettinelli and Grégory Miermont for allowing me to use Figure 8 from [15].

Références

- [1] C. ABRAHAM, Rescaled bipartite planar maps converge to the Brownian map. *Ann. Inst. Henri Poincaré Probab. Stat.* 52, 575–595 (2016)
- [2] L. ADDARIO-BERRY, M. ALBENQUE, The scaling limit of random simple triangulations and random simple quadrangulations. *Ann. Probab.* 45, 2767–2825 (2017)
- [3] L. ADDARIO-BERRY, M. ALBENQUE, Convergence of odd-angulations via symmetrization of labeled trees. *Ann. Henri Lebesgue* 4, 653 – 683 (2021)
- [4] D. ALDOUS, The continuum random tree I. *Ann. Probab.* 19, 1–28 (1991)
- [5] D. ALDOUS, The continuum random tree III. *Ann. Probab.* 21, 248–289 (1993)
- [6] M. ANG, G. RÉMY, X. SUN, The moduli of annuli in random conformal geometry. Preprint, [arXiv:2203.12398](https://arxiv.org/abs/2203.12398)
- [7] O. ANGEL, Growth and percolation on the uniform infinite planar triangulation. *Geomet. Funct. Anal.* 3, 935–974 (2003)
- [8] O. ANGEL, B. KOLESNIK, G. MIERMONT, Stability of geodesics in the Brownian map. *Ann. Probab.* 45, 3451–3479 (2017)
- [9] O. ANGEL, O. SCHRAMM, Uniform infinite planar triangulations. *Comm. Math. Phys.* 241, 191–213 (2003)

- [10] E. BAUR, G. MIERMONT, G. RAY, Classification of scaling limits of quadrangulations with a boundary. *Ann. Probab.* 47, 3397–3477 (2019)
- [11] J. BETTINELLI, Scaling limit of random planar quadrangulations with a boundary. *Ann. Inst. Henri Poincaré Probab. Stat.* 51, 432–477 (2015)
- [12] J. BETTINELLI, Geodesics in Brownian surfaces (Brownian maps). *Ann. Inst. Henri Poincaré Probab. Stat.* 52, 432–477 (2016)
- [13] J. BETTINELLI, E. JACOB, G. MIERMONT, The scaling limit of uniform random plane maps, via the Ambjørn-Budd bijection. *Electron. J. Probab.* 19, paper no. 74, 16pp. (2014)
- [14] J. BETTINELLI, G. MIERMONT, Compact Brownian surfaces I. Brownian disks. *Probab. Theory Related Fields* 167, 555–614 (2017)
- [15] J. BETTINELLI, G. MIERMONT, Compact Brownian surfaces II. Orientable surfaces. Preprint, [arXiv:2212.12511](https://arxiv.org/abs/2212.12511)
- [16] J. BOUTTIER, P. DI FRANCESCO, E. GUITTER, Planar maps as labeled mobiles. *Electronic J. Combinatorics* 11 #R69 (2004)
- [17] J. BOUTTIER, E. GUITTER, The three-point function of planar quadrangulations. *J. Stat. Mech. Theory Exp.* P07020, 39 pp. (2008)
- [18] J. BOUTTIER, E. GUITTER, Statistics in geodesics in large quadrangulations. *Phys. A* 41, no. 14, 145001, 30 pp. (2008)
- [19] J. BOUTTIER, E. GUITTER, Confluence of geodesic paths and separating loops in large planar quadrangulations. *J. Stat. Mech. Theory Exp.* no. 3, P03001, 44 pp. (2009)
- [20] T. BUDZINSKI, The hyperbolic Brownian plane. *Probab. Theory Related Fields* 171, 503–541 (2018)
- [21] D. BURAGO, Y. BURAGO, S. IVANOV, *A Course in Metric Geometry*. Graduate Studies in Mathematics, vol. 33. Amer. Math. Soc., Boston, 2001.
- [22] P. CHASSAING, B. DURHUUS Local limit of labeled trees and expected volume growth in a random quadrangulation. *Ann. Probab.* 34, 879–917 (2006)
- [23] A. CARACENI, N. CURIEN, Geometry of the uniform infinite half-planar quadrangulation, *Random Structures Algorithms* 52, 454–494 (2018)
- [24] G. CHAPUY, On tessellations of random maps and the t_g -recurrence. *Probab. Theory Related Fields* 174, 477–500 (2019)
- [25] P. CHASSAING, G. SCHAEFFER, Random planar lattices and integrated super-Brownian excursion. *Probab. Theory Related Fields* 128, 161–212 (2004)
- [26] P. CHASSAING, B. DURHUUS, Local limits of labeled trees and expected volume growth in a random quadrangulation. *Ann. Probab.* 34, 879–917 (2006)
- [27] R. CORI, B. VAUQUELIN, Planar maps are well labeled trees. *Canad. J. Math.* 33, 1023–1042 (1981)
- [28] N. CURIEN, J.-F. LE GALL, The Brownian plane. *J. Theoret. Probab.* 27, 1240–1291 (2014)
- [29] N. CURIEN, J.-F. LE GALL, First-passage percolation and local modifications of distances in random triangulations. *Ann. Sci. Éc. Norm. Supér.* 52, 631–701 (2019)

- [30] N. CURIEN, L. MÉNARD, G. MIERMONT, A view from infinity of the uniform infinite quadrangulation. *ALEA Lat. Am. J. Probab. Math. Stat.* 10, 45–88 (2013)
- [31] N. CURIEN, G. MIERMONT, Uniform infinite planar quadrangulation with a boundary. *Random Structures Algorithms* 47, 30–58 (2015)
- [32] F. DAVID, Planar diagrams, two-dimensional quantum gravity, and surface models. *Nuclear Phys. B* 257, 45–58 (1985)
- [33] F. DAVID, A. KUPIAINEN, R. RHODES, V. VARGAS, Liouville quantum gravity on the Riemann sphere. *Comm. Math. Phys.* 342, 869–907 (2016)
- [34] J.-F. DELMAS, Computation of moments for the length of the one-dimensional ISE support. *Elect. Journ. of Probab.* 8 (17), pp. 1-15 (2003)
- [35] J. DING, J. DUBÉDAT, A. DUNLAP, H. FALCONET, Tightness of Liouville first passage percolation for $\gamma \in (0, 2)$. *Publ. Math. Inst. Hautes Études Sci.* 132, 353–403 (2020)
- [36] T. DUQUESNE, J.-F. LE GALL, Probabilistic and fractal aspects of Lévy trees. *Probab. Theory Related Fields* 131, 553–603 (2005)
- [37] E. FUSY, E. GUITTER The three-point function of general planar maps. *J. Stat. Mech. Theory Exp.* no.9, p09012, 39 pp. (2014)
- [38] C. GUILLARMOU, R. RHODES, V. VARGAS, Polyakov’s formulation of $2d$ bosonic string theory. *Publ. Math. Inst. Hautes Études Sci.* 130, 11–185 (2019)
- [39] E. GUITTER, On a conjecture by Chapuy about Voronoï cells in large maps. *J. Stat. Mech. Theory Exp.* no.10, 103401, 33 pp. (2017)
- [40] O. GUREL-GUREVICH, A. NACHMIAS, Recurrence of planar graph limits. *Ann. Math.* 177, 761–781 (2013)
- [41] E. GWYNNE, J. MILLER, Scaling limit of the uniform infinite half-plane quadrangulation in the Gromov-Hausdorff-Prokhorov-uniform topology. *Electron. J. Probab.* 22, paper no 84, 47pp. (2017)
- [42] E. GWYNNE, J. MILLER, Convergence of the free Boltzmann quadrangulation with simple boundary to the Brownian disk. *Ann. Inst. Henri Poincaré Probab. Stat.* 55, 551–589 (2019)
- [43] E. GWYNNE, J. MILLER, Existence and uniqueness of the Liouville quantum gravity metric for $\gamma \in (0, 2)$. *Invent. Math.* 223, 213–333 (2021)
- [44] V.G. KNIZHNIK, A.M. POLYAKOV, A.B. ZAMOLODCHIKOV, Fractal structure of $2D$ -quantum gravity. *Modern Phys. Lett. A* 3, 819–826 (1988)
- [45] M. KRIKUN, Local structure of random quadrangulations. Preprint, [arXiv:math/0512304](https://arxiv.org/abs/math/0512304)
- [46] A. KUPIAINEN, R. RHODES, V. VARGAS, Integrability of Liouville theory, proof of the DOZZ formula. *Ann. Math.* (2) 191, 81–166 (2020)
- [47] J.-F. LE GALL, The topological structure of scaling limits of large planar maps. *Invent. Math.* 169, 621–670 (2007)
- [48] J.-F. LE GALL, Geodesics in large planar maps and in the Brownian map. *Acta Math.* 205, 287–360 (2010)

- [49] J.-F. LE GALL, Uniqueness and universality of the Brownian map. *Ann. Probab.* 41, 2880–2960 (2013)
- [50] J.-F. LE GALL, Brownian disks and the Brownian snake. *Ann. Inst. Henri Poincaré Probab. Stat.* 55, 237–313 (2019)
- [51] J.-F. LE GALL, The Brownian disk viewed from a boundary point. *Ann. Inst. Henri Poincaré Probab. Stat.* 58, 1091–1119 (2022)
- [52] J.-F. LE GALL, Geodesic stars in random geometry, *Ann. Probab.* 50, 1013–1058 (2022)
- [53] J.-F. LE GALL, T. LEHÉRICY, Separating cycles and isoperimetric inequalities in the uniform infinite planar quadrangulation. *Ann. Probab.* 47, 1498–1540 (2019)
- [54] J.-F. LE GALL, G. MIERMONT, Scaling limits of random planar maps with large faces. *Ann. Probab.* 39, 1–69 (2011)
- [55] J.-F. LE GALL, F. PAULIN, Scaling limits of bipartite planar maps are homeomorphic to the 2-sphere. *Geomet. Funct. Anal.* 18, 893–918 (2008)
- [56] J.-F. LE GALL, A. RIERA, Spatial Markov property in Brownian disks. Preprint, [arXiv:2302.01138](https://arxiv.org/abs/2302.01138)
- [57] J.-F. MARCKERT, A. MOKKADEM, Limit of normalized quadrangulations. The Brownian map. *Ann. Probab.* 34, 2144–2202 (2006)
- [58] C. MARZOUK, Scaling limits of random bipartite planar maps with a prescribed degree sequence. *Random Structures Algorithms* 53, 448–503 (2018)
- [59] C. MARZOUK, On scaling limits of random trees and maps with a prescribed degree sequence. *Annales Henri Lebesgue* 5, 317–386 (2022)
- [60] L. MÉNARD, The two uniform infinite quadrangulations of the plane have the same law. *Ann. Inst. Henri Poincaré Probab. Stat.* 46, 190–208 (2010)
- [61] A. METZ-DONNADIEU, Drilling holes in the Brownian disk. In preparation.
- [62] G. MIERMONT, On the sphericity of scaling limits of random planar quadrangulations. *Electron. Comm. Probab.* 13, 248–257 (2008)
- [63] G. MIERMONT, Tessellations of random maps of arbitrary genus. *Ann. Sci. École Norm. Sup.* 42, 735–781 (2009)
- [64] G. MIERMONT, The Brownian map is the scaling limit of uniform random plane quadrangulations. *Acta Math.* 210, 319–401 (2013)
- [65] J. MILLER, W. QIAN, Geodesics in the Brownian map : Strong confluence and geometric structure. *Mem. Amer. Math. Soc.*, to appear. [arXiv:2008.02242](https://arxiv.org/abs/2008.02242)
- [66] J. MILLER, S. SHEFFIELD, An axiomatic characterization of the Brownian map. *J. Éc. polytech. Math.* 8, 609–631 (2021)
- [67] J. MILLER, S. SHEFFIELD, Liouville quantum gravity and the Brownian map I : The QLE(8/3,0) metric. *Invent. Math.* 219, 75–152 (2020)
- [68] J. MILLER, S. SHEFFIELD, Liouville quantum gravity and the Brownian map II : Geodesics and continuity of the embedding. *Ann. Probab.* 49, 2732–2829 (2021)
- [69] J. MILLER, S. SHEFFIELD, Liouville quantum gravity and the Brownian map III : the conformal structure is determined. *Probab. Theory Rel. Fields* 179, 1183–1211 (2021)

- [70] A.M. POLYAKOV, Quantum geometry of bosonic strings. *Phys.Lett. B* 103, 207–210 (1981)
- [71] O. SCHRAMM, Conformally invariant scaling limits : an overview and a collection of problems, in Proceedings of the International Congress of Mathematicians (Madrid 2006), Vol.I, pp. 513–543. European Math. Soc., Zürich, 2007.