

Chapter 1

Introduction

1.1 Context

Random maps, seen as discrete models of random two-dimensional geometries, have generated a sustained interest in the last couple of decades. An important instance of this line of research are the results by Le Gall [67] and the second author [80], showing that a uniform random quadrangulation of the sphere with n faces, seen as a random finite metric space by endowing its vertex set with the usual graph metric renormalized by $n^{-1/4}$, converges in distribution toward the so-called *Brownian sphere*, or *Brownian map*. The aim of the present work is to generalize this result to the case of general compact orientable surfaces. Let us start with some elements of context.

Random surfaces as scaling limits of random maps. While the idea that continuum random geometries should be obtained as scaling limits of random maps originates from the physics literature on two-dimensional quantum gravity [44, 59, 88], this question was first approached in the mathematical literature in the pioneering work of Chassaing and Schaeffer [37], who studied the model of uniformly chosen random quadrangulations of the sphere, and found in particular that the proper scaling factor in this case was $n^{-1/4}$. Marckert and Mokkadem [76] then constructed a candidate limiting space today called the *Brownian sphere*, and showed the convergence toward it in another topology than the Gromov–Hausdorff topology. Le Gall [65] later showed that the sequence of rescaled metric spaces associated with uniform random quadrangulations of the sphere was relatively compact. Finally, Le Gall [67] and the second author [80] showed by two independent approaches that the previous sequence converges toward the Brownian sphere.

It is known that the Brownian sphere arises as a universal scaling limit for many models of planar maps *that are uniformly chosen in a certain class, given their face degrees*, and provided that face degrees are typically all of the same order of magnitude; see [1, 3, 4, 16, 24, 41, 67, 77]. See also [72] for models of maps that fall out of this universality class.

The scaling limits of quadrangulations on surfaces that are more general than the sphere were considered by the first author in [21, 22], who showed similar results to the above, but only up to extraction of appropriate subsequences, leaving a gap that amounts to uniquely characterize the limit. This gap was filled in the particular case of the disk topology in our previous work [25]. In particular, we showed that a uniform quadrangulation of genus 0 with one boundary component having n internal

faces and perimeter $2l_n$ weakly converges, once scaled by the factor $n^{-1/4}$ and when $l_n \sim L\sqrt{2n}$, toward a random metric space called the *Brownian disk of perimeter L* . Two alternate constructions of Brownian disks were proposed by Le Gall [68, 70], allowing in particular to show that Brownian disks arise as connected components of the complement of metric balls in the Brownian sphere, conditionally given their areas and boundary lengths. See also [18, 74, 83].

Besides the case of the sphere and the disk, only a few results have been obtained for maps on compact surfaces. Namely, it has been shown that uniform quadrangulations of a given compact surface with a boundary exhibit scaling limits [19, 20, 22], all of the same topology as the considered surface, and geodesics to a uniformly chosen points were studied [22]. More recently, it was shown that uniformly distributed essentially simple toroidal triangulations (that is, triangulations of the torus without contractible loops or double edges forming cycles that are homotopic to 0) also exhibit scaling limits [15], which are believed to be the same as for random quadrangulations. See also [10] for a scaling limit result of Boltzmann random maps with annular topology.

There has also been a growing interest in noncompact versions of these models, especially as they bridge some Brownian surfaces with so-called *uniform infinite random maps*, which are maps with infinitely many faces that first arose in a work by Angel and Schramm [12], as *local limits* of random finite maps. Three main models of noncompact Brownian surfaces have been identified: the *Brownian plane* [40], the *Brownian half-plane* [13, 52], and the *infinite-volume Brownian disk* [13], which can be thought of as noncompact versions of the Brownian sphere and Brownian disks, either with unbounded or bounded boundary. See [75] for a framework unifying those objects. The first two of these models will play an important role in the current work.

This whole line of research crucially depends on strong combinatorial techniques, and in particular on bijective approaches [6, 28, 90] that allow to give very detailed quantitative information on the geodesic paths in random maps and their scaling limits. The present work is no exception. See, for instance, [11, 66, 71, 79, 81] for results related to the structure of geodesics in the Brownian sphere, [69] for a recent survey, and [39] for another approach called *peeling*. We note, however, that, so far, these methods are restricted to models of maps chosen uniformly, conditionally given their face degrees, as alluded to above.

Random surfaces via Liouville quantum gravity metrics. A line of research parallel to the above consists in building the limiting spaces directly as continuum random metrics in planar domains or Riemann surfaces. This approach also finds its roots in the physics theory of Liouville quantum gravity [88]. In the case of Brownian surfaces, this has first been implemented by Miller and Sheffield in a series of works [82–85], where they use a growth model called *quantum Loewner evolution* (QLE) to define a random metric on the plane, whose metric balls are described by QLE,

and whose law as an abstract metric space is equal to that of the Brownian plane. Local variants of the construction allow to define the Brownian sphere in this way. The Miller–Sheffield metric is in fact a special element of a one-parameter family of *Liouville quantum gravity* (LQG) metrics, that have been defined as scaling limits of first-passage percolation models in mollified exponentiated Gaussian free fields landscapes [46, 55]. See [47] for an overview of LQG metrics.

These constructions operate entirely in the continuum, and naturally ask whether canonical embeddings of random maps in the sphere are compatible with the convergence toward the Brownian sphere, in the sense that the metrics induced by the embedding converge to the random metric of Miller–Sheffield. Such a result was recently obtained by Holden and Sun [58] (which is the last piece of a vast research project, described in details in this reference), who showed the joint convergence of the metric and the area measure generated by a uniform plane triangulation embedded via the Cardy–Smirnov embedding in an equilateral triangle. We refer to the overview article [51].

The existence of a canonical conformal structure for Brownian surfaces was also approached in a more direct way by Gwynne, Miller and Sheffield in [56, 57]. Their method, which has been implemented so far for the plane, half-plane, sphere and disk topologies, consists in taking limits of discrete embeddings obtained directly from the continuum limit by considering Poisson–Voronoi tessellations with a finer and finer mesh, and showing that the random walk on the discrete approximation converges to Brownian motion in the plane. In passing, this allows one to define Brownian motion on the Brownian surfaces under consideration.

Random surfaces and conformal field theories. While the definition of LQG metrics applies to any field that “locally looks like” the Gaussian free field, the exact law of the latter is of crucial importance to obtain the exact law of random surfaces that arise as scaling limits of maps, and this law can be obtained from Liouville conformal field theory [88]. Here, rather than dealing with random metrics, one is rather interested in the computation of partition functions defined from the field, and it has been shown recently in a rich body of work – see [45, 50, 61] and references therein – that this theory has a probabilistic interpretation in terms of Gaussian multiplicative chaoses, which are random measures defined in terms of the Gaussian free field. This approach has unveiled fundamental integrability properties for planar Gaussian multiplicative chaoses, which can be used to provide exact distributions for various quantities related to the LQG metrics, hence to the scaling limits of random maps. For instance, in [10], the authors compute the law of the conformal modulus of a Brownian annulus, which is a member of the family of Brownian surfaces described in the present work.

The interplay between these approaches provides a wealth of methods to prove various properties of random surfaces [91], and the geometric properties of the

Brownian surfaces, as well as the other LQG metrics, are the object of intensive current research.

1.2 Generalities and terminology on maps

Surface with a boundary. Recall that a *surface with a boundary* is a nonempty Hausdorff topological space in which every point has an open neighborhood homeomorphic to some open subset of $\mathbb{R} \times \mathbb{R}_{\geq 0}$. Its *boundary* is the set of points having a neighborhood homeomorphic to a neighborhood of $(0, 0)$ in $\mathbb{R} \times \mathbb{R}_{\geq 0}$. When it is nonempty, this set forms a one-dimensional topological manifold. In this work, we will only consider *orientable* compact connected surfaces with a (possibly empty) boundary. By the classification theorem, these are characterized up to homeomorphisms by two nonnegative integers, the genus g and the number b of connected components of the boundary. We denote by $\Sigma_b^{[g]}$ the compact orientable surface of genus g with b boundary components, which is unique up to homeomorphisms. It can be obtained from the connected sum of g tori, or from the sphere in the case $g = 0$, by removing b disjoint open disks whose boundaries are pairwise disjoint circles.

Map. A *map* is a proper cellular embedding of a finite graph, possibly with multiple edges and loops, into a compact connected orientable surface *without* boundary. Here, the word *proper* means that edges can intersect only at vertices, and *cellular* means that the connected components of the complement of the edges, which are called the *faces* of the map, are homeomorphic to two-dimensional open disks. Maps will always be considered up to orientation-preserving homeomorphisms of the surface into which they are embedded. The *genus* of a map is defined as the genus of the surface into which it is embedded; we speak of *plane* maps when the genus is 0. We call *half-edge* an oriented edge in a map. With every half-edge, we may associate in a one-to-one way a *corner*, which is the angular sector lying to its left at the origin of the half-edge. Note that this makes sense because the surfaces we are considering are orientable. We say that a corner, or the corresponding half-edge, is *incident* to a face f if it lies into f . We also say that the face is incident to the corner or the half-edge in this case. The number of half-edges (or equivalently, of corners) incident to a face is called its *degree*.

A map is *rooted* if it comes with a distinguished corner – or, equivalently, a half-edge – called the *root*. Rooting is a very useful notion as it allows to kill the symmetries of a map. In fact, when dealing with nonrooted maps, we will systematically count them by weighting each map \mathbf{m} by a factor $1/\text{Aut}(\mathbf{m})$, where $\text{Aut}(\mathbf{m})$ denotes the number of automorphisms of \mathbf{m} . The latter is also equal to $2|E(\mathbf{m})|/R(\mathbf{m})$, where $E(\mathbf{m})$ is the edge set of \mathbf{m} , and $R(\mathbf{m})$ is the number of distinct rooted maps that can be obtained from the nonrooted map \mathbf{m} . Therefore, with this convention, the weighted

number of nonrooted maps in a given family of maps with a given number e of edges is simply the cardinality of the set of rooted maps from this same family, divided by $2e$.

Map with holes. We will consider maps with pairwise distinct distinguished elements, generically denoted by h_1, h_2, \dots, h_k , that can be either *faces* or *vertices*. These distinguished elements are called the *holes* of the map, a given hole being called either an *external face* or an *external vertex*, depending on its nature. The nondistinguished faces and vertices are called the *internal faces* and *internal vertices*. The degree of a hole, also called its *perimeter*, is defined as 0 for an external vertex or as the degree of the face for an external face. Beware that the boundaries of the external faces are in general neither simple curves, nor pairwise disjoint. As a result, the object obtained by removing them from the surface in which the map is embedded is not necessarily a surface. Note that, however, removing from every external face an open disk whose closure is included in the (open) face results in a surface with a boundary.

Bipartite map. Finally, we say that a map is *bipartite* if its vertex set can be partitioned into two subsets such that no edge links two vertices of the same subset.

Tuples. The many tuples considered in this work will conventionally be denoted by a boldface font letter (possibly with a subscript) and their coordinates with the same letter in a normal font, with the index written as a superscript, as in $\mathbf{x} = (x^1, \dots, x^r)$, for instance. When \mathbf{x} is a tuple of real nonnegative numbers, we set $\|\mathbf{x}\| = \sum_{i=1}^r x^i$. We denote by $\mathbf{x}\mathbf{y}$ the concatenation of \mathbf{x} with \mathbf{y} . Finally, when concatenating with a 1-tuple, we often identify it with its unique coordinate, writing, for instance, $\mathbf{x}0 = (x^1, \dots, x^r, 0)$.

1.3 The Gromov–Hausdorff–Prokhorov topology

In this memoir, a *metric measure space* is a triple $(\mathcal{X}, d_{\mathcal{X}}, \mu_{\mathcal{X}})$, where $(\mathcal{X}, d_{\mathcal{X}})$ is a nonempty *compact* metric space and $\mu_{\mathcal{X}}$ is a finite Borel measure on \mathcal{X} . We say that two metric measure spaces $(\mathcal{X}, d_{\mathcal{X}}, \mu_{\mathcal{X}})$ and $(\mathcal{Y}, d_{\mathcal{Y}}, \mu_{\mathcal{Y}})$ are *isometry-equivalent* if there exists an isometry ϕ from $(\mathcal{X}, d_{\mathcal{X}})$ onto $(\mathcal{Y}, d_{\mathcal{Y}})$ such that $\mu_{\mathcal{Y}} = \phi_*\mu_{\mathcal{X}}$. This defines an equivalence relation on the class of all metric measure spaces. If $(\mathcal{X}, d_{\mathcal{X}}, \mu_{\mathcal{X}})$ and $(\mathcal{Y}, d_{\mathcal{Y}}, \mu_{\mathcal{Y}})$ are two metric measure spaces, the *Gromov–Hausdorff–Prokhorov metric* (*GHP metric* for short) is defined by

$$\begin{aligned} d_{\text{GHP}}((\mathcal{X}, d_{\mathcal{X}}, \mu_{\mathcal{X}}), (\mathcal{Y}, d_{\mathcal{Y}}, \mu_{\mathcal{Y}})) \\ = \inf_{\substack{\phi: \mathcal{X} \rightarrow \mathcal{Z} \\ \psi: \mathcal{Y} \rightarrow \mathcal{Z}}} \{d_{\mathcal{Z}}^{\text{H}}(\phi(\mathcal{X}), \psi(\mathcal{Y})) \vee d_{\mathcal{Z}}^{\text{P}}(\phi_*\mu_{\mathcal{X}}, \psi_*\mu_{\mathcal{Y}})\}, \end{aligned} \quad (1.1)$$

where the infimum is taken over all choices of compact metric spaces $(\mathcal{Z}, d_{\mathcal{Z}})$, and all isometric maps ϕ, ψ from \mathcal{X}, \mathcal{Y} to \mathcal{Z} , where $d_{\mathcal{Z}}^{\text{H}}$ is the *Hausdorff metric* on compact subsets of \mathcal{Z} , and $d_{\mathcal{Z}}^{\text{P}}$ is the *Prokhorov metric* on finite positive measures on \mathcal{Z} , defined as follows. First, for any $\varepsilon > 0$ and any closed subset $A \subseteq \mathcal{Z}$, we denote by

$$A^\varepsilon = \{z \in \mathcal{Z} : \inf_{y \in A} d_{\mathcal{Z}}(z, y) < \varepsilon\}$$

its ε -enlargement. Then, for any compact subsets $A, B \subseteq \mathcal{Z}$,

$$d_{\mathcal{Z}}^{\text{H}}(A, B) = \inf\{\varepsilon > 0 : A \subseteq B^\varepsilon \text{ and } B \subseteq A^\varepsilon\},$$

and, for any finite Borel measures μ, ν on \mathcal{Z} ,

$$d_{\mathcal{Z}}^{\text{P}}(\mu, \nu) = \inf\{\varepsilon > 0 : \text{for all closed } A \subseteq \mathcal{Z}, \mu(A) \leq \nu(A^\varepsilon) + \varepsilon \\ \text{and } \nu(A) \leq \mu(A^\varepsilon) + \varepsilon\}.$$

Equation (1.1) defines a metric on the set \mathbb{M} of isometry-equivalence classes of metric measure spaces, making it a complete and separable metric space. The references [93, Chapter 27] as well as [2, 68] discuss relevant aspects of the GHP topology, with some variations, as the exact definition of the metric may differ from place to place.

More generally, for $\ell, m \geq 0$, we will consider ℓ -marked, m -measured metric spaces of the form $(\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \boldsymbol{\mu}_{\mathcal{X}})$, where

- $(\mathcal{X}, d_{\mathcal{X}})$ is a nonempty compact metric space,
- \mathbf{A} is an ℓ -tuple, called *marking*, of nonempty compact subsets of \mathcal{X} , called *marks*,
- $\boldsymbol{\mu}_{\mathcal{X}}$ is an m -tuple of finite Borel measures on \mathcal{X} .

We often consider marks that are singletons; in this case, we identify the singleton with the point it contains. We define the ℓ -marked, m -measured *Gromov–Hausdorff–Prokhorov metric* (still *GHP metric* for short) on such spaces by

$$d_{\text{GHP}}^{(\ell, m)}((\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \boldsymbol{\mu}_{\mathcal{X}}), (\mathcal{Y}, d_{\mathcal{Y}}, \mathbf{B}, \boldsymbol{\mu}_{\mathcal{Y}})) \\ = \inf_{\substack{\phi: \mathcal{X} \rightarrow \mathcal{Z} \\ \psi: \mathcal{Y} \rightarrow \mathcal{Z}}} \left\{ d_{\mathcal{Z}}^{\text{H}}(\phi(\mathcal{X}), \psi(\mathcal{Y})) \vee \max_{1 \leq i \leq \ell} d_{\mathcal{Z}}^{\text{H}}(\phi(A^i), \psi(B^i)) \right. \\ \left. \vee \max_{1 \leq j \leq m} d_{\mathcal{Z}}^{\text{P}}(\phi_* \mu_{\mathcal{X}}^j, \psi_* \mu_{\mathcal{Y}}^j) \right\},$$

where the infimum is taken over the same family as in (1.1). Again, this defines a complete and separable metric on the set $\mathbb{M}^{(\ell, m)}$ of isometry-equivalence classes of ℓ -marked, m -measured metric spaces, where $(\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \boldsymbol{\mu}_{\mathcal{X}})$ and $(\mathcal{Y}, d_{\mathcal{Y}}, \mathbf{B}, \boldsymbol{\mu}_{\mathcal{Y}})$ are isometry-equivalent if there exists an isometry ϕ from \mathcal{X} onto \mathcal{Y} such that $\phi(A^i) = B^i$ for $1 \leq i \leq \ell$ and $\phi_* \mu_{\mathcal{X}}^j = \mu_{\mathcal{Y}}^j$ for $1 \leq j \leq m$. Note that we have $d_{\text{GHP}}^{(0, 1)} = d_{\text{GHP}}$. Finally,

the space $(\mathbb{M}^{(\ell)}, d_{\text{GH}}^{(\ell)}) = (\mathbb{M}^{(\ell,0)}, d_{\text{GHP}}^{(\ell,0)})$ of ℓ -marked compact metric spaces without measures is the so-called ℓ -marked Gromov–Hausdorff metric (GH metric for short).

As a first example, we will sometimes use in the present work the *point space* $\{\varrho\}$ consisting of a single point, seen as the element $(\{\varrho\}, (\varrho, \dots, \varrho), (0, \dots, 0)) \in \mathbb{M}^{(\ell,m)}$ for any values of ℓ and m .

In what follows, we will often simply use the terminology “marked” or “measured” instead of “ ℓ -marked” or “ m -measured” if the numbers ℓ or m are clear from the context. Furthermore, when $m \geq 2$, we will often single out the first measure by writing it as a separate coordinate, writing $(\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \mu_{\mathcal{X}}, \nu_{\mathcal{X}})$, for instance. The reason is that this first measure will often be an *area* measure whereas the other will be *boundary* measures, and these have different natural scales, as we will see shortly.

1.4 The main convergence result

Brownian surfaces. For $k \geq 0$, a *quadrangulation with k holes* is a *bipartite* map having k holes h_1, \dots, h_k and whose internal faces are all of degree 4. For¹ $n \in \mathbb{Z}_{\geq 0}$ and $\mathbf{l} = (l^1, \dots, l^k) \in (\mathbb{Z}_{\geq 0})^k$ (with the convention that $(\mathbb{Z}_{\geq 0})^0 = \{\emptyset\}$), we define the set $\vec{\mathbf{Q}}_{n,\mathbf{l}}^{[g]}$ of all genus g *rooted* quadrangulations with k holes having n internal faces, and whose holes h_1, \dots, h_k are of respective degrees $2l^1, \dots, 2l^k$; see Figure 1.1 for an example.

Likewise, we denote by $\mathbf{Q}_{n,\mathbf{l}}^{[g]}$ the set of *nonrooted* quadrangulations of genus g with n internal faces and half-perimeters given by \mathbf{l} . Since maps are counted with an inverse factor given by the number of automorphisms, the weighted cardinality of this set is

$$\sum_{\mathbf{q} \in \mathbf{Q}_{n,\mathbf{l}}^{[g]}} \frac{1}{\text{Aut}(\mathbf{q})} = \frac{|\vec{\mathbf{Q}}_{n,\mathbf{l}}^{[g]}|}{4n + 2\|\mathbf{l}\|}, \quad (1.2)$$

where $|\vec{\mathbf{Q}}_{n,\mathbf{l}}^{[g]}|$ is the cardinality of $\vec{\mathbf{Q}}_{n,\mathbf{l}}^{[g]}$, and $4n + 2\|\mathbf{l}\|$ is the number of oriented edges, hence of potential roots, in any element of $\mathbf{Q}_{n,\mathbf{l}}^{[g]}$.

It will be useful to notice for further reference that the quadrangulations with k holes in $\vec{\mathbf{Q}}_{n,\mathbf{l}}^{[g]}$ or in $\mathbf{Q}_{n,\mathbf{l}}^{[g]}$ all have the same number of internal vertices, namely

$$n + \|\mathbf{l}\| + 2 - 2g - k. \quad (1.3)$$

Indeed, let us consider such a map, and denote by v, e, f , its number of vertices, edges, faces. The number of external faces is thus $f - n$ so that the desired number

¹We will write $\mathbb{Z}_{\geq 0} = \{0, 1, 2, \dots\}$ the set of nonnegative integers, as well as $\mathbb{N} = \{1, 2, 3, \dots\}$ the set of positive integers.

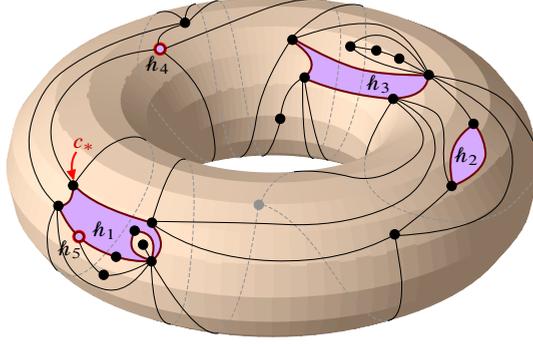


Figure 1.1. A quadrangulation from $\overline{\mathcal{Q}}_{19,(4,1,2,0,0)}^{[1]}$. The root is the corner c_* . Here, h_1 , h_2 , and h_3 are external faces, while h_4 and h_5 are external vertices.

is $v - k + f - n$. Furthermore, counting the corners yields $2e = 4n + 2\|I\|$, and the result follows from the Euler characteristic formula $v - e + f = 2 - 2g$.

If \mathbf{q} is a quadrangulation with k holes, we can view it as a k -marked, $(k + 1)$ -measured metric space, in the following way. We let $V(\mathbf{q})$ be the vertex set of \mathbf{q} , and $d_{\mathbf{q}}$ the graph metric on this set. We let

$$\partial\mathbf{q} = (V(h_1), \dots, V(h_k)),$$

where for $1 \leq i \leq k$, $V(h_i)$ is either $\{h_i\}$ if h_i is an external vertex, or the set of vertices incident to h_i if it is an external face. We let $\mu_{\mathbf{q}}$ and $\nu_{\partial\mathbf{q}}$ be the measures on $V(\mathbf{q})$ and the elements of $\partial\mathbf{q}$ defined by

$$\mu_{\mathbf{q}} = \sum_{v \in V(\mathbf{q})} \delta_v, \quad \nu_{\partial\mathbf{q}}^i = \sum_{v \in V(h_i)} m_v \delta_v,$$

where m_v , the *multiplicity* of v , is the number of corners of the face h_i that are incident to v (by convention, we set $m_{h_i} = 1$ for an external vertex h_i). These are respectively called the *area measure* and *boundary measures*. While we believe that our results also hold when $\nu_{\partial\mathbf{q}}^i$ is replaced by the counting measure on $V(h_i)$ (without multiplicities), it turns out that the above definition makes matter simpler. We associate with the quadrangulation \mathbf{q} the space

$$(V(\mathbf{q}), d_{\mathbf{q}}, \partial\mathbf{q}, \mu_{\mathbf{q}}, \nu_{\partial\mathbf{q}}) \in \mathbb{M}^{(k, k+1)}.$$

Our main result exhibits a family

$$\mathbb{S}_L^{[g]}, \quad g \geq 0, \quad L \in \bigsqcup_{k \geq 0} [0, \infty)^k,$$

of random marked measured metric spaces, where $\mathbb{S}_{\mathbf{L}}^{[g]}$ will be called the *Brownian surface of genus g with boundary perimeter vector \mathbf{L} and unit area*. The latter family describes the scaling limits of uniform random elements of $\vec{\mathbf{Q}}_{n, \mathbf{l}_n}^{[g]}$, in the following sense. Define the scaling operator Ω_n by

$$\Omega_n(\mathbf{q}) = \left(V(\mathbf{q}), \left(\frac{9}{8n} \right)^{1/4} d_{\mathbf{q}}, \partial \mathbf{q}, \frac{1}{n} \mu_{\mathbf{q}}, \frac{1}{\sqrt{8n}} \nu_{\partial \mathbf{q}} \right). \quad (1.4)$$

The scaling constants $(8/9)^{1/4}$ and $\sqrt{8}$ are here to make the upcoming description of $\mathbb{S}_{\mathbf{L}}^{[g]}$ simpler in Sections 3.4 to Chapter 6. Our main result is the following.

Theorem 1.1. *Fix $g, k \geq 0$. Let $\mathbf{L} = (L^1, \dots, L^k)$ be a k -tuple of nonnegative real numbers and, for $n \geq 1$, let $\mathbf{l}_n = (l_n^1, \dots, l_n^k) \in (\mathbb{Z}_{\geq 0})^k$ be such that $l_n^i / \sqrt{2n} \rightarrow L^i$ as $n \rightarrow \infty$ for $1 \leq i \leq k$. Let Q_n be a random variable that is uniformly distributed over $\vec{\mathbf{Q}}_{n, \mathbf{l}_n}^{[g]}$. Then*

$$\Omega_n(Q_n) \xrightarrow[n \rightarrow \infty]{(d)} \mathbb{S}_{\mathbf{L}}^{[g]},$$

where the convergence holds in distribution in the space $(\mathbb{M}^{(k, k+1)}, d_{\text{GHP}}^{(k, k+1)})$.

By our discussion on nonrooted maps, note that the same statement holds if Q_n is rather distributed over the set $\mathbf{Q}_{n, \mathbf{l}_n}^{[g]}$ of nonrooted maps, with a probability proportional to the inverse of the number of automorphisms. Note however that this automorphism number is equal to 1 for the vast majority of maps [89], so we expect that our results also hold for genuine uniform random nonrooted maps.

If $\mathbb{S}_{\mathbf{L}}^{[g]} = (\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \mu_{\mathcal{X}}, \nu_{\mathcal{X}})$, we will call $\mu_{\mathcal{X}}$ the *area measure*, and $\nu_{\mathcal{X}}$ the *boundary measures*. Note that $\mu_{\mathcal{X}}$ is a probability measure, since (1.3) implies that $|V(Q_n)| \sim n$ as $n \rightarrow \infty$, while $\nu_{\mathcal{X}}^i$ has total mass L^i for $1 \leq i \leq k$, so $\nu_{\mathcal{X}}^i$ is the trivial zero measure if $L^i = 0$.

Note that, for $(g, k) = (0, 0)$, the above result amounts to the aforementioned convergence of plane quadrangulations to the Brownian sphere [67, 80], while for $(g, k) = (0, 1)$ with $L^1 > 0$, it corresponds to the convergence of quadrangulations with a boundary to the Brownian disk [25]. Note however that the statement of the present memoir is slightly stronger, since it is formulated in terms of the marked GHP topology rather than the weaker GH topology. In the case $(g, k) = (0, 0)$ of the Brownian sphere, it amounts to the GHP topology since there are no marks and only one measure; this stronger form appears, for instance, in [5, Theorem 1.2] and [68, Theorem 7].

Topology and Hausdorff dimension. Let us also list some basic properties of the limiting metric spaces, which justify the terminology of *Brownian surfaces*. We say that a metric space is *locally of Hausdorff dimension d* if any nontrivial ball has Hausdorff dimension d .

Proposition 1.2. *Let $L = (L^1, \dots, L^k)$ be fixed and let b denote the number of positive coordinates of L . Almost surely, the random metric space $\mathcal{S}_L^{[g]}$ is homeomorphic to $\Sigma_b^{[g]}$, is locally of Hausdorff dimension 4, and, if $b > 0$, each of the b connected components of its boundary, considered as a metric space by restriction of the metric on $\mathcal{S}_L^{[g]}$, is locally of Hausdorff dimension 2.*

This statement is an immediate corollary of a result from [22], showing that, in the case $b = k$, any subsequential limit in distribution of $n^{-1/4}Q_n$ satisfies the stated properties. The case $b < k$ is easily obtained from there by the observation concerning null perimeters at the end of this section. However, our method of proof of Theorem 1.1 will also provide an alternative and rather transparent proof of Proposition 1.2, once an analogous statement has been established for the noncompact analogs of the cases of the sphere and disk, namely, the Brownian plane and the Brownian half-plane [13, 40, 52] (see Section 3.8). We also mention that the case $(g, k, b) = (0, 1, 0)$ was obtained in [21] (see also [29]).

A comment on notation. Throughout this memoir, we will often work in fixed topology and consistently use the following pieces of notation, as in the above statement:

- g for the genus of the surface,
- k for the size of the boundary perimeter vector, that is the number of holes in the discrete maps,
- b , as in *boundary*, for the number of nonzero coordinates in the boundary perimeter vector,
- p , as in *puncture*, for the number of null coordinates in the boundary perimeter vector.

Beware that the latter two numbers do not always correspond to the numbers of external faces and external vertices in the discrete maps, since we only require that $L_n/\sqrt{2n} \rightarrow L$. However, for n sufficiently large, the b holes corresponding to the b nonzero coordinates in the boundary perimeter vector are external faces; each of the p remaining holes can be either a vertex or a face but, in the latter case, it should be thought of as a “small face” in the sense that its perimeter is of order $o(\sqrt{n})$, and we will see that this implies a diameter of order $o(n^{1/4})$.

Method of proof. We prove Theorem 1.1 by some surgical methods, and from the known cases $g = 0$ and $k \in \{0, 1\}$. Heuristically, we will cut Q_n along well-chosen geodesics into a *finite number* of elementary pieces of planar topology, to which we can apply a variant of the cases $(g, k) \in \{(0, 0), (0, 1)\}$ of Theorem 1.1. The idea of cutting quadrangulations along geodesics into so-called slices appears in Bouttier and Guitter [29, 30]. The use of these slices and the study of their scaling limits play an important role in Le Gall’s proof [67] of the uniqueness of the Brownian sphere (they are called *maps with a piecewise geodesic boundary* in this reference)

and are crucial to our study [25] in the case of the disk. More specifically, in the latter reference, we view Brownian disks as a continuum version of the slice decomposition.

The proof of Theorem 1.1 relies on similar but yet different ideas, and will require the introduction of other types of surgeries on objects that we call (*composite*) *slices* and *quadrilaterals (with geodesic sides)*. The core of the proof of Theorem 1.1 consists in showing scaling limit results for these elementary pieces, as stated in Theorems 2.6 and 2.8. We believe that these results are of independent interest, as elementary pieces and their scaling limits might serve as building blocks in other models of random surfaces. In order to prove this result, it turns out that it is simpler to view the discrete and continuum elementary pieces as embedded into noncompact version of the Brownian sphere and disk, namely the Brownian plane and half-plane defined in [13, 40, 52]. We stress that the description of the Brownian half-plane in terms of gluing of composite slices considered in Section 4.3 below is related to the *slice decomposition of metric bands* property used by Miller and Qian [81] for studying geodesic stars in the Brownian sphere.

Theorem 1.1 generalizes the case of the sphere at two different levels, one given by the positive genus and one given by the addition of a boundary. Although these two levels of generalization rely to some extent on similar ideas, the difficulties that they generate are of quite different nature. The case of the disk, which was the focus of [25], relied on relatively well-understood objects, but required gluing an infinite number of such objects, which in principle could create problems in the limit. On the other hand, the surgery involved in the general case consists in gluing a bounded number of objects, but the objects themselves will turn out to be of a more complicated nature.

Null perimeter coordinates. We end this section by the following observation relating Brownian surfaces in case of null perimeter coordinates. The operations of adding or removing a mark used in the following proposition are given by Lemmas 3.1 and 3.4 in Chapter 3.

Proposition 1.3. *Let $\mathbf{L} = (L^1, \dots, L^k) \in [0, \infty)^k$, and $\mathbf{L}0 = (L^1, \dots, L^k, 0)$. Then the space $\mathbb{S}_{\mathbf{L}0}^{[g]}$ has same distribution as the space $\mathbb{S}_{\mathbf{L}}^{[g]}$, where, denoting by μ the area measure of the latter space, a random μ -distributed point has been added to the set of marks of $\mathbb{S}_{\mathbf{L}}^{[g]}$ in $(k + 1)$ -th position (and the zero measure has been added as a trivial $(k + 1)$ -th boundary measure).*

Consequently, if $L^i = 0$ for some given $i \in \{1, 2, \dots, k\}$, and if $\hat{\mathbf{L}}$ denotes the vector \mathbf{L} with i -th coordinate removed, then $\mathbb{S}_{\hat{\mathbf{L}}}^{[g]}$ has same distribution as the space $\mathbb{S}_{\mathbf{L}}^{[g]}$ with its i -th mark and (trivial) i -th boundary measure removed.

Proof. Let us fix $\mathbf{l}_n = (l_n^1, \dots, l_n^k) \in (\mathbb{Z}_{\geq 0})^k$ such that $l_n^j \sim \sqrt{2n}L^j$ for $1 \leq j \leq k$, and let Q_n be uniformly distributed over $\vec{\mathbb{Q}}_{n, \mathbf{l}_n}^{[g]}$. Setting $\mathbf{l}_n 0 = (l_n^1, \dots, l_n^k, 0)$,

a uniformly distributed random variable Q'_n in $\vec{\mathbf{Q}}_{n, \mathbf{l}_n 0}^{[g]}$ may be obtained by choosing an extra distinguished external vertex h_{k+1} uniformly at random among the internal vertices of Q_n , that is, according to the measure μ_{Q_n} conditioned on the set of internal vertices. Since the number of distinguished vertices in Q_n is at most k , while the total number of vertices is asymptotically equivalent to n , the GHP limit of the quadrangulation Q'_n rescaled as in Theorem 1.1 is the same as if we had chosen h_{k+1} uniformly at random among the set of all vertices of Q_n . By Theorem 1.1 applied to Q_n and Lemma 3.1 below, we obtain the result. The second part of the statement is obtained by permuting or removing marks and measures appropriately, as discussed in Lemma 3.4 below. ■

As an example, the Brownian sphere $\mathbb{S}_\emptyset^{[0]}$ can be seen as $\mathbb{S}_{(0,0)}^{[0]}$ by forgetting its two marks. Anticipating on the construction of the Brownian surfaces in Section 3.4, this provides a nontrivially equivalent construction of the Brownian sphere as the gluing of one quadrilateral with geodesic sides, rather than the one from [67, 80].

1.5 Scaling limits of Boltzmann quadrangulations

We may also consider scaling limits for models of quadrangulations with holes in which the area and perimeters are not fixed, but rather weighted by Boltzmann factors. We introduce the following sets of nonrooted maps:

$$\mathbf{Q}_l^{[g]} = \bigsqcup_{n \geq 0} \mathbf{Q}_{n, l}^{[g]} \quad \text{for } g \geq 0, \quad l \in \bigsqcup_{k \geq 0} (\mathbb{Z}_{\geq 0})^k,$$

and

$$\mathbf{Q}^{[g]}(b, p) = \bigsqcup_{l \in \mathbb{N}^b} \mathbf{Q}_{l \mathbf{0}^p}^{[g]} \quad \text{for } g, b, p \geq 0,$$

where $\mathbf{l0}^p$ stands for the sequence l to which we append p terms equal to 0.

We then let \mathcal{W} be the σ -finite measure on the set of nonrooted quadrangulations with an arbitrary number of holes and arbitrary genus, given by

$$\mathcal{W}(\mathbf{q}) = \frac{1}{\text{Aut}(\mathbf{q})} 12^{-|\mathbf{q}|} 8^{-\|\partial \mathbf{q}\|},$$

where $|\mathbf{q}|$ is the number of internal faces of \mathbf{q} , and $\|\partial \mathbf{q}\|$ is the sum of the perimeters of its holes. The reason for the choice of the weights $1/12$ and $1/8$ for the internal faces and perimeters comes from the following enumeration result, which will be proved in an extended form in Proposition B.1, in Appendix B.

Proposition 1.4. Fix $b \geq 0$ and $\mathbf{L} \in (0, \infty)^b$. Let $(\mathbf{l}_n, n \geq 0)$ be a sequence of integers such that $l_n^i \sim \sqrt{2n}L^i$ as $n \rightarrow \infty$ for $1 \leq i \leq b$. Then there exists a continuous function t_g of \mathbf{L} such that

$$\mathcal{W}(\mathbf{Q}_{n, \mathbf{l}_n}^{[g]}) \underset{n \rightarrow \infty}{\sim} t_g(\mathbf{L}) n^{(5g-7)/2+3b/4}.$$

The function $t_g(\mathbf{L})$ is related to the so-called *double scaling limit* of maps, as described in [48, Chapter 5], and its Laplace transform can be computed by solving Eynard and Orantin’s topological recursion. The method presented in Appendix B is based on the bijections presented in Chapter 2.

For any $g \geq 0$, $p \geq 0$, $\mathbf{L} \in [0, \infty)^p$ and $A > 0$, if $(\mathcal{X}, d, \mathbf{A}, \mu)$ is a random variable with same law as $\mathbb{S}_{\mathbf{L}/\sqrt{A}}^{[g]}$, we define the Brownian surface of genus g with boundary perimeter vector \mathbf{L} and area A as a random variable $\mathbb{S}_{A, \mathbf{L}}^{[g]}$ with same law as $(\mathcal{X}, A^{1/4}d, \mathbf{A}, A\mu)$. If $\mathbf{L} \in \bigsqcup_{b \geq 0} (0, \infty)^b$ and $p \geq 0$, we let $\mathbf{L}\mathbf{0}^p \in [0, \infty)^{b+p}$ be the sequence \mathbf{L} to which we append p terms equal to 0.

For integers $g, b, p \geq 0$, and for $\mathbf{L} \in (0, \infty)^b$, setting $k = b + p$, we define a σ -finite measure on $\mathbb{M}^{(k, k+1)}$ by the formula

$$\mathbb{S}_{\mathbf{L}, p}^{[g]}(\cdot) = \int_{(0, \infty)} dA A^{(5g-7)/2+3b/4+p} t_g\left(\frac{\mathbf{L}}{\sqrt{A}}\right) \mathbb{P}(\mathbb{S}_{A, \mathbf{L}\mathbf{0}^p}^{[g]} \in \cdot).$$

The measure $\mathbb{S}_{\mathbf{L}, p}^{[g]}$ is a σ -finite measure that “randomizes” the area measure of the Brownian surface of genus g with b boundary components of lengths given by \mathbf{L} , as well as p marked vertices, in the sense that its conditional law given having total area A is that of $\mathbb{S}_{A, \mathbf{L}\mathbf{0}^p}^{[g]}$.

Recall that the scaling operator Ω_n is defined by (1.4); here, we use it for any $n \in (0, \infty)$.

Theorem 1.5. Let $g, b, p \in \mathbb{Z}_{\geq 0}$, let $k = b + p$, let $\mathbf{L} \in (0, \infty)^b$, let $K > 0$, and let $F: \mathbb{M}^{(k, k+1)} \rightarrow \mathbb{R}$ be a continuous and bounded function that is supported on the set of spaces $(\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \mu_{\mathcal{X}}, \nu_{\mathcal{X}})$ such that $\mu_{\mathcal{X}}(\mathcal{X}) \in [1/K, K]$. Let $(\mathbf{l}_a, a > 0)$ be a family where $\mathbf{l}_a \in \mathbb{N}^b$ is such that $l_a^i \sim \sqrt{2/a}L^i$ for $1 \leq i \leq b$. Then, it holds that

$$a^{5(g-1)/2+3b/4+p} \mathcal{W}(F(\Omega_{a^{-1}}(Q)) \mathbf{1}_{\mathbf{Q}_{\mathbf{l}_a \mathbf{0}^p}^{[g]}}) \xrightarrow{a \downarrow 0} \mathbb{S}_{\mathbf{L}, p}^{[g]}(F).$$

Note that our main result, Theorem 1.1, can be seen as a “local limit” version of Theorem 1.5, in the sense that it gives the conditional statement of this last result given $\mathbf{Q}_{a^{-1}, \mathbf{l}_a \mathbf{0}^p}^{[g]}$, taking $a = 1/n$. There is also a version of this theorem, where the perimeters given by \mathbf{L} are left free as well. For $g, b, p \geq 0$, we define the σ -finite measure

$$\mathbb{S}_{b, p}^{[g]}(\cdot) = \int_{(0, \infty)^b} d\mathbf{L} \mathbb{S}_{\mathbf{L}, p}^{[g]}(\cdot).$$

Corollary 1.6. *Let $g, b, p \in \mathbb{Z}_{\geq 0}$, $k = b + p$, $K > 0$, and let $F: \mathbb{M}^{(k, k+1)} \rightarrow \mathbb{R}$ be a continuous and bounded function that is supported on the set of spaces $(\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \mu_{\mathcal{X}}, \nu_{\mathcal{X}})$ such that $\mu_{\mathcal{X}}(\mathcal{X})$ and $\nu_{\mathcal{X}}^i(A^i)$, $1 \leq i \leq b$, all lie in $[1/K, K]$. Then it holds that*

$$2^{b/2} a^{5(g-1)/2+5b/4+p} \mathcal{W}(F(\Omega_{a^{-1}}(Q))) \mathbf{1}_{\mathcal{Q}^{[g]}(b,p)} \xrightarrow{a \downarrow 0} \mathcal{S}_{b,p}^{[g]}(F).$$

Interestingly, the measure $\mathcal{S}_{L,p}^{[g]}$ is finite in the particular cases $g = 0, b = 1$ and $p \in \{0, 1\}$, or $g = 0, b = 2$ and $p = 0$; it can be checked that it is infinite in all other cases. By computing the functions $t_0(L)$ in the case $b \in \{1, 2\}$, we obtain three probability distributions by normalizing the measures $\mathcal{S}_{(L),0}^{[0]}$, $\mathcal{S}_{(L),1}^{[0]}$, $\mathcal{S}_{(L,L'),0}^{[0]}$. Those are the law of the *free Brownian disk* of perimeter $L \in (0, \infty)$:

$$\text{FBD}_L = \int_0^\infty dA \frac{L^3}{\sqrt{2\pi A^5}} \exp\left(-\frac{L^2}{2A}\right) \mathbb{P}(\mathcal{S}_{(L),A}^{[0]} \in \cdot),$$

the law of the *free pointed Brownian disk* of perimeter $L \in (0, \infty)$:

$$\text{FBD}_L^\bullet = \int_0^\infty dA \frac{L}{\sqrt{2\pi A^3}} \exp\left(-\frac{L^2}{2A}\right) \mathbb{P}(\mathcal{S}_{(L,0),A}^{[0]} \in \cdot),$$

and the law of the *free Brownian annulus* of boundary perimeters L, L' :

$$\text{FBA}_{L,L'} = \int_0^\infty dA \frac{(L + L')}{\sqrt{2\pi A^3}} \exp\left(-\frac{(L + L')^2}{2A}\right) \mathbb{P}(\mathcal{S}_{(L,L'),A}^{[0]} \in \cdot).$$

Note in particular that

$$\text{FBD}_L^\bullet = \lim_{\varepsilon \downarrow 0} \text{FBA}_{L,\varepsilon}.$$

These laws, as well as the associated σ -finite measures $\mathcal{S}_{1,0}^{[0]}$, $\mathcal{S}_{1,1}^{[0]}$, $\mathcal{S}_{2,0}^{[0]}$, play an important role in [10].

In the case $b = 0$, the two previous statements are in fact the same, since $\mathcal{S}_{\emptyset,p}^{[g]} = \mathcal{S}_{0,p}^{[g]}$. This measure describes the scaling limit of quadrangulations with no boundary, p marked vertices, and free area measure. In this case, the quantity $t_g(\emptyset)$ is equal to the classical universal constant t_g arising in map enumeration; see [17, 62]. Explicitly, the numbers $\tau_g = 2^{5g-2} \Gamma((5g-1)/2) t_g$ satisfy $\tau_0 = -1$ and the recursion

$$\tau_{g+1} = \frac{(5g+1)(5g-1)}{3} \tau_g + \frac{1}{2} \sum_{h=1}^g \tau_h \tau_{g+1-h}, \quad g \geq 0.$$

In this case, we thus have the following formula:

$$\mathcal{S}_{0,p}^{[g]}(\cdot) = t_g \int_{(0,\infty)} dA A^{(5g-7)/2+p} \mathbb{P}(\mathcal{S}_{A,0^p}^{[g]} \in \cdot).$$

1.6 Perspectives

A natural question, which we plan to investigate in future works, is to derive the analog of Theorem 1.1 for bipartite quadrangulations on *nonorientable compact surfaces*, using the bijective techniques developed in [23, 34]. The first step of showing the existence of subsequential limits for nonorientable quadrangulations without boundary has been taken in [34]. Addressing this question would complete the catalog of compact Brownian surfaces.

As mentioned in the first section of this introduction, an important aspect is that of *universality* of the spaces $S_L^{[g]}$. In fact, we expect these spaces to be the scaling limits of many other models of random maps on surfaces. In the case of the Brownian sphere $S_\emptyset^{[0]}$, this was indeed verified for several models; see the references mentioned above. In the case of Brownian disks, we showed in [25] that the spaces $S_{(L,1)}^{[0]}$ appear as scaling limits of many conditioned Boltzmann models. This approach to universality should generalize to our context, at the price of some specific technicalities. We will not address this question here, but will comment more on this in Chapter 6.

It would be most interesting to complete the bridge between Brownian surfaces and LQG metrics and CFT (conformal field theory). As was pointed to us by Jason Miller, in order to define a canonical conformal structure and Brownian motion on Brownian surfaces, it would be natural to investigate whether the construction of general Brownian surfaces given in the present memoir, by gluing elementary pieces of disk topologies along geodesic boundaries, can be made compatible with the approach of [56, 57] mentioned in the introduction. Knowing that such a structure exists, one can try to delve even further into its integrability properties. The works [45, 50] state precise conjectures linking Liouville CFT with scaling limits of the area measure of random maps (without boundary) after suitable uniformization. In a nutshell, the LQG metrics are local objects that can be defined globally by using charts and atlases on general Riemann surfaces. However, fixing a surface amounts to fixing the conformal modulus of the LQG metric, while Brownian surfaces have a random modulus. Hence, the computation of the law of this modulus is an important question, which has been solved by [10] in the case of the annular topology. It seems that the case of general compact surfaces should be approachable as well given the recent developments on conformal bootstrap in Liouville CFT [49, 94].

We also mention that random surfaces with boundaries of the type studied in this memoir are related to the study of self-avoiding paths in random geometries. See [53, 54] for more on this in the case of the gluing of two Brownian half-planes or disks. It would be interesting to explicitly describe the scaling limits of self-avoiding paths and loops on maps of fixed topologies as gluings of Brownian surfaces along boundaries. As Nina Holden pointed to us, this would involve presumably difficult computations of the partition functions for self-avoiding loops in fixed classes of the

fundamental group of the surface, although this problem simplifies in the case of the self-avoiding loop on a Brownian sphere [9].

1.7 Organization of the memoir

In Chapter 2, we present the extension of the famous Cori–Vauquelin–Schaeffer bijection allowing to encode a quadrangulation with a simpler tree-like structure carrying integer labels on its vertices. We also present a variant of the bijection, which leads to the definition of the elementary pieces into which we decompose a quadrangulation. We finally state the relevant scaling limit results for these elementary pieces. In Chapter 3, we present the surgical operation we need in order to reconstruct a metric space from its elementary pieces, namely gluing along geodesic segments. The proofs of Theorem 1.1 and Proposition 1.2 are given in Section 3.4. In Chapters 4 and 5, we present the metric spaces forming the continuum elementary pieces into consideration and explain how they are natural building blocks of the Brownian plane and half-plane, which are the noncompact analogs of the Brownian sphere and disk, and we tweak known convergence results to these noncompact Brownian surfaces to prove that the continuum elementary pieces are the scaling limits of the discrete elementary pieces. Finally, we give in Chapter 6 an alternate description of Brownian surfaces that does not involve gluing operations and that is closer to the usual definition of the Brownian sphere and disks.