

Brownian geometry*

Jean-François Le Gall**

Received: 5 October 2018 / Revised: 3 February 2019 / Accepted: 13 February 2019

Published online: 27 May 2019

© The Mathematical Society of Japan and Springer Japan KK, part of Springer Nature 2019

Communicated by: Takashi Kumagai

Abstract. We present different continuous models of random geometry that have been introduced and studied in recent years. In particular, we consider the Brownian sphere (also called the Brownian map), which is the universal scaling limit of large planar maps in the Gromov–Hausdorff sense, and the Brownian disk, which appears as the scaling limit of planar maps with a boundary. We discuss the construction of these models, and we emphasize the role played by Brownian motion indexed by the Brownian tree.

Keywords and phrases: Random geometry, Brownian sphere, Brownian disk, random planar map, scaling limit, Gromov–Hausdorff convergence, tree-indexed Brownian motion, Brownian tree

Mathematics Subject Classification (2010): 05C80, 05C10, 60C05, 60D05

Contents

1. Introduction	136
2. Discrete and continuous models of random geometry	140
2.1. Planar maps	140
2.2. The Gromov–Hausdorff distance	141
2.3. Convergence to the Brownian sphere	142
3. The construction of the Brownian sphere	144
3.1. The Brownian tree	144

* This article is based on the 21st Takagi Lectures that the author delivered at Research Institute for Mathematical Sciences, Kyoto University on June 23, 2018.

** Supported by the ERC Advanced Grant 740943 GEOBROWN

J.-F. LE GALL

Institut de Mathématique d’Orsay-Bâtiment 307, Université Paris-Sud, 91405 Orsay Cedex, France

(e-mail: jean-francois.legall@math.u-sud.fr)

3.2. Snake trajectories	146
3.3. Constructing a compact metric space from a snake trajectory	147
3.4. Measures on snake trajectories	149
4. Discrete bijections with trees	151
4.1. Schaeffer's bijection	151
4.2. Ideas of the proof of Theorem 1	154
5. Infinite-volume models and the Brownian plane	155
6. Planar maps with a boundary and Brownian disks	158
7. Excursion theory for Brownian motion indexed by the Brownian tree	160
8. Constructing Brownian disks from the positive Brownian snake excursion measure	165
9. Slicing Brownian disks at heights	169

1. Introduction

The goal of this work is to survey a number of recent developments concerning the continuous models of planar random geometry that have been studied extensively in the last ten years, and their connections with discrete models. A very important feature of the continuous models that we will present is their universality, meaning that they appear in the scaling limit of many different discrete models. This is similar of course to the universality of standard Brownian motion, which is the scaling limit of all random walks satisfying mild moment conditions. Partly because of this analogy, and also because Brownian motion plays a crucial role in the construction of our basic objects of study, we use the name Brownian geometry for the general area of continuous models of random geometry that are discussed below. In the present article, we stress the role played by Brownian motion indexed by the Brownian tree, which is the main ingredient of the construction of the random metric space called the Brownian sphere (or Brownian map) and of other models, and which in our opinion is also an important object worth of study in its own. Many properties of Brownian motion indexed by the Brownian tree, in particular the excursion theory presented in Sect. 7 below, have direct applications to Brownian geometry.

The discrete models of random geometry that we will consider are planar maps, which are finite connected graphs embedded in the two-dimensional sphere and viewed up to orientation-preserving homeomorphisms (see Sect. 2 below for a more precise definition). The faces of a planar map are the connected components of the complement of the union of edges, and important particular cases of planar maps are triangulations, respectively, quadrangulations, where all faces are bounded by 3 edges, respectively by 4 edges. We note that many of the results that follow can be extended to graphs embedded in surfaces of higher genus, but we will not discuss these extensions here. Planar maps are important objects of study in combinatorics, and random planar maps have been used for a long time by theoretical physicists as models of random geometry, in the setting of two-dimensional quantum gravity, see in particular [64] and the

book [6]. From the mathematical point of view, a natural question is to consider a planar map chosen at random in a suitable class, say the class of all triangulations with a fixed number n of faces, and to investigate the properties of this object when $n \rightarrow \infty$. One expects, in a way similar to the convergence of rescaled random walks to Brownian motion, that, when its size tends to infinity, the random planar map, suitably rescaled, will be close to a certain continuous model. It turns out that this vague idea can be made precise in the framework of the Gromov–Hausdorff convergence of compact metric spaces (see e.g. [20] for basic facts about the Gromov–Hausdorff distance). Starting from a random planar map M_n uniformly distributed over the class of all triangulations with n faces (or quadrangulations with n faces), one shows [43], [54] that the vertex set $V(M_n)$ equipped with the graph distance rescaled by the factor $n^{-1/4}$ converges in distribution in the Gromov–Hausdorff sense to a limiting random compact metric space which we call the Brownian sphere, see Theorem 1 below (the case of triangulations had been conjectured by Schramm [62]). We note that in previous work the Brownian sphere was called the Brownian map after Markert and Mokkadem [51]. The proof of the latter convergence was strongly motivated by earlier results concerning asymptotics for the two-point function [23] or the three-point function [18] of random quadrangulations. The preceding convergence to the Brownian sphere has been extended to many classes of random planar maps, always with the same limiting space, up to unimportant scaling factors on the distance: This is the universality property of the Brownian sphere, which was already mentioned above.

The construction of the Brownian sphere, and the relevance of Brownian motion indexed by the Brownian tree, are best understood from purely combinatorial considerations about planar maps. Perhaps surprisingly, various classes of planar maps are in one-to-one correspondence with certain classes of discrete trees whose vertices are assigned integer labels. A common feature of these bijections is the fact that labels assigned to the vertices of the tree are closely related to graph distances from a distinguished vertex in the associated planar map. Therefore, a good understanding of the labeled tree associated with a random planar map yields useful information about the metric properties of the vertex set of the planar map equipped with the graph distance. In Sect. 4 below, we present the simplest example of the bijections between planar maps and labeled trees, in the case of quadrangulations.

It turns out that the tree associated with a large random planar map is close, modulo a suitable rescaling, to the continuous random tree which we call the Brownian tree (this is essentially the CRT introduced and studied by Aldous [4], [5]). Furthermore, labels on the tree behave like Brownian motion indexed by the Brownian tree when the size of the planar map goes to infinity. At least informally, these observations explain the construction of the Brownian sphere which is presented in Sect. 3: Following [2], we introduce the concept of a

snake trajectory—this is a convenient framework for studying the Brownian snake driven by a Brownian excursion [40], which is basically the same object as Brownian motion indexed by the Brownian tree—and explain how to associate a compact metric space with a snake trajectory. If the snake trajectory is chosen at random according to the normalized Brownian snake excursion measure, the associated compact metric space is the Brownian sphere, and the “labels” (the values of the tree-indexed Brownian motion) are related to distances from a distinguished point of the Brownian sphere.

The Brownian sphere is by no means the only interesting model in our Brownian geometry. In Sect. 5, we briefly present the Brownian plane, which is an infinite-volume version of the Brownian sphere and can be obtained as the scaling limit of the infinite random lattices called the UIPT (for uniform infinite planar triangulation) and the UIPQ (for uniform infinite planar quadrangulation). In Sect. 6, we introduce Brownian disks as scaling limits of planar maps with a boundary, when the boundary size tends to infinity [14], [16]. In contrast with the Brownian sphere, which is homeomorphic to the two-dimensional sphere, Brownian disks are homeomorphic to the closed unit disk. We pay special attention to the free Brownian disk, which has a fixed boundary size or perimeter but a random volume.

Sect. 8 presents a construction of Brownian disks from a continuous random tree equipped with Brownian labels, which is analogous to the construction of the Brownian sphere, with the difference that the labels now correspond to distances from the boundary (this is in contrast with the previous constructions of [14], [16], which also used labeled trees, but with a different interpretation of labels). Our construction relies on an excursion theory for Brownian motion indexed by the Brownian tree, which is developed in Sect. 7 and is of independent interest. Roughly speaking, if \mathcal{T}_ξ denotes the Brownian tree and $(Z_a)_{a \in \mathcal{T}_\xi}$ denotes Brownian motion indexed by \mathcal{T}_ξ , we describe the distribution of “excursions” of Z away from 0, each excursion corresponding to the restriction of Z to one connected component of $\{a \in \mathcal{T}_\xi : Z_a \neq 0\}$. We obtain that these excursions are independent conditionally given their “boundary sizes”, and distributed according to a certain excursion measure on snake trajectories.

The construction of Sect. 8 makes it possible to identify certain subsets of the Brownian sphere as Brownian disks. In particular, Theorem 17 shows that connected components of the complement of the ball of radius r centered at the distinguished point in the Brownian sphere are independent Brownian disks, conditionally on their boundary sizes and volumes. A similar result holds for the free Brownian disk \mathbb{D} : If $r > 0$ and $H(x)$ denotes the distance from a point $x \in \mathbb{D}$ to the boundary, connected components of the set $\{x \in \mathbb{D} : H(x) > r\}$ are independent free Brownian disks conditionally on their boundary sizes. Finally, in Sect. 9, we present the very recent results of [50] studying the sequence of boundary sizes of the connected components of $\{x \in \mathbb{D} : H(x) > r\}$ as a pro-

cess parameterized by r . We show that this process is a growth-fragmentation process whose distribution is completely determined. The latter result is very closely related to the recent papers [12], [13] investigating scaling limits for a similar process associated with triangulations with a boundary.

Even if it was not possible to provide detailed proofs in this survey, we have tried to sketch the main ideas underlying several important results. We give a detailed presentation of Schaeffer’s bijection between quadrangulations and labeled trees, and, at the end of Sect. 4, we explain informally why the construction of the Brownian sphere, which may appear rather involved at first glance, is a continuous counterpart of this bijection. Similarly in Sect. 8, we emphasize that the study of connected components of the complement of balls in the Brownian sphere can be reduced to the study of excursions of Brownian motion indexed by the Brownian tree.

Let us briefly mention several recent articles that are related to the present work. The paper [26] discusses the Gromov–Hausdorff convergence of rescaled planar maps when the graph distance is replaced by a “local modification”, and shows that the scaling limit is still the Brownian sphere. The study of the UIPQ and the UIPQ has given rise to a number of interesting developments: See in particular [31] for a proof of the recurrence of simple random walk on these infinite random lattices. Hyperbolic versions of the Brownian plane have been studied by Budzinski [19]. The Brownian half-plane, which also appears as the scaling limit of quadrangulations with a boundary when the volume and the boundary size tend to infinity in a suitable way, is discussed in [34] and [9]—a presumably equivalent construction had been given earlier by Caraceni and Curien [21]. The paper [9] provides an exhaustive study of possible scaling limits of quadrangulations with a boundary, leading to new models of Brownian geometry in addition to the Brownian sphere, the Brownian plane or the Brownian disk. In a series of recent papers, Miller and Sheffield [57]–[60] have developed a completely new approach to the Brownian sphere, showing also that this random compact metric space can be equipped with a conformal structure which is linked to Liouville quantum gravity. An important step in this approach [57] was the derivation of an axiomatic characterization of the Brownian sphere. The paper [57] uses a definition of Brownian disks which is different from the one in [14] but which can be shown to be consistent with the latter thanks to the results of [45]. Brownian disks play an important role in the recent work [32], [34], [35] of Gwynne and Miller motivated by the study of statistical physics models on random planar maps. We also mention the paper [36] showing that certain discrete *conformal* embeddings of random planar maps converge to their continuous counterparts. Finally we refer to [48], [55] for pedagogical presentations of random planar maps and the convergence to the Brownian sphere.

2. Discrete and continuous models of random geometry

2.1. Planar maps

The basic discrete model of random geometry that we will consider is a random planar map. Let us start with a precise definition.

Definition 1. *A planar map is a proper embedding of a finite connected graph in the two-dimensional sphere \mathbb{S}^2 . Two planar maps are identified if they correspond via an orientation-preserving homeomorphism of the sphere.*

In this definition, “proper” means that edges are not allowed to cross. The identification modulo homeomorphisms is interpreted by saying that we are only interested in the shape of the embedding, and not in its precise details.

In the preceding definition, we should in fact have written “multigraph” instead of graph, meaning that we allow self-loops and multiple edges. Many of the results that follow are expected to hold, and sometimes have been proved, also for simple planar maps where self-loops and multiple edges are forbidden, but the technicalities become more difficult. See Fig. 1 for an example with a self-loop and a double edge.




Fig. 1. A rooted triangulation with 20 faces

Thanks to the fact that the graph is embedded, we can define the notion of a face. 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 half-edges incident to this face: Note that we say half-edges instead of edges because if both sides of an edge are incident to the same face, this edge

is counted twice in the degree (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 , and we say *triangulation* when $p = 3$, *quadrangulation* when $p = 4$. Fig. 1 shows a triangulation with 20 faces.

We will deal with *rooted* planar maps, meaning that we distinguish an oriented edge, which is called the root edge. The origin of the root edge is called the root vertex, and the face lying to the left of the root edge (this makes sense because the root edge is oriented) is called the root face. See again Fig. 1. Notice that in order to identify two rooted planar maps via an orientation-preserving homeomorphism we require that this homeomorphism preserves the root edge. The reason for dealing with rooted maps comes from the fact that enumeration questions, or bijections between maps and simpler objects such as trees, become more tractable (rooting a map avoids problems related to the presence of symmetries). However, it is strongly believed that the results that follow hold as well for planar maps that are not rooted.

Let $p \geq 3$ and $n \geq 1$ be integers. The set of all rooted p -angulations with n faces will be denoted by \mathbb{M}_n^p . It is easy to see that \mathbb{M}_n^p is empty if p and n are both odd integers. So when p is odd, in particular when $p = 3$, we will implicitly restrict our attention to even values of n . Thanks to the identification in Definition 1, the set \mathbb{M}_n^p is finite, and so it makes sense to choose a rooted p -angulation with n faces uniformly at random.

If M is a planar map, we will denote the vertex set of M by $V(M)$. We equip $V(M)$ with the usual graph distance d_{gr}^M : If v and v' are two vertices of M , $d_{\text{gr}}^M(v, v')$ is the minimal number of edges on a path from v to v' . Our first goal is to study the metric space $(V(M), d_{\text{gr}}^M)$ when M is chosen uniformly at random in \mathbb{M}_n^p (for some fixed p) and when n is large. For this study, we will need a notion of convergence of a sequence of compact metric spaces.

2.2. The Gromov–Hausdorff distance

Let us 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 2 (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 [56] 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 (see [30], Lemma 5.8 for the analogous result in the slightly different setting of the Gromov–Prohorov convergence of metric measure spaces).

2.3. Convergence to the Brownian sphere

We will now discuss the convergence in distribution of $(M_n, n^{-1/4} d_{\text{gr}}^{M_n})$ when M_n is chosen uniformly at random in \mathbb{M}_n^p (for some fixed p). Note that we rescale the graph distance $d_{\text{gr}}^{M_n}$ by the factor $n^{-1/4}$: The need for such a rescaling is clear since one expects that the diameter of the graph blows up when the number of faces grows to infinity. The reason why the correct rescaling factor is $n^{-1/4}$ is more mysterious and will be best understood from the bijections between planar maps and labeled trees that are described below (see the beginning of Sect. 4.2).

The following theorem is proved in [43]. The particular case of quadrangulations $p = 4$ was obtained independently by Miermont [54]. The case $p = 3$ solves a problem of Schramm [62].

Theorem 1 (The scaling limit of p -angulations). *Suppose that either $p = 3$ (triangulations) or $p \geq 4$ is even. Set*

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

For every integer $n \geq 2$ (n even if $p = 3$), let M_n be uniformly distributed over \mathbb{M}_n^p . Then,

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

in the Gromov–Hausdorff sense. The limit $(\mathbf{m}_\infty, d_\infty)$ 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.

In most of the previous work in this area, the Brownian sphere is called the Brownian map after Marckert and Mokkadem [51] 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 3 below and of the related objects called the Brownian disk and the Brownian plane that we shall discuss later. We note that the role of the constants c_p in the theorem is only to ensure that the limit does not depend on p . It is expected that the result of the theorem holds for all values of $p \geq 3$, but the case of odd values $p \geq 5$ seems more difficult to handle for technical reasons.

The fact that the limit does not depend on p is a very important feature of Theorem 1. 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 1, always with the same limit $(\mathbf{m}_\infty, d_\infty)$ hold for general planar maps with a fixed number of edges [15], 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) [3], for planar maps with a prescribed degree sequence [52], etc. We also mention that results similar to Theorem 1 hold if the graph distance is replaced by a “local modification”: The paper [26] 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.

As a general principle, the scaling limit of large random planar maps is expected to be the Brownian sphere whenever some bound is assumed on the degree of faces. On the contrary, if one considers probability distributions on planar maps that favor the appearance of very large faces, different scaling limits may occur (the so-called stable maps of [47]), but we will not discuss this case here.

It is implicit in Theorem 1 that the limit $(\mathbf{m}_\infty, d_\infty)$ 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 2 ([41]). *The Hausdorff dimension of $(\mathbf{m}_\infty, d_\infty)$ is a.s. equal to 4.*

Theorem 3 ([49]). *The compact metric space $(\mathbf{m}_\infty, d_\infty)$ is a.s. 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 2 is in fact relatively easy, but that of Theorem 3 is more intricate and 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 3 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 [46].

3. The construction of the Brownian sphere

In this section, we present a construction of the limiting space $(\mathbf{m}_\infty, d_\infty)$ of Theorem 1. This construction relies on the notion of Brownian motion indexed by the Brownian tree. We start by a brief presentation of the Brownian tree.

3.1. 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 , which will be denoted by $[[a, b]]$, 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. This makes it possible to define a notion of genealogy in the tree \mathcal{T} : If $a, b \in \mathcal{T}$, we say that b is a descendant of a , or a is an ancestor of b , if $a \in [[\rho, b]]$.

In the present work, we will consider only compact \mathbb{R} -trees, and we will use the fact that such trees can be coded by continuous functions. Let $h : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ be a nonnegative continuous function on \mathbb{R}_+ such that $h(0) = 0$. We assume that h has compact support, so that

$$\sigma_h := \sup\{t \geq 0 : h(t) > 0\} < \infty.$$

Here and later we make the convention that $\sup \emptyset = 0$.

For every $s, t \in \mathbb{R}_+$, 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}_+ , and thus we may introduce the associated equivalence relation on \mathbb{R}_+ , 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 \mathbb{R}_+ / \sim_h .

Lemma 4 ([29]). *The quotient space $\mathcal{T}_h := \mathbb{R}_+ / \sim_h$ equipped with the distance d_h is a compact \mathbb{R} -tree.*

The \mathbb{R} -tree (\mathcal{T}_h, d_h) will be 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 $\rho = p_h(0)$. In addition, it is often convenient to equip \mathcal{T}_h with a *volume measure*, which is defined as the push forward of Lebesgue measure on $[0, \sigma_h]$ 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 , but we will not need this fact.

The coding by a function makes it possible to define “lexicographical” intervals on the tree. Let us explain this. If $s, t \geq 0$ and $s > t$, we make the convention that $[s, t] = [s, \infty) \cup [0, t]$ (of course, if $s \leq t$, $[s, t]$ is the usual interval). If $a, b \in \mathcal{T}$ with $a \neq b$, there is a smallest “interval” $[s, t]$ with $s, t \geq 0$ (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”. We note that $[a, b] \cup [b, a]$ is not the whole tree, unless a and b are both leaves (vertices whose removal does not disconnect the tree).

Let us now randomize h . We let $\mathbf{n}(dh)$ stand for Itô’s excursion measure of positive excursions of linear Brownian motion (see e.g. [61], Chap. XII) normalized so that, for every $\varepsilon > 0$,

$$\mathbf{n}(\max_{s \geq 0} h(s) > \varepsilon) = \frac{1}{2\varepsilon}.$$

Under $\mathbf{n}(dh)$, we will write $\sigma = \sigma_h$ for the duration of the excursion h . It will also be convenient to introduce the conditional probability measure $\mathbf{n}_{(s)} := \mathbf{n}(\cdot | \sigma = s)$, for every $s > 0$. In particular $\mathbf{n}_{(1)}$ is the law of the normalized excursion, and we have

$$\mathbf{n} = \int_0^\infty \mathbf{n}_{(s)} \frac{ds}{2\sqrt{2\pi s^3}}.$$

Definition 3. *The Brownian tree is the tree \mathcal{T}_h coded by h under $\mathbf{n}(dh)$.*

It is important to realize that \mathbf{n} is an infinite measure. We can also consider the tree \mathcal{T}_h under the probability measure $\mathbf{n}_{(1)}(dh)$, and this random tree is Aldous' continuum random tree, also called the CRT (our normalization is slightly different from the one in [4], [5]). However, it is often more convenient to argue under the infinite measure \mathbf{n} .

3.2. Snake trajectories

We now propose to discuss Brownian motion indexed by the Brownian tree of Definition 3. The fact that we are interested in a random process indexed by a random set creates some technical difficulties, which we will avoid here by introducing the concept of a *snake trajectory*.

A finite real path is a continuous mapping $w : [0, \zeta_{(w)}] \rightarrow \mathbb{R}$, where the number $\zeta_{(w)} \geq 0$ is called the lifetime of w . We let \mathcal{W} denote the space of all finite paths in \mathbb{R} . The set \mathcal{W} is a Polish space when equipped with the distance

$$d_{\mathcal{W}}(w, w') = |\zeta_{(w)} - \zeta_{(w')}| + \sup_{t \geq 0} |w(t \wedge \zeta_{(w)}) - w'(t \wedge \zeta_{(w')})|.$$

The endpoint or tip of the path w is denoted by $\widehat{w} = w(\zeta_{(w)})$. For every $x \in \mathbb{R}$, we set $\mathcal{W}_x = \{w \in \mathcal{W} : w(0) = x\}$. The trivial element of \mathcal{W}_x with zero lifetime is identified with the point x —in this way we view \mathbb{R} as the subset of \mathcal{W} consisting of all finite paths with zero lifetime.

Definition 4. Let $x \in \mathbb{R}$. A *snake trajectory with initial point x* is a continuous mapping

$$\begin{aligned} \omega : \mathbb{R}_+ &\longrightarrow \mathcal{W}_x \\ s &\longmapsto \omega_s \end{aligned}$$

which satisfies the following two properties:

- (i) We have $\omega_0 = x$ and the number $\sigma(\omega) := \sup\{s \geq 0 : \omega_s \neq x\}$, called the duration of the snake trajectory ω , is finite.
- (ii) For every $0 \leq s \leq s'$, we have

$$\omega_s(t) = \omega_{s'}(t), \quad \text{for every } 0 \leq t \leq \min_{s \leq r \leq s'} \zeta_{(\omega_r)}.$$

Property (i) implies in particular that the function $s \mapsto \zeta_{(\omega_s)}$ has compact support.

Important Remark. A snake trajectory ω is completely determined by the knowledge of the lifetime function $s \mapsto \zeta_{(\omega_s)}$ and the tip function $s \mapsto \widehat{\omega}_s = \omega_s(\zeta_{(\omega_s)})$. Indeed, for any $s \geq 0$ and $r \in [0, \zeta_{(\omega_s)}]$, if $\theta_s(r) = \inf\{u \geq s : \zeta_{(\omega_u)} = r\}$, property (ii) implies that $\omega_s(r) = \widehat{\omega}_{\theta_s(r)}$.

We write \mathcal{S}_x for the set of all snake trajectories with initial point x , and

$$\mathcal{S} := \bigcup_{x \in \mathbb{R}} \mathcal{S}_x$$

for the set of all snake trajectories.

Let $\omega \in \mathcal{S}$. Then the real function $s \mapsto \zeta_{(\omega_s)}$ satisfies the conditions required to define the tree coded by this function (cf. Sect. 3.1) and we will write \mathcal{T}_ζ for this tree, and p_ζ for the canonical projection from \mathbb{R}_+ onto \mathcal{T}_ζ . We sometimes say that \mathcal{T}_ζ is the genealogical tree of the snake trajectory ω . Property (ii) in Definition 4 implies that $\omega_s = \omega_{s'}$ whenever $p_\zeta(s) = p_\zeta(s')$. In other words, ω_s only depends on the equivalence class of s in the quotient space \mathcal{T}_ζ , and the mapping $s \mapsto \omega_s$ induces a function defined on the genealogical tree \mathcal{T}_ζ . We should think of the collection $(\omega_s)_{s \geq 0}$ as forming a “tree of paths” whose genealogy is prescribed by \mathcal{T}_ζ (see the left side of Fig. 6 below for an illustration).

Notation In what follows, we will consider snake trajectories ω that may be deterministic or chosen according to a measure on \mathcal{S} , and we will use the notation $W_s = W_s(\omega) = \omega_s$, and $\zeta_s = \zeta_s(\omega) = \zeta_{(\omega_s)}$.

3.3. Constructing a compact metric space from a snake trajectory

The Brownian sphere of Theorem 1 is constructed from a random snake trajectory distributed according to a certain probability measure. To explain this construction, it is best to consider first the case of a deterministic snake trajectory ω .

So we fix $\omega \in \mathcal{S}_0$ and we recall that \mathcal{T}_ζ is the tree coded by $(\zeta_s)_{s \geq 0}$ (we use the notation explained at the end of Sect. 3.2). If $a \in \mathcal{T}_\zeta$, we set $Z_a = \widehat{W}_s$ if s is such that $p_\zeta(s) = a$ and we also say that W_s is the historical path of a (by preceding observations, this does not depend on the choice of s such that $p_\zeta(s) = a$). We view $(Z_a)_{a \in \mathcal{T}_\zeta}$ as a collection of labels assigned to the points of \mathcal{T}_ζ . Note that the function $a \mapsto Z_a$ is continuous on \mathcal{T}_ζ .

We will now associate a metric space with the space trajectory ω , and roughly speaking this metric space will be obtained from the genealogical tree \mathcal{T}_ζ by gluing together certain pairs of points. Let us turn to a precise definition. For every $a, b \in \mathcal{T}_\zeta$, we set

$$D^\circ(a, b) = Z_a + Z_b - 2 \max\left(\min_{c \in [a, b]} Z_c, \min_{c \in [b, a]} Z_c\right), \quad (1)$$

where we recall that $[a, b]$ stands for the lexicographical interval from a to b in \mathcal{T}_ζ . We note that $D^\circ(a, b) = 0$ if and only if

$$Z_a = Z_b = \max\left(\min_{c \in [a, b]} Z_c, \min_{c \in [b, a]} Z_c\right), \quad (2)$$

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 $a, b \in \mathcal{T}_\xi$,

$$D(a, b) = \inf \left\{ \sum_{i=1}^k D^\circ(a_{i-1}, a_i) \right\}, \quad (3)$$

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}_ξ such that $a_0 = a$ and $a_k = b$. We note that

$$D(a, b) \geq |Z_a - Z_b| \quad (4)$$

as an immediate consequence of the similar bound for D° .

We now observe that D is a pseudo-metric on \mathcal{T}_ξ , and we let \mathcal{M} be the associated quotient space, which is the quotient of \mathcal{T}_ξ for the equivalence relation $a \approx b$ if and only if $D(a, b) = 0$. We equip \mathcal{M} with the distance induced by D , for which we keep the same notation. We note that (\mathcal{M}, D) is a compact metric space, and we let Π denote the canonical projection from \mathcal{T}_ξ onto \mathcal{M} . By abuse of notation, for every $x \in \mathcal{M}$, we write $Z_x = Z_a$ if $x = \Pi(a)$ (by (4) this does not depend on the choice of a such that $x = \Pi(a)$). So labels can also be viewed as defined on the quotient space \mathcal{M} . Later it will be convenient to have a volume measure $\mathbf{v}(dx)$ on \mathcal{M} , which is defined as the push forward of the volume measure on \mathcal{T}_ξ under Π .

The preceding construction obviously depends on the choice of ω , which was fixed in the beginning of this section. We claim that it does so in a measurable way.

Lemma 5. *The mapping $\omega \mapsto (\mathcal{M}, D)$ defined above, with values in the space (\mathbb{K}, d_{GH}) , is measurable.*

We refer to [45], Lemma 6 for the proof of a more precise statement.

Let us mention some properties of D that will play a role later. Let a_* be any point of \mathcal{T}_ξ such that

$$Z_{a_*} = \inf_{a \in \mathcal{T}_\xi} Z_a.$$

The existence of such a point follows from a compactness argument (notice that a_* may not be unique, but it will follow from (5) below that $\Pi(a_*)$ is uniquely determined). Then we have, for every $a \in \mathcal{T}_\xi$,

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

The lower bound $D(a_*, a) \geq Z_a - Z_{a_*}$ is immediate from (4). The corresponding upper bound is also trivial since it is clear that $D^\circ(a_*, a) = Z_a - Z_{a_*}$. So setting $Z_* = Z_{a_*}$ and $x_* = \Pi(a_*)$, we get that, for every $x \in \mathcal{M}$,

$$D(x_*, x) = Z_x - Z_*. \quad (5)$$

We interpret this by saying that \mathcal{M} has a distinguished point x_* such that labels exactly correspond to distances from x_* , up to the shift by Z_* .

3.4. Measures on snake trajectories

We start with a key lemma.

Lemma 6. *Let $h : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ be a continuous function with compact support such that $h(0) = 0$. Assume that h is Hölder continuous, meaning that there exist positive constants $\delta \in (0, 1]$ and C such that $|h(s) - h(s')| \leq C|s - s'|^\delta$ for every $s, s' \geq 0$. Then there exists a random snake trajectory $W^h = (W_s^h)_{s \geq 0}$ with initial point 0 such that:*

- (i) $\zeta_{(W_s^h)} = h(s)$, for every $s \geq 0$, a.s.
- (ii) The process $(\widehat{W}_s^h)_{s \geq 0}$ is a centered Gaussian process with covariance

$$\text{cov}(\widehat{W}_s^h, \widehat{W}_{s'}^h) = \min_{s \wedge s' \leq r \leq s \vee s'} h(r).$$

The process $(W_s^h)_{s \geq 0}$ is called the Brownian snake driven by the function h .

We note that the distribution of W^h is completely determined by properties (i) and (ii), thanks to the remark following Definition 4. The intuition underlying the preceding definition is as follows: For every $s \geq 0$, W_s^h is a Brownian path with lifetime $h(s)$, when $h(s)$ decreases the path W_s^h is erased from its tip and when $h(s)$ increases the path W_s^h is extended by adding “little pieces of Brownian paths” at its tip.

The proof of Lemma 6 is straightforward: See [47], Sect. 4.2 for a detailed argument. Note that the Hölder continuity assumption of h is used to warrant the existence of a continuous modification of a process satisfying properties (i) and (ii) of Lemma 6.

As a consequence of (ii), we have $W_0^h = 0$ and

$$\mathbb{E}[(\widehat{W}_s^h - \widehat{W}_{s'}^h)^2] = d_h(s, s'),$$

where the pseudo-metric d_h was defined in Sect. 3.1. Since we already noted that the snake trajectory W^h can be viewed as indexed by the tree \mathcal{T}_h , the last display justifies the fact that \widehat{W}^h is interpreted as Brownian motion indexed by \mathcal{T}_h . In fact, if $\varphi : [0, u] \rightarrow \mathcal{T}_h$ is an isometry mapping the interval $[0, u]$ onto a

line segment of \mathcal{T}_h , we immediately see that $(\widehat{W}_{\varphi(r)}^h - \widehat{W}_{\varphi(0)}^h)_{0 \leq r \leq u}$ is a linear Brownian motion.

If h satisfies the properties in Lemma 6, we let \mathbf{P}^h stand for the distribution of W^h , which is thus a probability measure on the space \mathcal{S}_0 . We now introduce Brownian snake excursion measures, which will play a major role in what follows.

Definition 5. *The Brownian snake excursion measure \mathbb{N}_0 is the σ -finite measure on \mathcal{S}_0 defined by*

$$\mathbb{N}_0(d\omega) = \int \mathbf{n}(dh) \mathbf{P}^h(d\omega).$$

Similarly, the normalized Brownian snake excursion measure is the probability measure on \mathcal{S}_0 defined by

$$\mathbb{N}_0^{(1)}(d\omega) = \int \mathbf{n}_{(1)}(dh) \mathbf{P}^h(d\omega).$$

In other words, to construct a random snake trajectory distributed according to \mathbb{N}_0 (resp. according to $\mathbb{N}_0^{(1)}$) we just pick a Brownian excursion h distributed according to \mathbf{n} (resp. a normalized Brownian excursion) and consider the Brownian snake driven by h . This makes sense because we know that h is Hölder continuous, $\mathbf{n}(dh)$ or $\mathbf{n}_{(1)}(dh)$ a.e.

Lemma 5 now allows us to set the following definition.

Definition 6. *The Brownian sphere is the random compact metric space (\mathcal{M}, D) obtained via the construction of Sect. 3.3 from a snake trajectory ω distributed according to $\mathbb{N}_0^{(1)}$.*

One can prove [41] that $\mathbb{N}_0^{(1)}$ a.s., for every $a, b \in \mathcal{T}_\xi$ the property $D(a, b) = 0$ holds if and only if $D^\circ(a, b) = 0$ (the fact that $D^\circ(a, b) = 0$ implies $D(a, b) = 0$ is obvious since $D \leq D^\circ$). So the construction of the Brownian sphere can be summarized by saying that we start from the CRT \mathcal{T}_ξ equipped with “Brownian labels” $(Z_a)_{a \in \mathcal{T}_\xi}$, and we identify points a and b of the CRT if and only if $D^\circ(a, b) = 0$, which has a simple interpretation as explained above after (2) (furthermore the metric D is the largest metric bounded above by D°).

It is often useful to consider also the *free Brownian sphere*, which is just the metric space (\mathcal{M}, D) under the measure \mathbb{N}_0 . Many properties of the free Brownian sphere are “nicer” than those of the “standard Brownian sphere” because there is no constraint on the total volume, but the price to pay is to work under a σ -finite measure.

4. Discrete bijections with trees

4.1. Schaeffer's bijection

In this section, we explain a bijection between quadrangulations and (discrete) labeled trees, which can be found in [23] and is in some sense a discrete analog of the construction of the Brownian sphere that was given in the previous section. In fact this discrete bijection (and its generalizations) plays a major role in the proof of Theorem 1, and helps in the understanding of the definition of the Brownian sphere and of the metric D . We restrict our attention to the case of quadrangulations because the description is simpler in that case, but we immediately mention that similar bijections exist for more general planar maps (see in particular [17]).

We first need to introduce the class of discrete trees that will be relevant. First recall that a *plane tree* τ is a (finite) rooted ordered tree. A plane tree can be specified by representing each vertex 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 (ignore for the moment the circled figures). To make the connection with planar maps, we 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.




Fig. 2. Schaeffer's 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 quadrangulations with 8 faces (case $\varepsilon = -1$)

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.

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.

We then claim that there is a one-to-one correspondence between the sets $\mathbb{M}_n^{4,\bullet}$ and $\mathbb{T}_n \times \{-1, 1\}$ (this correspondence is called Schaeffer's bijection). 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 (for instance, in the tree of the left side of Fig. 2, the root \emptyset has 2 corners, the vertex 21 has 3 corners, and the vertex 221 has only one corner). 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 (see the middle part of Fig. 2). We agree that every corner inherits the label of the vertex to which it is incident.

With the labeled tree $(\tau, (\ell_v)_{v \in V(\tau)})$, we associate a quadrangulation M by the following procedure. 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 (in 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), 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 where, for instance, there are edges of M connecting c_0 to c_1 , c_1 to v_* , c_2 to c_5 , c_3 to c_4 , etc.).

We still have to define the root of the quadrangulation 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 7. *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 $v \in V(\tau)$,*

$$d_{\text{gr}}^M(v_*, v) = \ell_v - \min_{u \in V(\tau)} \ell_u + 1. \quad (6)$$

Let us explain why property (6) 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 Schaeffer's 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 Schaeffer's 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 Schaeffer's bijection.

Property (6) is useful when studying the metric properties of M (in view of proving the case $p = 4$ of Theorem 1). However, (6) only gives information about distances from the distinguished vertex v_* , which is far from sufficient if one is interested in the Gromov–Hausdorff convergence. If v and v' are two arbitrary vertices of M , there is however a very 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 that was already used in Schaeffer's bijection. For every $i \in \{0, 1, \dots, 2n-1\}$, let v_i be the vertex corresponding to the corner c_i . Then, if $0 \leq i < j \leq 2n-1$, we have

$$d_{\text{gr}}^M(v_i, v_j) \leq \ell_{v_i} + \ell_{v_j} - 2 \max\left(\min_{k \in [i, j]} \ell_{v_k}, \min_{k \in [j, 2n-1] \cup [0, i]} \ell_{v_k}\right) + 2. \quad (7)$$

The proof of this bound is easy. Consider the geodesic path γ from the corner c_i to v_* constructed as in the proof of (6), and the similar geodesic 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 (7) minus 1. The concatenation of these two geodesic paths up to their coalescence time thus gives a path from v_i to v_j whose length is the right-hand side of (7).

4.2. Ideas of the proof of Theorem 1

Schaeffer’s bijection allows us to sketch the main ideas of the proof of Theorem 1 in the case of quadrangulations. We start from a uniformly distributed rooted and pointed quadrangulation M_n with n faces (the fact that we consider a rooted and pointed quadrangulation rather than a rooted quadrangulation as in Theorem 1 is unimportant since by “forgetting” the distinguished vertex of M_n we get a uniformly distributed rooted quadrangulation), and we let $(\tau_n, (\ell_v^n)_{v \in V(\tau_n)})$ be the associated labeled tree. We note that τ_n is uniformly distributed over the set of all plane trees with n edges, because for every such tree there is the same number 3^n of possible assignments of labels. It is well known that the height of the tree τ_n is of order \sqrt{n} when n is large, and, from the central limit theorem, one may guess that the maximal and the minimal label in τ_n are of order $\sqrt{\sqrt{n}} = n^{1/4}$ (just note that conditionally given τ_n , the increments of labels along the different edges of τ_n are independent and uniformly distributed over $\{-1, 0, 1\}$). Recalling (6), we see that the diameter of M_n must be of order $n^{1/4}$, which explains the rescaling in Theorem 1.

Then, a well-known result of Aldous shows that the tree τ_n viewed as a metric space for the graph distance rescaled by the factor $1/\sqrt{2n}$ converges in distribution to the CRT—with our particular normalization of the CRT. This convergence can be stated in a more precise form using the so-called “contour functions” which keep track of the lexicographical order on the trees. Furthermore, using the fact that the variance of the uniform distribution on $\{-1, 0, 1\}$ is $2/3$, one gets that the labels rescaled by $(2/3)^{-1/2}(2n)^{-1/4}$ converge to Brownian motion indexed by the CRT (we do not make the meaning of this convergence precise here). This suggests that the scaling limit of M_n can be described in terms of the CRT equipped with Brownian labels. However, in contrast with the discrete picture, we need to perform some identification of vertices of the CRT. Let us explain this. Writing again c_0, \dots, c_{2n-1} for the sequence of corners of the tree τ_n , we note that for $i < j$, the corner c_i is connected to the corner c_j by an edge of M_n as soon as

$$\ell_{c_j} = \ell_{c_i} - 1 \text{ and } \ell_{c_k} \geq \ell_{c_i} \text{ for every } k \in \{i, i + 1, \dots, j - 1\}.$$

The point is now that, even for large values of n , this property will hold for certain pairs (i, j) such that $(j - i)$ is of order n . Because of the rescaling of the graph distance by $n^{-1/4}$, which informally implies that two adjacent vertices are identified in the scaling limit, this means that certain pairs of distinct points of the CRT must be glued together.

Finally, a tightness argument relying on the bound (7) can be used to verify that sequential limits of $(V(M_n), n^{-1/4}d_{\text{gr}}^{M_n})$ exist in the Gromov–Hausdorff sense, and are represented as quotient spaces of the CRT (equipped with Brownian labels) for a certain pseudo-metric D . The discrete bound (7) implies that

the pseudo-metric D satisfies $D \leq D^\circ$, where D° is defined in (1). It immediately follows that D must be bounded above by the right-hand side of (3). The remaining part of the argument, which unfortunately is much harder, is to verify that (3) indeed holds.

5. Infinite-volume models and the Brownian plane

The random planar maps discussed in the preceding sections are finite (random) graphs embedded in the sphere. It turns out that one can also define infinite random lattices that are limits in a certain sense of uniformly distributed triangulations or quadrangulations with a fixed number of faces (one could consider more general planar maps, see in particular [63]). A pioneering work of Angel and Schramm [8], which (together with the companion paper [7] and the Chassaing–Schaeffer paper [23]) motivated much of the subsequent research about random planar maps, introduced the so-called uniform infinite planar triangulation or UIPT as the local limit of uniformly distributed triangulations with a fixed number of faces—in fact, Angel and Schramm considered “type II triangulations” where self-loops are not allowed, but the analogous construction for general triangulations can be found in [63]. Let us present the analog of the Angel–Schramm construction for quadrangulations, which is due to Krikun [39].

If M is a rooted planar map with root vertex ρ , and $r \geq 1$ is an integer, the ball of radius r in M , which is denoted by $B_r(M)$, is the rooted planar map obtained by keeping only those faces of M that are incident to a vertex whose graph distance from ρ is at most $(r - 1)$. See Fig. 3 for an illustration in the case of a quadrangulation. This definition of balls can be extended to infinite (rooted) planar lattices, meaning infinite (rooted) connected graphs properly embedded in the plane.

For every $n \geq 1$, let Q_n be uniformly distributed over the set \mathbb{M}_n^4 of all rooted quadrangulations with n faces. Then one proves [39] that there exists an infinite random rooted planar lattice Q_∞ such that, for every integer $r \geq 1$ and for every rooted planar map M , we have

$$\mathbb{P}(B_r(Q_n) = M) \xrightarrow{n \rightarrow \infty} \mathbb{P}(B_r(Q_\infty) = M).$$

The infinite random lattice Q_∞ is called the uniform infinite planar quadrangulation or UIPQ. It is the *local limit* of Q_n as $n \rightarrow \infty$, meaning that the distribution of what one sees in Q_n in a fixed neighborhood of the root vertex “stabilizes” when $n \rightarrow \infty$ to the distribution of the corresponding neighborhood of the root vertex in Q_∞ . We emphasize that this convergence is very different from the convergence in Theorem 1 (which also dealt with uniformly distributed quadrangulations): Here there is no rescaling of the graph distance, and the limit




Fig. 3. A large quadrangulation Q near its root vertex ρ and in grey the ball $B_2(Q)$

is an infinite random lattice instead of a random compact metric space. Both the Krikun paper [39] and the Angel–Schramm work [8] for triangulations relied on enumeration techniques, but a different approach to the UIPQ based on bijections with labeled trees was proposed by Chassaing and Durhuus [22] (the equivalence between this approach and Krikun’s one was later established by Ménard [53]). A simple construction of the UIPQ, relying on the version of Schaeffer’s bijection presented in Sect. 4.1, can be found in [27].

The UIPQ is an infinite-volume limit of finite quadrangulations. In the same way, one may ask about the existence of an infinite-volume version of the Brownian sphere. This is the Brownian plane, which was constructed in the joint paper [24] with Nicolas Curien, and appears in the following theorem as a scaling limit for the UIPQ. Before stating this theorem, recall that a metric space is called boundedly compact if every closed bounded set is compact. Write $d_{\text{gr}}^{Q_\infty}$ for the graph distance on the vertex set $V(Q_\infty)$, and view $(V(Q_\infty), d_{\text{gr}}^{Q_\infty})$ as a pointed metric space, where the distinguished point is the root vertex.

Theorem 8 ([24]). *There exists a random boundedly compact pointed metric space (\mathcal{P}, D_∞) such that*

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

where the convergence holds in distribution in the local Gromov–Hausdorff sense.

The local Gromov–Hausdorff convergence (in distribution) means that, for every real $r > 0$, the closed ball of radius r centered at the distinguished point of $V(Q_\infty)$ converges (in distribution) to the same ball centered at the distinguished point in the limiting space \mathcal{P} , in the sense of the Gromov–Hausdorff distance for compact spaces. Just like the Brownian sphere, the Brownian plane is believed to be a universal object, and in fact a version of the preceding theorem for the UIPT has been proved by Budzinski [19] with the same limiting space.

We refer to [24] for the construction of the Brownian plane, which is a continuous analog of the construction of the UIPQ in [27] (a slightly different approach to the Brownian plane is given in [25]). The construction of [24] is very similar to the construction of the Brownian sphere in Sect. 3. The key ingredient is now Brownian motion indexed by the infinite Brownian tree, which can be understood as the Brownian tree conditioned on non-extinction.

One may obtain the Brownian plane as a limiting object in a variety of different ways. For instance, starting from the Brownian sphere $(\mathbf{m}_\infty, d_\infty)$ of Theorem 1 and viewing \mathbf{m}_∞ as a pointed space with distinguished point x_* (cf. the end of Sect. 3.3), one checks that (\mathcal{P}, D_∞) is the limit of $(\mathbf{m}_\infty, \lambda d_\infty)$ when $\lambda \rightarrow \infty$, in the local Gromov–Hausdorff sense. In the terminology of [20], one may say that the Brownian plane is the tangent cone (in distribution) of the Brownian sphere at x_* . Alternatively one can start from the uniformly distributed quadrangulation Q_n and scale the distance by a factor ε_n tending to 0 less fast than $n^{-1/4}$. Fig. 4 gives a diagram taken from [24] that summarizes these convergences in distribution, together with those of Theorems 1 and 8.




Fig. 4. Convergence to the Brownian plane

In a way similar to Theorem 3, the Brownian plane is homeomorphic to the usual plane. On the other hand, the Brownian plane shares the same local properties as the Brownian sphere (in fact in a strong sense, since one can couple the Brownian plane and the Brownian sphere so that the respective balls of suffi-

ciently small radius centered at the distinguished point are isometric, see [24]). In particular, the Hausdorff dimension of the Brownian plane is also equal to 4. Furthermore, the Brownian plane enjoys an additional property of scaling invariance: For every $\lambda > 0$, the space $(\mathcal{P}, \lambda D_\infty)$ has the same distribution as (\mathcal{P}, D_∞) . This makes certain calculations more tractable in the Brownian plane than in the Brownian sphere: See [25] for several remarkable distributions related to the Brownian plane.

6. Planar maps with a boundary and Brownian disks

In this section we introduce Brownian disks as scaling limits of quadrangulations with a boundary. Brownian disks are models of random geometry which unlike the Brownian sphere are homeomorphic to the closed disk. Nonetheless, Brownian disks are very closely related to the Brownian sphere, and, as we will discuss later, various subsets of the Brownian sphere can be identified as Brownian disks.

Let us start with a basic definition. Recall that the root face of a rooted planar map is the face lying to the left of the root edge.

Definition 7. *A quadrangulation with a (general) boundary is a rooted planar map Q such that all faces but the root face have degree 4. The root face is also called the outer face and the other faces are called inner faces. The degree of the outer face, which is an even integer, is called the boundary size or the perimeter of Q .*

See Fig. 5 for an example. One could also consider p -angulations with a boundary (in particular triangulations with a boundary) but for the sake of simplicity we restrict our attention to quadrangulations.

For every integer $k \geq 1$, we denote the set of all pointed quadrangulations with a boundary of 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

$$\#\mathbb{Q}_n^{\partial,k} \underset{n \rightarrow \infty}{\sim} b_k 12^n n^{-5/2}.$$

See formula (4) in [28].

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.

The following result, which is analogous to Theorem 1, is a special case of [16], Theorem 8.




Fig. 5. A quadrangulation with a boundary of size 14

Theorem 9. For every integer $k \geq 1$, let B_k be a Boltzmann quadrangulation with a boundary of size $2k$. Then,

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

where the convergence holds in distribution for the Gromov–Hausdorff topology. The limiting random compact metric space (\mathbb{D}, D^∂) is called the free Brownian disk with perimeter 1.

The factor $\sqrt{3/2}$ in the convergence of the theorem is present only to allow a simpler description of the limit in the next section.

In contrast with Theorem 1, we notice that the number of faces of B_k is not fixed, but only its perimeter. One can verify that the number of faces of B_k is typically of order k^2 , and so the scaling factor $k^{-1/2}$ in Theorem 9 corresponds to the factor $n^{-1/4}$ in Theorem 1. One can prove versions of Theorem 9 for quadrangulations where both the boundary size and the volume (number of faces) are fixed and grow to infinity simultaneously in such a way that the volume stays proportional to the square of the boundary size: This leads to the definition of Brownian disks with given perimeter and volume. See [9] for a discussion of all possible scaling limits of quadrangulations with a boundary, and [33] for an analog of Theorem 9 in the case of quadrangulations with a *simple* boundary.

For every $a > 0$, the free Brownian disk with perimeter a may be defined as the random metric space $(\mathbb{D}, \sqrt{a} D^\partial)$.

One proves [14] that the Brownian disk is homeomorphic to the closed unit disk, and this makes it possible to define the boundary $\partial\mathbb{D}$ as the set of all points in \mathbb{D} that have no neighborhood homeomorphic to the open unit disk.

7. Excursion theory for Brownian motion indexed by the Brownian tree

In this section, which is mostly taken from the joint paper [2] with Céline Abraham, we discuss an excursion theory for Brownian motion indexed by the Brownian tree. An important motivation is to derive a construction of Brownian disks which is analogous to the construction of the Brownian sphere explained in Sect. 3. However, we believe that this excursion theory is interesting in its own and should have many other applications. There is of course a strong analogy with the classical Itô theory [37] but also important differences due to the fact that the parameter set is a tree, and so connected components of the complement of the zero set of Brownian motion are \mathbb{R} -trees instead of intervals, as appear in the classical setting.

Recall from Definition 5 the σ -finite measure $\mathbb{N}_0(d\omega)$ on the space of snake trajectories with initial point 0, and the notation $W_s(\omega) = \omega_s$, $\zeta_s(\omega) = \zeta_{(\omega_s)}$ for $s \geq 0$, and $\sigma(\omega) = \sup\{s \geq 0 : \zeta_s \neq 0\}$. The “Brownian tree” \mathcal{T}_ξ is the tree coded by the function $(\zeta_s)_{s \geq 0}$ as explained in Sect. 3.1, and we use the notation $Z_a = \widehat{W}_s$ if $a = p_\xi(s)$, where we recall that \widehat{W}_s denotes the tip of the finite path W_s , and p_ξ stands for the canonical projection from \mathbb{R}_+ onto \mathcal{T}_ξ . The collection $(Z_a)_{a \in \mathcal{T}_\xi}$ is thus our Brownian motion indexed by the Brownian tree.

In a way very similar to classical excursion theory, our aim is to describe the process Z restricted to a connected component of $\{b \in \mathcal{T}_\xi : Z_b \neq 0\}$. To this end we first introduce the notion of an *excursion debut*. We say that $a \in \mathcal{T}_\xi$ is an excursion debut if

- (i) $Z_a = 0$;
- (ii) a has a strict descendant a' such that $Z_b \neq 0$ for every $b \in \llbracket a, a' \rrbracket$.

In (ii), we use the obvious notation $\llbracket a, a' \rrbracket = \llbracket a, a' \rrbracket \setminus \{a\}$. We then observe that connected components of $\{b \in \mathcal{T}_\xi : Z_b \neq 0\}$ are in one-to-one correspondence with excursion debuts: The connected component \mathcal{C}_a associated with an excursion debut a is just the set of all strict descendants a' of a such that the property $Z_b \neq 0$ for every $b \in \llbracket a, a' \rrbracket$ holds.

We will now explain how the values of Z on a given connected component can be represented by a snake trajectory. In what follows, we fix an excursion debut a but note that the discussion applies to *every* excursion debut, almost surely. The fact that a has strict descendants implies that there are exactly two times $u < v$ such that $p_\xi(u) = p_\xi(v) = a$ (there could be three such times if a were a branching point of \mathcal{T}_ξ , but this case is excluded because branching points have nonzero labels, almost surely). Recall that $W_u = W_v$ is called the historical path of a . We note that the descendants of a are exactly the points $p_\xi(s)$ for $s \in [u, v]$. We can then define a snake trajectory $\tilde{W}^{(a)} = (\tilde{W}_s^{(a)})_{s \geq 0}$ in \mathcal{S}_0 , which describes the labels of descendants of a , by setting for every $s \geq 0$,

$$\tilde{W}_s^{(a)}(t) := W_{(u+s) \wedge v}(\zeta_u + t), \quad \text{for } 0 \leq t \leq \tilde{\zeta}_s^{(a)} := \zeta_{(u+s) \wedge v} - \zeta_u.$$

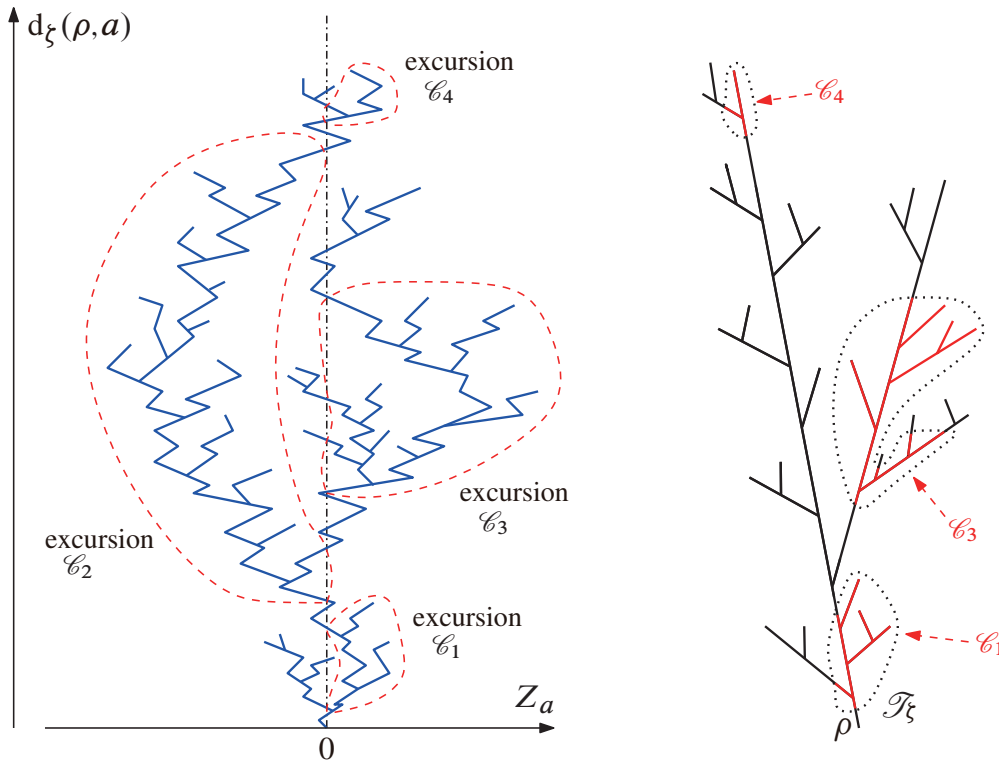


Fig. 6. A schematic representation of excursions. The right side shows the tree \mathcal{T}_ξ , and the parts of the tree inside the dotted lines are a few connected components of the set $\{a \in \mathcal{T}_\xi : Z_a \neq 0\}$. The left side shows the values of Z_a for $a \in \mathcal{T}_\xi$, or equivalently the paths W_s which form a “tree of Brownian paths”, and the parts inside the dashed lines are a few excursions away from 0

In fact we are not interested in all descendants of a , but only in those that lie in the associated connected component \mathcal{C}_a . For this reason, we introduce the time change

$$\eta_s^{(a)} := \inf \left\{ r \geq 0 : \int_0^r dt \mathbf{1}_{\{\tau_0^*(\tilde{W}_t^{(a)}) \geq \tilde{\zeta}_t^{(a)}\}} > s \right\},$$

where we use the notation $\tau_0^*(w) = \inf\{t \in (0, \zeta_{(w)}] : w(t) = 0\}$ for $w \in \mathcal{W}$, with the usual convention $\inf \emptyset = +\infty$. The effect of this time change will be to disregard the paths $\tilde{W}_s^{(a)}$ that return to 0 and then survive for a positive amount of time. Setting for every $s \geq 0$,

$$W_s^{(a)} := \tilde{W}_{\eta_s^{(a)}}^{(a)}$$

defines another snake trajectory in \mathcal{S}_0 , which accounts for the labels on the connected component \mathcal{C}_a . We sometimes call $W^{(a)}$ the excursion associated with the excursion debut a .

Let $(a_i)_{i \in I}$ be the (countable) collection of all excursion debuts. For every $i \in I$, we write l_i for the total local time at 0 accumulated by the historical path of a_i (this makes sense because historical paths behave like one-dimensional Brownian paths), and we note that l_i is also the total local time at 0 for the historical path of any point in the component \mathcal{C}_{a_i} .

Theorem 10 ([2]). *There exists a σ -finite measure \mathbb{M}_0 on \mathcal{S}_0 such that, for any nonnegative measurable function Φ on $\mathbb{R}_+ \times \mathcal{S}_0$, we have*

$$\mathbb{N}_0\left(\sum_{i \in I} \Phi(l_i, W^{(a_i)})\right) = \int_0^\infty dl \mathbb{M}_0(\Phi(l, \cdot)).$$

For symmetry reasons, we may write

$$\mathbb{M}_0 = \frac{1}{2}(\mathbb{N}_0^* + \check{\mathbb{N}}_0^*)$$

where \mathbb{N}_0^* is supported on nonnegative snake trajectories, and $\check{\mathbb{N}}_0^*$ is the push forward of \mathbb{N}_0^* under the mapping $\omega \mapsto -\omega$. Under \mathbb{N}_0^* , the paths W_s form a “tree of Brownian paths” starting from 0, which take positive values until the first time when they return to 0 (if they do return to 0) and are stopped at that time if not earlier. See Fig. 7 for a schematic illustration.




Fig. 7. A schematic representation of the paths W_s under \mathbb{N}_0^* . The quantity \mathcal{L}_0^* measures the “number” of circled points corresponding to returns of certain paths W_s to 0

Theorem 10 provides a first-moment formula for the collection of excursions $(W^{(a_i)})_{i \in I}$, but, in contrast with the classical excursion theory, this result does not say anything about the independence of these excursions. To discuss independence properties, we first need to introduce the “boundary size” of an excursion, which roughly speaking measures the quantity of paths W_s that return to 0.

Proposition 11. *The limit*

$$\mathcal{L}_0^* := \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon^2} \int_0^\sigma \mathbf{1}_{\{0 < |\widehat{W}_s| < \varepsilon\}} ds$$

exists \mathbb{M}_0 a.e.

Using scaling arguments, it is not hard to define the conditional probability measures $\mathbb{M}_0(\cdot | \mathcal{L}_0^* = z)$ for every $z > 0$.

In order to state the main result of this section, we still need to introduce a process $(\Lambda_r)_{r>0}$ defined under the excursion measure \mathbb{N}_0 , such that, for every $r > 0$, Λ_r “counts the number” of paths W_s that accumulate a total local time r at 0 and are stopped when they have accumulated that amount of local time. The precise definition of Λ_r fits in the general framework of exit measures as presented in [40], Chapter V, but can also be given via the following approximation:

$$\Lambda_r := \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \int_0^\sigma \mathbf{1}_{\{\chi_r(W_s) < \xi_s < \chi_r(W_s) + \varepsilon\}} ds, \quad \mathbb{N}_0 \text{ a.e.},$$

where $\chi_r(W_s) = \inf\{t \geq 0 : L_t^0(W_s) > r\}$, if $(L_t^0(W_s))_{0 \leq t \leq \xi_s}$ denotes the local time at 0 of the path W_s .

Thanks to the special Markov property of the Brownian snake (see the appendix of [44]), one can prove that the process $(\Lambda_r)_{r>0}$ is Markovian under \mathbb{N}_0 (this makes sense even though \mathbb{N}_0 is an infinite measure because $\mathbb{N}_0(\Lambda_r \neq 0) < \infty$ for every $r > 0$) with the transition kernels of the continuous-state branching process with branching mechanism $\psi(u) = \sqrt{8/3} u^{3/2}$ (see e.g. [40], Section II.1 for the definition and basic facts about continuous-state branching processes). In particular, $(\Lambda_r)_{r>0}$ has a càdlàg modification with only positive jumps, which we consider in the next statement.

Recall that l_i denotes the total local time at 0 accumulated by the historical path of a_i .

Theorem 12 ([2]). *The numbers l_i , $i \in I$, are exactly the jump times of the process $(\Lambda_r)_{r>0}$. Furthermore, conditionally on the process $(\Lambda_r)_{r>0}$, the excursions $W^{(a_i)}$, $i \in I$, are independent and, for every $j \in I$, the conditional distribution of $W^{(a_j)}$ is $\mathbb{M}_0(\cdot | \mathcal{L}_0^* = \Delta \Lambda_{l_j})$.*

In particular, the boundary size of the excursion $W^{(a_j)}$ is $\Delta \Lambda_{l_j}$.

In the applications developed below, we will be interested mainly in positive excursions and in the measure \mathbb{N}_0^* , which we call the positive Brownian snake excursion measure. As in the case of \mathbb{M}_0 we can define the conditional probability measures

$$\mathbb{N}_0^{*,z} := \mathbb{N}_0^*(\cdot | \mathcal{L}_0^* = z),$$

for every $z > 0$.

Interestingly, a number of explicit distributions can be computed under \mathbb{N}_0^* . In particular the joint distribution of the pair $(\mathcal{L}_0^*, \sigma)$ (boundary size and volume) under \mathbb{N}_0^* has a density given by

$$f(z, s) = \frac{\sqrt{3}}{2\pi} \sqrt{z} s^{-5/2} \exp\left(-\frac{z^2}{2s}\right).$$

Consequently, for every fixed $z > 0$, the density of σ under $\mathbb{N}_0^{*,z}$ is

$$g_z(s) = \frac{1}{\sqrt{2\pi}} z^3 s^{-5/2} \exp\left(-\frac{z^2}{2s}\right).$$

The latter density also appears as the density of the asymptotic distribution of the rescaled volume (number of faces) of a Boltzmann quadrangulation with perimeter $2k$, when $k \rightarrow \infty$. This will be explained by the results of the next section.

In the classical setting of excursions away from 0 for a standard linear Brownian motion starting from 0, it is well known that the process can be reconstructed by concatenating the different excursions (some care is required since there are infinitely many excursions on any interval $[0, t]$, $t > 0$). In our setting of a tree-indexed process, things are more complicated since excursions are no longer ordered linearly, but have a certain genealogical structure induced by the genealogy of their debuts: In the example of Fig. 6, the excursion \mathcal{C}_1 is an ancestor of both \mathcal{C}_3 and \mathcal{C}_4 , but \mathcal{C}_3 is not ancestor of \mathcal{C}_4 . Still this genealogical structure can be described in the following way.

For every $a, a' \in \mathcal{T}_\xi$, we let $\delta(a, a')$ be the total local time at 0 accumulated by the process Z along the line segment $\llbracket a, a' \rrbracket$ of the tree \mathcal{T}_ξ . This makes sense since we know that Z evolves like a linear Brownian motion along any segment of the tree. Then $\delta(\cdot, \cdot)$ is a pseudo-metric on \mathcal{T}_ξ , and we can define the associated equivalence relation by setting $a \simeq a'$ if and only if $\delta(a, a') = 0$. Obviously $a \simeq a'$ holds if a and a' belong to the same connected component of $\{b \in \mathcal{T}_\xi : Z_b \neq 0\}$ (because then Z does not vanish on $\llbracket a, a' \rrbracket$). The quotient space \mathcal{T}_ξ / \simeq can thus be seen as obtained from \mathcal{T}_ξ by gluing each excursion component into a single point. It turns out [44] that this quotient space has a remarkable probabilistic structure.

Theorem 13 ([44]). *Under \mathbb{N}_0 , the quotient space \mathcal{T}_ξ / \simeq equipped with the distance induced by δ is a stable Lévy tree with index $3/2$.*

We refer to [29] for the definition and main properties of Lévy trees (note that these trees are typically defined under an infinite measure). See Fig. 8 for a simulation. The stable Lévy tree appearing in the theorem has infinitely many points of infinite multiplicity, and to each such point one can assign a “mass” corresponding informally to the degree of the point in the tree. Then one can

check that points of infinite multiplicity of the tree \mathcal{T}_ξ / \simeq are in one-to-one correspondence with excursions of $(Z_a)_{a \in \mathcal{T}_\xi}$, and that the mass of every point of infinite multiplicity coincides with the boundary size of the associated excursion.




Fig. 8. A simulation of the stable tree with index $3/2$ (simulation: I. Kortchemski)

The previous lines then suggest the following possible “reconstruction” method (which we will not attempt to make rigorous here). Starting from a stable Lévy tree with index $3/2$, associate independently with each point a of infinite multiplicity a snake trajectory W^a distributed according to $\mathbb{M}_0(\cdot | \mathcal{L}_0^* = m_a)$, where m_a is the mass of a , then “glue” at the location of a the genealogical tree of the snake trajectory W^a , with the corresponding labels inherited from W^a . The resulting random \mathbb{R} -tree equipped with labels should be the Brownian tree equipped with Brownian labels.

8. Constructing Brownian disks from the positive Brownian snake excursion measure

The results of this section are taken from [45]. The first naive idea to construct a free Brownian disk is to imitate the construction of Sect. 3.3, replacing the measure \mathbb{N}_0 by \mathbb{N}_0^* . This does not give the desired result, but yields another object of interest, namely the (free) Brownian disk with glued boundary.

Recall our notation (\mathbb{D}, D^∂) for the free Brownian disk with perimeter 1, and $\partial\mathbb{D}$ for the boundary of \mathbb{D} . We define a pseudo-metric on \mathbb{D} by setting, for every $x, y \in \mathbb{D}$,

$$D^\dagger(x, y) = \min\{D^\partial(x, y), D^\partial(x, \partial\mathbb{D}) + D^\partial(y, \partial\mathbb{D})\}.$$

Clearly $D^\dagger(x, y) = 0$ if and only if $x = y$, or both x and y belong to $\partial\mathbb{D}$. Write \mathbb{D}^\dagger for the set obtained from \mathbb{D} by identifying all points of the boundary $\partial\mathbb{D}$ to a single point. Then D^\dagger induces a metric on \mathbb{D}^\dagger , which we still denote by D^\dagger . The compact metric space $(\mathbb{D}^\dagger, D^\dagger)$ is called the free Brownian disk with perimeter 1 and glued boundary. The case of a perimeter equal to z is treated analogously.

Proposition 14. *The random metric space (\mathcal{M}, D) defined via the construction of Sect. 3.3 from a snake trajectory ω distributed according to $\mathbb{N}_0^{*,z}$ is a free Brownian disk with perimeter z and glued boundary.*

The problem is then to recover the free Brownian disk from the same object with glued boundary. This can indeed be achieved by a slight modification of the construction of Sect. 3.3.

From now on, we argue under the measure $\mathbb{N}_0^{*,z}(d\omega)$ for some fixed $z > 0$. Recalling that \mathcal{T}_ξ denotes the genealogical tree of the snake trajectory ω , we use the same notation Z_a for the “labels” on \mathcal{T}_ξ ($Z_a = \widehat{W}_s$ if $a = p_\xi(s)$). In contrast with the case of \mathbb{N}_0 , labels are now nonnegative reals, and we define the “boundary” $\partial\mathcal{T}_\xi$ by

$$\partial\mathcal{T}_\xi := \{a \in \mathcal{T}_\xi : Z_a = 0\}.$$

Recalling the definition of D° in (1), we set, for every $a, b \in \mathcal{T}_\xi \setminus \partial\mathcal{T}_\xi$,

$$\Delta^\circ(a, b) = \begin{cases} D^\circ(a, b) & \text{if } \max(\min_{c \in [a, b]} Z_c, \min_{c \in [b, a]} Z_c) > 0, \\ \infty & \text{otherwise.} \end{cases}$$

Roughly speaking, the condition in the first line of the last display means that we can go from a to b “around” the tree \mathcal{T}_ξ without visiting a vertex of $\partial\mathcal{T}_\xi$. We then define $\Delta(a, b)$ for every $a, b \in \mathcal{T}_\xi \setminus \partial\mathcal{T}_\xi$ by the exact analog of formula (3):

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

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 $\mathcal{T}_\xi \setminus \partial\mathcal{T}_\xi$ such that $a_0 = a$ and $a_k = b$. One easily verifies that the mapping $(a, b) \mapsto \Delta(a, b)$ takes finite values and is continuous on $(\mathcal{T}_\xi \setminus \partial\mathcal{T}_\xi) \times (\mathcal{T}_\xi \setminus \partial\mathcal{T}_\xi)$.

Theorem 15. *With probability one under $\mathbb{N}_0^{*,z}$, the function $(a, b) \mapsto \Delta(a, b)$ has a continuous extension to $\mathcal{T}_\xi \times \mathcal{T}_\xi$, which is a pseudo-metric on \mathcal{T}_ξ . We let Θ stand for the associated quotient space, and we equip Θ with the induced metric, which is still denoted by $\Delta(a, b)$. Then, the random metric space (Θ, Δ) is a free Brownian disk with perimeter z under $\mathbb{N}_0^{*,z}$, and its boundary $\partial\Theta$ is the image of $\partial\mathcal{T}_\xi$ under the canonical projection from \mathcal{T}_ξ onto Θ . Furthermore, if $x \in \Theta$ is the image of $a \in \mathcal{T}_\xi$ under the canonical projection, we have*

$$\Delta(x, \partial\Theta) = Z_a.$$

We note that we can define a volume measure $\mathbf{V}(dx)$ on Θ as the image of the volume measure on \mathcal{T}_ξ under the canonical projection. In particular the total mass of \mathbf{V} is $\mathbf{V}(\Theta) = \sigma$ (recall our notation σ for the duration of the snake trajectory ω , which is also the total mass of the volume measure on \mathcal{T}_ξ). Hence we may define the Brownian disk with perimeter z and volume v as the random metric space (Θ, Δ) under the conditional probability measure $\mathbb{N}_0^{*,z}(\cdot|\sigma = v)$. This is consistent with the construction of [14],[16] using scaling limits of quadrangulations with a boundary with fixed perimeter and volume.

A nice feature of the construction of Theorem 15 (in contrast with the previous constructions in [14],[16]) is the fact that labels Z_a now correspond to distances from the boundary. This also makes it possible to construct a natural “length measure” on the boundary. The following proposition is closely related to the approximation of \mathcal{L}_0^* in Proposition 11.

Proposition 16. *Almost surely under $\mathbb{N}_0^{*,z}$, there exists a finite measure ν on $\partial\Theta$ with total mass z , such that, for every bounded continuous function φ on Θ ,*

$$\langle \nu, \varphi \rangle = \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon^2} \int_{\Theta} \mathbf{V}(dx) \varphi(x) \mathbf{1}_{\{\Delta(x, \partial\Theta) < \varepsilon\}}.$$

We will now exhibit certain particular subsets of the Brownian sphere that are Brownian disks. So we now argue under the measure $\mathbb{N}_0^{(1)}(d\omega)$ and consider the metric space (\mathcal{M}, D) constructed in Sect. 3.3. Recall from the end of this section that \mathcal{M} has a distinguished point x_* such that distances from x_* exactly correspond to the labels Z_x up to a shift (see (5) above). We will discuss properties of the connected components of the complement of balls centered at x_* . At this point, we should mention that the point x_* does not play a special role, and that the re-rooting invariance properties of the Brownian sphere [42], Sect. 8 show that the same properties hold if x_* is replaced by a point chosen according to the volume measure on the Brownian sphere. We recall that this volume measure, which is denoted by $\mathbf{v}(dx)$, is the push forward of the volume measure on \mathcal{T}_ξ , and that \mathbf{v} is a probability measure under $\mathbb{N}_0^{(1)}(d\omega)$.

We note that the Brownian sphere is a length space (as a Gromov–Hausdorff limit of length spaces) and that, if O is an open subset of \mathcal{M} , we can define an

intrinsic metric D_{intr}^O on O by declaring that $D_{\text{intr}}^O(x, y)$ is the minimal length of a continuous path connecting x to y in O (see [20], Chap. 2).

For every $z > 0$ and $v > 0$, we let $\mathbb{F}_{z,v}$ be the distribution of the Brownian disk with perimeter z and volume v . The following statement can be found in [45], Theorem 3 (see also [38] for a related work).

Theorem 17. *Let $r > 0$ and let $B(x_*, r)$ stand for the closed ball of radius r centered at x_* in (\mathcal{M}, D) . Then, $\mathbb{N}_0^{(1)}$ a.s. for every connected component \mathbf{C} of $\mathcal{M} \setminus B(x_*, r)$, the limit*

$$|\partial\mathbf{C}| := \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon^2} \int_{\mathbf{C}} \mathbf{v}(dx) \mathbf{1}_{\{D(x, \partial\mathbf{C}) < \varepsilon\}} \tag{8}$$

exists and is called the boundary size of \mathbf{C} . On the event $\{\mathcal{M} \setminus B(x_, r) \neq \emptyset\}$, write $\mathbf{C}^{r,1}, \mathbf{C}^{r,2}, \dots$, for the connected components of $\mathcal{M} \setminus B(x_*, r)$ ranked in decreasing order of their boundary sizes. Let $D_{\text{intr}}^{r,j}$ be the intrinsic distance on $\mathbf{C}^{r,j}$. Then, $\mathbb{N}_0^{(1)}$ a.s. on the event $\{\mathcal{M} \setminus B(x_*, r) \neq \emptyset\}$, for every $j = 1, 2, \dots$, the metric $D_{\text{intr}}^{r,j}$ has a continuous extension to the closure $\overline{\mathbf{C}}^{r,j}$ of $\mathbf{C}^{r,j}$, and this extension is a metric on $\overline{\mathbf{C}}^{r,j}$. Furthermore, under $\mathbb{N}_0^{(1)}(\cdot | \mathcal{M} \setminus B(x_*, r) \neq \emptyset)$ and conditionally on the sequence*

$$(|\partial\mathbf{C}^{r,1}|, \mathbf{v}(\mathbf{C}^{r,1})), (|\partial\mathbf{C}^{r,2}|, \mathbf{v}(\mathbf{C}^{r,2})), \dots,$$

the metric spaces $(\overline{\mathbf{C}}^{r,j}, D_{\text{intr}}^{r,j})$, $j = 1, 2, \dots$, are independent Brownian disks with respective distributions $\mathbb{F}_{|\partial\mathbf{C}^{r,j}|, \mathbf{v}(\mathbf{C}^{r,j})}$, $j = 1, 2, \dots$

Let us briefly explain why Theorem 17 is related to the excursion theory developed in Sect. 7. The key point is the fact that distances from x_* are given (up to the shift by $-Z_*$) by the labels Z_x . Assuming that $r > -Z_*$ for simplicity, it is then not too hard to verify that connected components of the complement of $B(x_*, r)$ correspond—via the construction presented in Sect. 3—to excursions of Brownian motion indexed by the Brownian tree above the (random) level $r + Z_*$. The distribution of these excursions can be analysed thanks to Theorem 12 and we also use Theorem 15 to relate the positive Brownian snake excursion measure to the law of Brownian disks. There are however two significant technical difficulties, because on one hand we have to deal with excursions above a *random* level, instead of level 0 in Theorem 12, and on the other hand, we argue under $\mathbb{N}_0^{(1)}$ instead of \mathbb{N}_0 in Sect. 7.

Informally, Theorem 17 says that connected components of the complement of a ball centered at a “typical point” in the Brownian sphere are independent Brownian disks conditionally on their boundary sizes and volumes. A similar result [45], Theorem 18 holds for the connected components of the complement of the Brownian net, which is a particular subset of the free Brownian

sphere playing an important role in the axiomatic characterization of Miller and Sheffield [57]. At this point, we mention that we could have stated a version of Theorem 17 for the free Brownian sphere, which is nicer in the sense that we do not need to condition on the volumes: We get that the connected components of the complement of a ball centered at x_* are independent free Brownian disks conditionally on their perimeters. In the next section, we discuss a similar statement for the free Brownian disk, where distances from x_* are replaced by distances from the boundary.

9. Slicing Brownian disks at heights

In this section, which is based on a joint work with Armand Riera [50], we consider the random metric space (Θ, Δ) defined in Theorem 15, which is a free Brownian disk with perimeter z under the probability measure $\mathbb{N}_0^{*,z}$. For every $x \in \Theta$, define the *height* of x by

$$H(x) = \Delta(x, \partial\mathbb{D}).$$

We also consider the maximal height

$$H^* = \max_{x \in \Theta} H(x).$$

Recall the notation $\mathbf{V}(dx)$ for the volume measure on Θ .

Theorem 18. *Let $r > 0$. Then, $\mathbb{N}_0^{*,z}$ a.s., for every connected component \mathcal{C} of $\{x \in \Theta : H(x) > r\}$, the limit*

$$|\partial\mathcal{C}| = \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon^2} \int_{\mathcal{C}} \mathbf{V}(dx) \mathbf{1}_{\{H(x) < r + \varepsilon\}}$$

exists and is called the perimeter of \mathcal{C} . On the event $\{H^ > r\}$, let $\mathcal{C}^{r,1}, \mathcal{C}^{r,2}, \dots$, be the connected components of $\{x \in \Theta : H(x) > r\}$ ranked in decreasing order of their perimeters. Then, a.s. on the event $\{H^* > r\}$, for every $j = 1, 2, \dots$, the intrinsic metric on $\mathcal{C}^{r,j}$ has a continuous extension to the closure $\overline{\mathcal{C}^{r,j}}$ of $\mathcal{C}^{r,j}$, which is a metric on $\overline{\mathcal{C}^{r,j}}$, and conditionally on the perimeters $|\partial\mathcal{C}^{1,r}|, |\partial\mathcal{C}^{2,r}|, \dots$, the resulting metric spaces $\overline{\mathcal{C}^{r,1}}, \overline{\mathcal{C}^{r,2}}, \dots$, are independent free Brownian disks.*

As explained for Theorem 17 at the end of the previous section, Theorem 18 can be derived from the excursion theory developed in Sect. 7. The difficulty now comes from the fact that we must argue under $\mathbb{N}_0^{*,z}$ instead of \mathbb{N}_0 in Sect. 7.

With the notation of Theorem 18, an obvious question is to describe the distribution of the process

$$\mathbf{X}(r) = (|\partial\mathcal{C}^{1,r}|, |\partial\mathcal{C}^{2,r}|, \dots)$$

giving for every $r > 0$ the perimeters of all connected components of $\{x \in \Theta : H(x) > r\}$ (by convention $\mathbf{X}(r) = (0, 0, \dots)$ if $H^* \leq r$). We also take $\mathbf{X}(0) = (z, 0, 0, \dots)$ and view $(\mathbf{X}(r))_{r \geq 0}$ as a random process taking values in the space of nonincreasing sequences of nonnegative real numbers. Theorem 18 then suggests that this process enjoys properties similar to those of the growth-fragmentation processes that have been studied recently by several authors. In fact, Bertoin, Curien and Kortchemski [13] (see also [12] for extensions) have considered a process analogous to \mathbf{X} for triangulations with a boundary and showed that the scaling limit of this process (when the boundary size tends to infinity) is a well-identified growth-fragmentation process. Still it does not seem easy to apply the results of [13] in order to identify the distribution of the process \mathbf{X} , but the excursion theory of Sect. 7 can be used instead to compute this distribution.

Before stating our last result, we need to recall a few basic facts about growth-fragmentation processes (see [11] for more details). The starting ingredient is a positive self-similar Markov process $(Y_t)_{t \geq 0}$ with only negative jumps. Suppose that $Y_0 = z$, and view $(Y_t)_{t \geq 0}$ as the evolution in time of the mass of an initial particle also called the Eve particle. At each time t where the process Y has a jump, we consider that a new particle with mass $-\Delta Y_t$ (a child of the Eve particle) is born, and the mass of this new particle evolves (from time t) again according to the law of the process Y , but independently of the evolution of the mass of the Eve particle. Then each child of the Eve particle has children at discontinuity times of its mass process, and so on. Under suitable assumptions (see [11]), we can make sense of the process $(\mathbf{Y}(t))_{t \geq 0}$ giving for every time t the sequence (in decreasing order) of masses of all particles alive at that time. The process \mathbf{Y} is Markovian and is called the growth-fragmentation process with Eve particle process Y .

Theorem 19. *Under $\mathbb{N}_0^{*,z}$, the process $(\mathbf{X}(r))_{r \geq 0}$ is a growth-fragmentation process, which is constructed from an Eve particle process X whose distribution starting from 1 is specified as follows:*

$$X_t = \exp(\xi_{\tau(t)}),$$

where

$$\tau(t) = \inf \left\{ u \geq 0 : \int_0^u e^{\xi_s/2} ds > t \right\}$$

and ξ is the spectrally negative Lévy process such that, for every $q > 0$, $\mathbb{E}[\exp(q \xi_t)] = \exp(t\psi(q))$, with

$$\psi(q) = \sqrt{\frac{3}{2\pi}} \left(-\frac{8}{3}q + \int_{1/2}^1 (x^q - 1 + q(1-x)) (x(1-x))^{-5/2} dx \right). \quad (9)$$

Remark. The process ξ drifts to $-\infty$ and the event $\{\tau(t) = \infty\}$ occurs with positive probability if $t > 0$: On this event, we of course make the convention that $\exp(\xi_\infty) = 0$.

The expression of the process X in terms of the Lévy process ξ is a special case of the classical Lamperti representation of a positive self-similar Markov process (here with index $1/2$) in terms of a spectrally negative Lévy process. The formula for ψ is the same as formula (1) in [13], except for the (unimportant) multiplicative constant $\sqrt{3/2\pi}$. This should not come as a surprise in view of preceding comments.

Although we have chosen to state them as properties of the free Brownian disk, Theorems 18 and 19 are really results about the tree-indexed Brownian motion $(Z_a)_{a \in \mathcal{T}_\xi}$ under $\mathbb{N}_0^{*,z}$. In particular, Theorem 19 relies on the identification of the distribution of the process giving, for each $r \geq 0$, the sequence of boundary sizes of all excursions above level r of the process $(Z_a)_{a \in \mathcal{T}_\xi}$ under $\mathbb{N}_0^{*,z}$. There is a striking analogy with the fragmentation process occurring when cutting the CRT at a fixed height: Precisely, it is shown in [10] that the sequence of volumes of the connected components of the complement of the ball of radius r centered at the root in the CRT is a self-similar fragmentation process whose dislocation measure has the form $(2\pi)^{-1/2}(x(1-x))^{-3/2} dx$, to be compared with the measure $(x(1-x))^{-5/2} dx$ appearing in formula (9).

As a consequence of Theorem 19 and known asymptotics [12], Corollary 4.5 for the distribution of the extinction time of a growth-fragmentation process, we derive the following corollary about the tail of the distribution of the maximal height in a Brownian disk.

Corollary 20. *There exist positive constants c_1 and c_2 such that, for every $r \geq 1$,*

$$c_1 r^{-6} \leq \mathbb{N}_0^{*,z}(H^* > r) \leq c_2 r^{-6}.$$

Acknowledgements. It is a pleasure to thank the organizers of the Takagi Lectures for giving me the opportunity to discuss the present work at this prestigious meeting.

References

- [1] C. Abraham, Rescaled bipartite planar maps converge to the Brownian map, *Ann. Inst. Henri Poincaré Probab. Stat.*, **52** (2016), 575–595.

- [2] C. Abraham and J.-F. Le Gall, Excursion theory for Brownian motion indexed by the Brownian tree, *J. Eur. Math. Soc. (JEMS)*, **20** (2018), 2951–3016.
- [3] L. Addario-Berry and M. Albenque, The scaling limit of random simple triangulations and random simple quadrangulations, *Ann. Probab.*, **45** (2017), 2767–2825.
- [4] D. Aldous, The continuum random tree. I, *Ann. Probab.*, **19** (1991), 1–28.
- [5] D. Aldous, The continuum random tree. III, *Ann. Probab.*, **21** (1993), 248–289.
- [6] J. Ambjørn, B. Durhuus and T. Jonsson, *Quantum Geometry. A Statistical Field Theory Approach*, Cambridge Monogr. Math. Phys., Cambridge Univ. Press, Cambridge, 1997.
- [7] O. Angel, Growth and percolation on the uniform infinite planar triangulation, *Geom. Funct. Anal.*, **13** (2003), 935–974.
- [8] O. Angel and O. Schramm, Uniform infinite planar triangulations, *Comm. Math. Phys.*, **241** (2003), 191–213.
- [9] E. Baur, G. Miermont and G. Ray, Classification of scaling limits of quadrangulations with a boundary, *Ann. Probab.*, to appear; preprint, arXiv:1608.01129.
- [10] J. Bertoin, Self-similar fragmentations, *Ann. Inst. H. Poincaré Probab. Statist.*, **38** (2002), 319–340.
- [11] J. Bertoin, Markovian growth-fragmentation processes, *Bernoulli*, **23** (2017), 1082–1101.
- [12] J. Bertoin, T. Budd, N. Curien and I. Kortchemski, Martingales in self-similar growth-fragmentations and their connections with random planar maps, *Probab. Theory Related Fields*, to appear; preprint, arXiv:1605.00581.
- [13] J. Bertoin, N. Curien and I. Kortchemski, Random planar maps and growth-fragmentations, *Ann. Probab.*, **46** (2018), 207–260.
- [14] J. Bettinelli, Scaling limit of random planar quadrangulations with a boundary, *Ann. Inst. Henri Poincaré Probab. Stat.*, **51** (2015), 432–477.
- [15] J. Bettinelli, E. Jacob and G. Miermont, The scaling limit of uniform random plane maps, via the Ambjørn–Budd bijection, *Electron. J. Probab.*, **19** (2014), no. 74.
- [16] J. Bettinelli and G. Miermont, Compact Brownian surfaces I. Brownian disks, *Probab. Theory Related Fields*, **167** (2017), 555–614.
- [17] J. Bouttier, P. Di Francesco and E. Guitter, Planar maps as labeled mobiles, *Electron. J. Combin.*, **11** (2004), R69.
- [18] J. Bouttier and E. Guitter, The three-point function of planar quadrangulations, *J. Stat. Mech. Theory Exp.*, **2008** (2008), P07020.
- [19] T. Budzinski, The hyperbolic Brownian plane, *Probab. Theory Related Fields*, **171** (2018), 503–541.
- [20] D. Burago, Y. Burago and S. Ivanov, *A Course in Metric Geometry*, Grad. Stud. Math., **33**, Amer. Math. Soc., Providence, RI, 2001.
- [21] A. Caraceni and N. Curien, Geometry of the uniform infinite half-planar quadrangulation, *Random Structures Algorithms*, **52** (2018), 454–494.
- [22] P. Chassaing and B. Durhuus, Local limit of labeled trees and expected volume growth in a random quadrangulation, *Ann. Probab.*, **34** (2006), 879–917.
- [23] P. Chassaing and G. Schaeffer, Random planar lattices and integrated superBrownian excursion, *Probab. Theory Related Fields*, **128** (2004), 161–212.
- [24] N. Curien and J.-F. Le Gall, The Brownian plane, *J. Theoret. Probab.*, **27** (2014), 1249–1291.
- [25] N. Curien and J.-F. Le Gall, The hull process of the Brownian plane, *Probab. Theory Related Fields*, **166** (2016), 187–231.
- [26] N. Curien and J.-F. Le Gall, First-passage percolation and local modifications of distances in random triangulations, *Ann. Sci. Éc. Norm. Supér. (4)*, to appear; preprint, arXiv:1511.04264.
- [27] N. Curien, L. Ménard and G. Miermont, A view from infinity of the uniform infinite planar quadrangulation, *ALEA Lat. Am. J. Probab. Math. Stat.*, **10** (2013), 45–88.

- [28] N. Curien and G. Miermont, Uniform infinite planar quadrangulations with a boundary, *Random Structures Algorithms*, **47** (2015), 30–58.
- [29] T. Duquesne and J.-F. Le Gall, Probabilistic and fractal aspects of Lévy trees, *Probab. Theory Related Fields*, **131** (2005), 553–603.
- [30] A. Greven, P. Pfaffelhuber and A. Winter, Convergence in distribution of random metric measure spaces (Λ -coalescent measure trees), *Probab. Theory Related Fields*, **145** (2009), 285–322.
- [31] O. Gurel-Gurevich and A. Nachmias, Recurrence of planar graph limits, *Ann. of Math. (2)*, **177** (2013), 761–781.
- [32] E. Gwynne and J. Miller, Scaling limit of the uniform infinite half-plane quadrangulation in the Gromov–Hausdorff–Prokhorov-uniform topology, *Electron. J. Probab.*, **22** (2017), no. 84.
- [33] E. Gwynne and J. Miller, Convergence of the free Boltzmann quadrangulation with simple boundary to the Brownian disk, *Ann. Inst. Henri Poincaré Probab. Stat.*, **55** (2019), 551–589.
- [34] E. Gwynne and J. Miller, Convergence of the self-avoiding walk on random quadrangulations to $\text{SLE}_{8/3}$ on $\sqrt{8/3}$ -Liouville quantum gravity, preprint, arXiv:1608.00956.
- [35] E. Gwynne and J. Miller, Convergence of percolation on uniform quadrangulations with boundary to SLE_6 on $\sqrt{8/3}$ -Liouville quantum gravity, preprint, arXiv:1701.05175.
- [36] E. Gwynne, J. Miller and S. Sheffield, The Tutte embedding of the mated-CRT map converges to Liouville quantum gravity, preprint, arXiv:1705.11161.
- [37] K. Itô, Poisson point processes attached to Markov processes, In: *Proc. Sixth Berkeley Symp. Math. Stat. Prob.*, **3**, Univ. California Press, Berkeley, CA, 1970, pp. 225–239.
- [38] E. Jacob and G. Miermont, The Brownian map contains Brownian disks, in preparation.
- [39] M. Krikun, Local structure of random quadrangulations, preprint, arXiv:math/0512304.
- [40] J.-F. Le Gall, *Spatial Branching Processes, Random Snakes and Partial Differential Equations*, Lectures Math. ETH Zürich, Birkhäuser Verlag, Basel, 1999.
- [41] J.-F. Le Gall, The topological structure of scaling limits of large planar maps, *Invent. Math.*, **169** (2007), 621–670.
- [42] J.-F. Le Gall, Geodesics in large planar maps and in the Brownian map, *Acta Math.*, **205** (2010), 287–360.
- [43] J.-F. Le Gall, Uniqueness and universality of the Brownian map, *Ann. Probab.*, **41** (2013), 2880–2960.
- [44] J.-F. Le Gall, Subordination of trees and the Brownian map, *Probab. Theory Related Fields*, **171** (2018), 819–864.
- [45] J.-F. Le Gall, Brownian disks and the Brownian snake, *Ann. Inst. Henri Poincaré Probab. Stat.*, **55** (2019), 237–313.
- [46] J.-F. Le Gall and T. Lehericy, Separating cycles and isoperimetric inequalities in the uniform infinite planar quadrangulation, *Ann. Probab.*, to appear; preprint, arXiv:1710.02990.
- [47] J.-F. Le Gall and G. Miermont, Scaling limits of random planar maps with large faces, *Ann. Probab.*, **39** (2011), 1–69.
- [48] J.-F. Le Gall and G. Miermont, Scaling limits of random trees and planar maps, In: *Probability and Statistical Physics in Two and More Dimensions*, Clay Math. Proc., **15**, Amer. Math. Soc., Providence, RI, 2012, pp. 155–211.
- [49] J.-F. Le Gall and F. Paulin, Scaling limits of bipartite planar maps are homeomorphic to the 2-sphere, *Geom. Funct. Anal.*, **18** (2008), 893–918.
- [50] J.-F. Le Gall and A. Riera, Growth-fragmentation processes in Brownian motion indexed by the Brownian tree, preprint, arXiv:1811.02825.
- [51] J.-F. Marckert and A. Mokkadem, Limit of normalized quadrangulations: The Brownian map, *Ann. Probab.*, **34** (2006), 2144–2202.

-
- [52] C. Marzouk, Scaling limits of random bipartite planar maps with a prescribed degree sequence, *Random Structures Algorithms*, **53** (2018), 448–503.
- [53] L. Ménard, The two uniform infinite quadrangulations of the plane have the same law, *Ann. Inst. Henri Poincaré Probab. Stat.*, **46** (2010), 190–208.
- [54] G. Miermont, The Brownian map is the scaling limit of uniform random plane quadrangulations, *Acta Math.*, **210** (2013), 319–401.
- [55] G. Miermont, Aspects of random maps, Lecture notes of the 2014 Saint-Flour Probability Summer School, available at <http://perso.ens-lyon.fr/gregory.miermont/coursSaint-Flour.pdf>.
- [56] G. Miermont, personal communication.
- [57] J. Miller and S. Sheffield, An axiomatic characterization of the Brownian map, preprint, arXiv:1506.03806.
- [58] J. Miller and S. Sheffield, Liouville quantum gravity and the Brownian map I: The QLE(8/3,0) metric, preprint, arXiv:1507.00719.
- [59] J. Miller and S. Sheffield, Liouville quantum gravity and the Brownian map II: Geodesics and continuity of the embedding, preprint, arXiv:1605.03563.
- [60] J. Miller and S. Sheffield, Liouville quantum gravity and the Brownian map III: the conformal structure is determined, preprint, arXiv:1608.05391.
- [61] D. Revuz and M. Yor, *Continuous Martingales and Brownian Motion*, Grundlehren Math. Wiss., **293**, Springer-Verlag, 1991.
- [62] O. Schramm, Conformally invariant scaling limits: an overview and a collection of problems, In: *Proceedings of the International Congress of Mathematicians. Vol. I*, Eur. Math. Soc., Zürich, 2007, pp. 513–543.
- [63] R. Stephenson, Local convergence of large critical multi-type Galton–Watson trees and applications to random maps, *J. Theoret. Probab.*, **31** (2018), 159–205.
- [64] Y. Watabiki, Construction of non-critical string field theory by transfer matrix formalism in dynamical triangulation, *Nuclear Phys. B*, **441** (1995), 119–163.