Geometry of random planar maps and genus-o hyperbolic surfaces

Timothy Budd

IMAPP, Radboud University, Nijmegen, The Netherlands.
t.budd@science.ru.nl

June 14, 2023

Abstract

These are notes for the mini-course at the ISM Discovery school on "Geometry and spectra of random hyperbolic surfaces" in Montréal.

1 Introduction

2 Generating functions of Weil–Petersson volumes

2.1 Weil–Petersson volumes

For $g, n \geq 0$ and a sequence $L = (L_1, \ldots, L_n)$ of non-negative real numbers, we consider the moduli space $\mathcal{M}_{g,n}(L)$ of genus-g hyperbolic surfaces with n geodesic boundaries of lengths $L_1, \ldots, L_n$. In the case $L_i = 0$, the ith boundary is a cusp.

In the previous courses we have seen that $\mathcal{M}_{g,n}(L)$ comes equipped with the Weil–Petersson symplectic structure ω, which gives rise to the Weil–Petersson volume form

$$\mu_{\text{WP}} = \frac{\omega^{3g-3+n}}{(3g-3+n)!}. \tag{1}$$

The total Weil-Petersson volume

$$V_{g,n}(L) = \int_{\mathcal{M}_{g,n}} \mu_{\text{WP}} \tag{2}$$

of the moduli space is finite and given by a polynomial in boundary lengths [15].

2.2 Generating functions

The goal of this section is to learn something about the Weil-Petersson volumes for arbitrary n and fixed genus, which we will soon restrict to $g = 0$. For this it is convenient to encode the volumes in a single generating function, such that knowledge of this generating function allows us to retrieve any of the volumes. Since the volumes depend on the boundary lengths, we need to incorporate a non-trivial

weighting to achieve this. To this end we consider a (Borel) measure q on $[0,\infty)$ and informally define the genus-g *partition function* to be

$$F_g^{\mathrm{WP}}(q) = \sum_{n=0}^{\infty} \frac{1}{n!} \int_{[0,\infty)^n} V_{g,n}(L)\mathrm{d}q(L_1)\cdots\mathrm{d}q(L_n). \tag{3}$$

When we are interested in doing probability theory, it makes sense to take this definition literally: suppose each of the (Lebesgue) integrals converges and the sum as well, then the right-hand side normalized by $1/F_g^{\mathrm{WP}}$ gives rise to a probability measure, that we call the *genus-g Boltzmann hyperbolic surface* with weight q. Indeed, we would sample n with probability proportional to the summand. Subsequently sample $L \in [0,\infty)^n$ with density proportional to $V_{g,n}(L)\mathrm{d}q(L_1)\cdots\mathrm{d}q(L_n)$ and finally a random hyperbolic surface from $\mathcal{M}_{g,n}(L)$ with density proportional to μ_{WP}.

As an example we could fix $\ell \geq 0$ and take the weight $q = x\delta_\ell$ to be a multiple of the delta measure at ℓ. If then $x > 0$ is chosen small enough such that $F_g^{(\mathrm{WP})}(q) = \sum_{n=0}^{\infty} \frac{x^n}{n!} V_{g,n}(\ell,\ldots,\ell) < 0$, then the associated Boltzmann hyperbolic surface would be a random hyperbolic surface with a random number of boundaries, each of length ℓ.

If we do not wish to worry about convergence issues (yet), it suffices to note that (3) depends only on the even moments of q, since the polynomials $V_{g,n}(L)$ are even. Hence, with some abuse of notation, we may write

$$F_g^{\mathrm{WP}}(q) = F_g^{\mathrm{WP}}(t_0(q), t_1(q), \ldots), \tag{4}$$

where the *times* $t_0(q), t_1(q), \ldots$ are the conveniently normalized even moments

$$t_k(q) = \frac{2}{4^k\, k!} \int_0^{\infty} L^{2k}\mathrm{d}q(L). \tag{5}$$

The same information about the Weil–Petersson volumes is thus contained in $F_g^{\mathrm{WP}}(t_0, t_1, \ldots)$, which we interpret as a formal power series in the variables $t_0, t_1, \ldots$.

Importantly, we can retrieve the Weil–Petersson volumes from F_g^{WP} by using the (formal) *functional derivative* $\frac{\delta}{\delta q(\ell)}$ defined via

$$\frac{\delta}{\delta q(\ell)} F(q) = \frac{\partial}{\partial x} F[q + x\delta_\ell]\Big|_{x=0}. \tag{6}$$

For instance

$$\frac{\delta}{\delta q(\ell)} F_g^{\mathrm{WP}}(q) = \sum_{n=0}^{\infty} \frac{1}{n!} \int_{[0,\infty)^n} V_{g,n+1}(L,\ell)\mathrm{d}q(L_1)\cdots\mathrm{d}q(L_n) \tag{7}$$

so $\frac{\delta}{\delta q(\ell)}$ serves to add a distinguished boundary of length ℓ, which does not receive a weight. In particular, the Weil-Petersson volume itself is retrieved via

$$V_{g,n}(L) = \left.\frac{\delta F_g^{\mathrm{WP}}(q)}{\delta q(L_1)\cdots\delta q(L_n)}\right|_{q=0}. \tag{8}$$

Note also that $\frac{\delta}{\delta q(\ell)}$ is related to partial derivatives with respect to the times via

$$\frac{\delta}{\delta q(\ell)} \equiv \sum_{k=0}^{\infty} \frac{2\ell^{2k}}{4^k k!} \frac{\partial}{\partial t_k}, \qquad \frac{\delta}{\delta q(0)} \equiv \frac{\partial}{\partial t_0}. \tag{9}$$

2.3 An explicit evaluation of the genus-0 partition function

From here we specialize to genus 0, although much of this section holds for higher genus as well. Since the work of Mirzakhani [16] it is known that the coefficients in the polynomials $V_{g,n}(L)$ are directly related to certain intersection numbers on the compactified moduli space $\overline{\mathcal{M}}_{g,n}$ of genus-g curves with n marked points. We will not delve into this topic, except to mention that there are canonical cohomology classes of 2-forms κ_1 and $\psi_1, \dots, \psi_n$ on $\overline{\mathcal{M}}_{g,n}$, whose integrals provide these intersection numbers. In particular,

$$F_0^{\mathrm{WP}}(t_0, t_1, \dots) = \sum_{n=3}^{\infty} \sum_{\substack{m, d_1, \dots, d_n \geq 0 \\ m+d_1+\dots+d_n = n-3}} \frac{\pi^{2m}}{m!} t_{d_1} \cdots t_{d_n} \int_{\overline{\mathcal{M}}_{0,n}} \kappa_1^m \psi_1^{d_1} \cdots \psi_n^{d_n}. \tag{10}$$

A relation between κ_1 and the ψ-classes [19] implies the identification [11]

$$F_0^{\mathrm{WP}}(t_0, t_1, \dots) = F_0(t_0, t_1, t_2 + \gamma_2, t_3 + \gamma_3, \dots), \qquad \gamma_k = \frac{(-1)^k}{(k-1)!} \pi^{2k-2} \mathbf{1}_{k \geq 2}, \tag{11}$$

where F_0 is the generating functions of pure ψ-class intersection numbers

$$\begin{aligned} F_0(t_0, t_1, \dots) &= \sum_{n=3}^{\infty} \sum_{\substack{d_1, \dots, d_n \geq 0 \\ d_1+\dots+d_n = n-3}} t_{d_1} \cdots t_{d_n} \int_{\overline{\mathcal{M}}_{0,n}} \psi_1^{d_1} \cdots \psi_n^{d_n} \\ &= \frac{1}{6} t_0^3 + \frac{1}{6} t_0^3 t_1 + \frac{1}{24} (t_0^4 t_2 + 4 t_0^3 t_1^2) + \cdots . \end{aligned}$$

The reason to introduce F_0 here is the famous conjecture of Witten [19], proved by Kontsevich [7] and others, showing that $\sum_g F_g$ is a τ-function of the KdV hierarchy and thus satisfies an infinite hierarchy of partial differential equations. We focus on a single one, sometimes referred to as the *string equation*,

$$\frac{\partial F_0}{\partial t_0} = \frac{t_0^2}{2} + \sum_{i=0}^{\infty} t_{i+1} \frac{\partial F_0}{\partial t_i}, \tag{12}$$

since it already uniquely determines $F_0(t_0, t_1, \dots)$, see e.g. [6].

Proposition 1. *Let $u_0(t_0, t_1, \dots)$ be the formal power series solution to $Z(u_0) = 0$, where*

$$Z(r) \equiv Z(r; t_0, t_1, \dots) = r - \sum_{k=0}^{\infty} t_k \frac{r^k}{k!}. \tag{13}$$

Then we have

$$u_0 = \frac{\partial^2 F_0}{\partial t_0^2} \quad \text{and} \quad F_0(t_0, t_1, \dots) = \frac{1}{2} \int_0^{u_0} Z(r)^2 \mathrm{d}r. \tag{14}$$

We have included the proof of this proposition in Appendix **??**. Combining with the relation (11) we thus obtain the following characterization of the partition function F_0^{WP}.

Theorem 2. *Let $R \equiv R(t_0, t_1, \dots) = u_0(t_0, t_1, t_2 + \gamma_2, \dots)$ be the formal power series solution to*

$$Z^{\mathrm{WP}}(R) = 0, \qquad Z^{\mathrm{WP}}(r) = r - \sum_{k=0}^{\infty} (t_k + \gamma_k) \frac{r^k}{k!}. \tag{15}$$

Then we have

$$R = \frac{\partial^2 F_0^{\mathrm{WP}}}{\partial t_0^2} \quad \text{and} \quad F_0^{\mathrm{WP}} = \frac{1}{2} \int_0^{R} Z^{\mathrm{WP}}(r)^2 \mathrm{d}r. \tag{16}$$

Substituting the definition of the times (5) and the shift (11), we obtain an equivalent expression of $Z^{\mathrm{WP}}(r)$ in terms of the weights,

$$Z^{\mathrm{WP}}(r) = \frac{\sqrt{r}}{\sqrt{2}\,\pi} J_1(2\pi\sqrt{2r}) - \int_0^\infty I_0(\ell\sqrt{2r})\mathrm{d}q(\ell), \tag{17}$$

where J_1 and I_0 are the modified Bessel functions with series expansions

$$J_1(2x) = \sum_{k=1}^{\infty} (-1)^{k+1} \frac{x^{2k-1}}{k!(k-1)!}, \qquad I_0(2x) = \sum_{k=0}^{\infty} \frac{x^{2k}}{k!^2}. \tag{18}$$

Question. *Why does the generating function $R = \frac{\delta^2 F_0^{\mathrm{WP}}}{\delta q(0)^2}$ of Weil–Petersson volumes with two marked cusps satisfy a closed equation like (15) or (17)? Can a geometric understanding provide an entrance to studying the geometry of (random) hyperbolic surfaces with many boundaries?*

Using that $Z^{\mathrm{WP}}(R) = 0$ and $\frac{\partial Z^{\mathrm{WP}}(r)}{\partial q(\ell)} = I_0(\ell\sqrt{2r})$, it is straightforward to determine further universal formulas for generating functions with marked boundaries, e.g.

$$\begin{aligned}
\frac{\delta F^{\mathrm{WP}}(q)}{\delta q(L_1)} &= \int_0^R I_0(L_1\sqrt{2r}) Z^{\mathrm{WP}}(r)\mathrm{d}r, \\
\frac{\delta^2 F^{\mathrm{WP}}(q)}{\delta q(L_1)\delta q(L_2)} &= \int_0^R I_0(L_1\sqrt{2r}) I_0(L_2\sqrt{2r})\mathrm{d}r, \\
\frac{\delta^3 F^{\mathrm{WP}}(q)}{\delta q(L_1)\delta q(L_2)\delta q(L_3)} &= I_0(L_1\sqrt{2R}) I_0(L_2\sqrt{2R}) I_0(L_2\sqrt{2R}) \frac{\delta R}{\delta q(0)}.
\end{aligned}$$

3 Planar maps

There is a close analogy between Weil–Petersson volumes of genus-g hyperbolic surfaces and the enumeration of genus-g maps (also known as ribbon diagrams). A connection between intersection numbers and maps already featured prominently in Kontsevich's proof [7] of Witten's conjecture. Here we will content ourselves with noting the similarity between the generating function $F_0^{\mathrm{WP}}(q)$ of Weil-Petersson volumes and the generating function of planar maps. But first we need to get some definitions out of the way.

3.1 Introducing maps

A *genus-g map* $\mathfrak{m}$ is a multigraph, i.e. a graph with multiple edges and loops allowed, that is properly embedded in a genus-g surface. An embedding is proper if the edges do not intersect, except possibly at their shared vertex, and if the complement of the graph in the surface is a collection of topological disks, called the *faces* of $\mathfrak{m}$. When viewed modulo orientation-preserving homeomorphisms of the surface, a genus-g map can be described purely combinatorially by its incidence relations. For instance we can think of $\mathfrak{m}$ as a gluing of n (regular) polygons into a surface of genus g by specifying the pairwise identifications of the sides of the polygons. The *degree* of a face or a vertex of $\mathfrak{m}$ is the number of edges incident to the face, where we count an edge double if it is incident on both sides to the same face or vertex. The sets of vertices, edges and faces of $\mathfrak{m}$ are denoted by $V(\mathfrak{m}), E(\mathfrak{m}), F(\mathfrak{m})$.

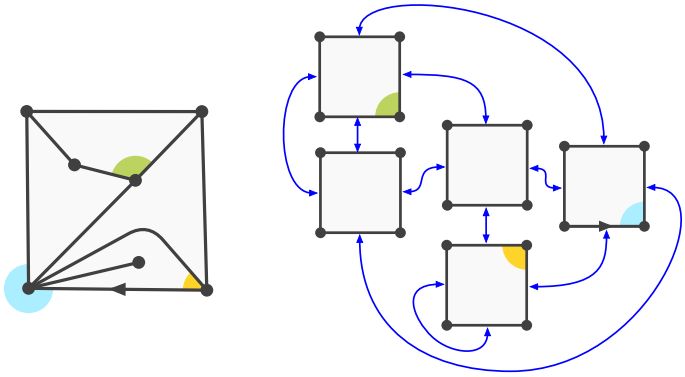

Figure 1: An example of a rooted planar map with all faces of degree 4 (also known as a *quadrangulation*). Equivalently it can be understood as a gluing of identical squares into a topological sphere. The colored corners are included to help see the relation between the two representations.

3.2 Planar map generating function

For our purposes we will restrict to *even maps*, i.e. maps for which each face has even degree. We denote the set of such maps with n faces of degrees $2d_1, \dots, 2d_n$ by $M_{g,n}(d)$, in analogy with the moduli space $\mathcal{M}_{g,n}(L)$ of hyperbolic surfaces with n boundaries. Then for $t, q_1, q_2, \dots \geq 0$ the *partition function of even genus-g maps* is defined as

$$F_g^{\mathrm{m}}(t,q) = \sum_{n=1}^{\infty} \frac{1}{n!} \sum_{d_1=1}^{\infty} q_{d_1} \cdots \sum_{d_n=1}^{\infty} q_{d_n} \sum_{\mathfrak{m} \in M_{g,n}(d)} \frac{t^{|V(\mathfrak{m})|}}{|\operatorname{Aut}(\mathfrak{m})|}, \tag{19}$$

where $\operatorname{Aut}(\mathfrak{m})$ is the group of face-preserving automorphisms of $\mathfrak{m}$ and the factor $1/|\operatorname{Aut}(\mathfrak{m})|$ is included to accommodate properly the maps that have non-trivial symmetries. Since it is often painful to keep track of these, we usually work with *rooted* maps, which are maps together with a distinguished oriented edge. This additional bit of information suffices to kill any symmetries. Denoting by $\vec{M}_g$ the set of all rooted genus-g maps, it is a little exercise to see that (19) can equivalently be expressed as

$$F_g^{\mathrm{m}}(t,q) = \sum_{\mathfrak{m} \in \vec{M}_g} \frac{1}{2|E(\mathfrak{m})|} t^{|V(\mathfrak{m})|} \prod_{f \in F(\mathfrak{m})} q_{\deg f/2}. \tag{20}$$

The goal of this section will be to derive the following characterization of $F_0^{\mathrm{m}}(q)$ in the planar case, which in one form or another goes back all the way to the work of Tutte [18] in the sixties.

Theorem 3. *Let $R(t,q) = \frac{t}{1-q_1} + O(t^2)$ be the formal power series solution to $g_q(R) = t$, where*

$$g_q(r) = r - \sum_{k=1}^{\infty} q_k \binom{2k-1}{k} r^k. \tag{21}$$

Then

$$R = \frac{\partial^2 F_0^{\mathrm{m}}}{\partial t \partial q_1} \quad \text{and} \quad F_0^{\mathrm{m}}(t,q) = \frac{1}{2} \int_0^R \left((g_q(r) - t)^2 - (r-t)^2 \mathbf{1}_{r<t} \right) \frac{\mathrm{d}r}{r}. \tag{22}$$

A specialized generating function that we will encounter later is the *pointed disk function* $W_\bullet^{(\ell)}(t,q)$ which enumerates rooted planar maps with a distinguished vertex and in which the root face, i.e. the face to the left of the root, is restricted to have degree 2ℓ. In this case the distinguished vertex does not receive weight t and the root face does not receive a weight q_ℓ, meaning

$$W_\bullet^{(\ell)}(t,q) = \sum_{\substack{\mathfrak{m}\in\vec{M}_0 \\ \text{root face deg. } 2\ell}} t^{|V(m)|-1} \prod_{\substack{f\in F(\mathfrak{m}) \\ \text{except root face}}} q_{\deg f/2}. \tag{23}$$

It is not difficult to see from (19) or (20) that it is related to F_0^{m} via

$$W_\bullet^{(\ell)}(t,q) = 2\ell \frac{\partial^2 F_0^{\mathrm{m}}}{\partial t \partial q_\ell}. \tag{24}$$

From (22) together with $g_q(R) = t$ and $R \geq t$ it follows that

$$\begin{aligned} W_\bullet^{(\ell)}(t,q) &= \ell \frac{\partial}{\partial t} \int_0^R \frac{\partial}{\partial q_\ell} \left((g_q(r)-t)^2 - (r-t)^2 \mathbf{1}_{r<t} \right) \frac{\mathrm{d}r}{r} \\ &= -2\ell \frac{\partial}{\partial t} \int_0^R (g_q(r)-t) \binom{2\ell-1}{\ell} r^{\ell-1} \mathrm{d}r \\ &= 2\ell \int_0^R \binom{2\ell-1}{\ell} r^{\ell-1} \mathrm{d}r \\ &= 2\binom{2\ell-1}{\ell} R^\ell = \binom{2\ell}{\ell} R^\ell. \end{aligned}$$

Vice versa, this expression implies (22). To see this we note that the right-hand side of (22) indeed determines a power series in $t, q_1, q_2, \ldots$ and vanishes when $t = 0$ or $q = 0$, as it should since every map has at least one vertex and one face. The next section will be devoted to proving $W_\bullet^{(\ell)} = \binom{2\ell}{\ell} R^\ell$ and thus Theorem 3.

3.3 Tree bijection

Even though its solution dates back 60 years, solving the enumeration problem of planar maps at the level of generating functions requires some ingenuity (see e.g. [18, 2, 5]). Compared to maps, trees are generally much simpler combinatorial objects to enumerate. This owes to the natural decomposition that many types of trees admit, often leading to the associated generating function satisfying a characteristic equation.

The simplest example is the *rooted plane tree*, which is a rooted planar map with a single face. If we include for convenience the tree consisting of a single vertex, then the generating function $G(x) = 1 + x + \cdots$ of rooted plane trees with a weight x per edge, satisfies the equation

$$G(x) = 1 + xG(x)^2.$$

This equation follows from the observation that a rooted plane tree either is a single vertex (with weight 1) or it decomposes into a pair of plane trees (with generating function $G(x)^2$) by removing the root edge (which has weight x):

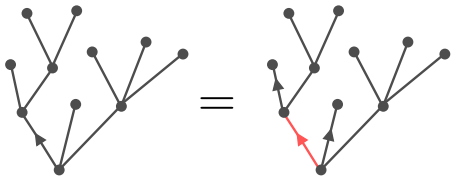

Of course, in this case the solution to the characteristic equation leads to the generating function of the Catalan numbers.

Since the sought-after generating function R of Theorem 3 satisfies an equation $g_q(R) = t$, one may rightfully suspect that R has an interpretation as a generating function of certain trees. The explanation is provided by the *Bouttier-Di Francesco-Guitter* (BDFG) bijection [3], which we will explain now.

Let us denote by $\vec{M}_0^\bullet$ the set of rooted bipartite planar maps with a distinguished vertex (we say the map is *pointed* and call this vertex the *origin*), with the additional restriction that the root edge is oriented away from the origin (in the sense that the graph distance to the origin increases along the root edge). See Fig. 2a for an example.

We will associate a tree with a map $\mathfrak{m} \in \vec{M}_0^\bullet$ via the following prescription.

1. Label the (white) vertices of $\mathfrak{m}$ by the graph distance to the pointed vertex.
2. Draw a new black vertex in each face of $\mathfrak{m}$.
3. For each edge e of $\mathfrak{m}$, let v be the endpoint of e with the largest label. Draw a new red edge starting at v and ending on the black vertex within the face to the left of e when facing v (Fig. 2b). If e is the root of $\mathfrak{m}$ we take the new edge to be the new root (oriented away from v).
4. Remove all original edges of $\mathfrak{m}$ as well as the origin vertex (Fig. 2d).
5. Shift all labels uniformly such that the root vertex receives label 0.

The claim is that the resulting map is a particular type of tree called a mobile. A *mobile* is a tree $\mathfrak{t}$, i.e. a rooted planar map with only one face, with unlabeled black vertices and integer-labeled white vertices satisfying the following properties (see Fig. 2d for an example):

(i) The endpoints of each edge have distinct color.

(ii) The root edge starts at a white vertex with label 0.

(iii) Around each black vertex, if a white neighbour has label ℓ then the next white neighbour in counterclockwise order around the black vertex must have label at most $\ell + 1$.

We denote the set of mobiles by Mob.

Theorem 4 (BDFG bijection). *This construction determines a bijection* $\Phi : \vec{M}_0^\bullet \to \mathrm{Mob}$ *satisfying*

- *each face of degree* $2k$ *in* $\mathfrak{m}$ *corresponds to a black vertex of degree* k *in* $\Phi(\mathfrak{m})$*;*
- *each edge of* $\mathfrak{m}$ *corresponds to an edge of* $\Phi(\mathfrak{m})$*;*
- *if* ℓ_0 *is the distance between the endpoint of the root edge and the origin in* $\mathfrak{m}$*, then the minimal label on* $\Phi(\mathfrak{m})$ *is* $-\ell_0 + 1$*;*

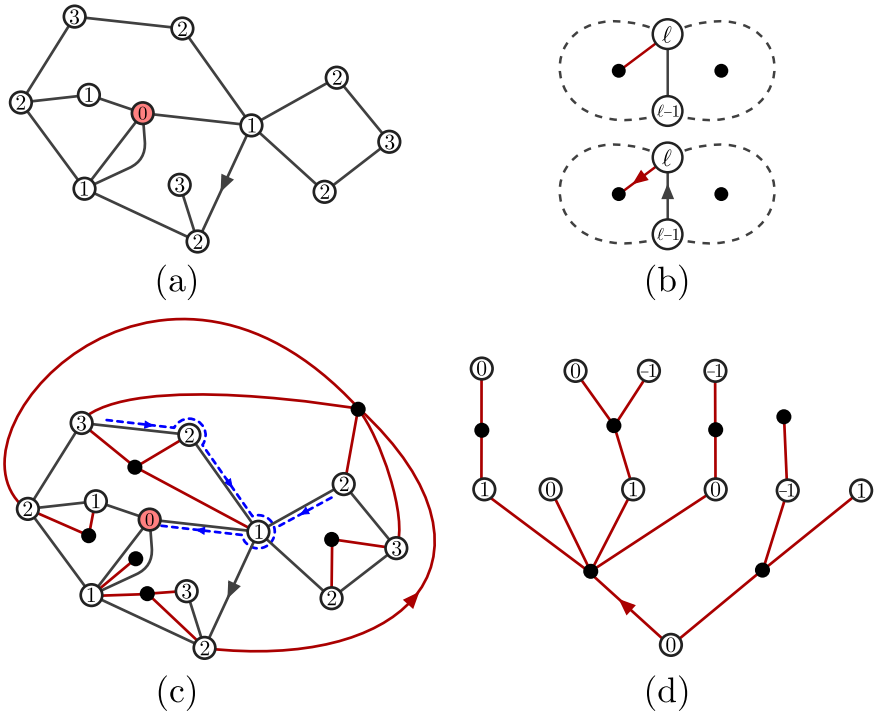

Figure 2: The Bouttier–DiFrancesco–Guitter bijection. (a) A rooted bipartite planar map $\mathfrak{m}$ with a distinguished vertex (shaded in red) together with its canonical labeling by the graph distance. (b) The prescription for drawing new (red) edges and a new root. (c) The result of applying the prescription to $\mathfrak{m}$. The (blue) dashed lines indicate two left-most geodesics. (d) After deleting the edges of $\mathfrak{m}$ and the origin and shifting the labels such that the start of the root edge receives label 0, one obtains the mobile $\mathfrak{t}$.

- *each vertex of* $\mathfrak{m}$ *at graph distance* $\ell > 0$ *from the origin corresponds to a white vertex of label* $\ell - \ell_0$ *in* $\Phi(\mathfrak{m})$.

We will only sketch the main ingredients of the proof.

The result $\Phi(\mathfrak{m})$ is a mobile. First we need to convince ourselves that the constructions yields a tree, in particular that $\mathfrak{t} = \Phi(\mathfrak{m})$ cannot contain cycles. The explanation is that for any edge e of $\mathfrak{m}$ one can find a plane curve, called the *left-most geodesic (to the origin)*, starting at (say, the midpoint of) e and ending at the origin that does not intersect $\mathfrak{t}$. This immediately implies the impossibility of cycles in $\mathfrak{t}$, because every cycle in $\mathfrak{t}$ would enclose the origin on one side and at least one edge on the other side, contradicting the existence of a left-most geodesic path from that edge to the origin.

The left-most geodesic is constructed as follows (see the dashed curve in Fig. 2c for an example): denote the endpoints of e by v_ℓ and $v_{\ell-1}$ with labels ℓ and $\ell - 1$ respectively. The curve starts by traversing e towards $v_{\ell-1}$. If $\ell = 1$, $v_{\ell-1}$ is the origin and we are done. Otherwise, the curve circles around $v_{\ell-1}$ in clockwise direction until it encounters an edge with endpoint at distance $\ell - 2$, that we denote $v_{\ell-2}$. Such an edge always exists due to the definition of the graph distance, and by construction of $\mathfrak{t}$ one encounters no edge of $\mathfrak{t}$ along the way. Traversing the edge to $v_{\ell-2}$ and iterating, one obtains a curve ending at the origin v_0, since that is the unique vertex with minimal label. The path $v_\ell, v_{\ell-1}, \ldots, v_0$

in $\mathfrak{m}$ is called the left-most geodesic, because it is a path of minimal length from v_ℓ to the origin and at each vertex it chooses the left-most option among such minimal paths.

In the absence of cycles, the number of connected components of $\mathfrak{t}$ is given by $|V(\mathfrak{t})| - |E(\mathfrak{t})|$. But by construction $|V(\mathfrak{t})| = |V(\mathfrak{m})| + |F(\mathfrak{m})| - 1$ and $|E(\mathfrak{t})| = |E(\mathfrak{m})|$, which together with Euler's formula implies that $\mathfrak{t}$ has a single connected component and is thus a tree.

It remains to check the constraints on the labels of $\mathfrak{t}$. To this end let us have a look at a single face of degree $2k$ in $\mathfrak{m}$, like in this example:

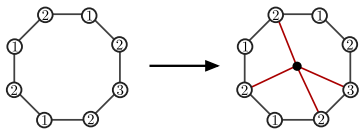

Examining the labels around the face in counterclockwise direction, we find exactly k increments and k decrements. Hence the black vertex inside this face will receive k red edges. Moreover, each red edge is followed by any number of decrements and exactly one increment before the next red edge is encountered. This explains the label constraint on the mobiles as well as the properties listed in Theorem 4.

The inverse construction. The angular region around a vertex that is delimited by two neighbouring edges incident is called a *corner* of that vertex. The *contour* of a tree is the cyclic sequence of corners one encounters while walking around the tree in clockwise direction (i.e. keeping the tree on the right-hand side). Given a mobile $\mathfrak{t}$ one may construct a (rooted pointed bipartite planar) map as follows.

1. Assuming $-\ell_0 + 1$ is the minimal label of $\mathfrak{t}$, insert a new white vertex (the origin) with label $-\ell_0$ in the face of $\mathfrak{t}$.

2. For each corner c of a white vertex with label ℓ in $\mathfrak{t}$, we draw a new edge from c to the next corner of a white vertex in the contour that has label $\ell - 1$ in case $\ell > -\ell_0 + 1$ or to the origin in case $\ell = -\ell_0 + 1$. If c is the corner of the root vertex that sits left of the root edge, then the new edge is taken to be the new root (oriented away from c).

3. Remove all (red) edges of $\mathfrak{t}$.

One can show [3] that the drawing of edges can be done uniquely in a planar fashion and that this construction yields exactly the inverse mapping $\Phi^{-1} : \mathrm{Mob} \to \vec{M}_0^\bullet$.

3.4 Generating function via mobiles

Recall that $W_\bullet^{(\ell)} = 2\frac{\partial^2 F_0^{\mathrm{m}}}{\partial t \partial q_1}$ is the generating function of rooted pointed bipartite maps with root face degree 2ℓ. Consider $\ell = 1$, subtracting the contribution $2t$ of the maps consisting of a single edge, gluing the two sides of the root face together and dividing by two to restrict the root edge to point away from the origin, we deduce that

$$\frac{1}{2}W_\bullet^{(1)} - t = R - t = \sum_{\mathfrak{m} \in \vec{M}_0^\bullet} t^{|V(\mathfrak{m})|-1} \prod_{f \in F(\mathfrak{m})} q_{\deg f/2} \tag{25}$$

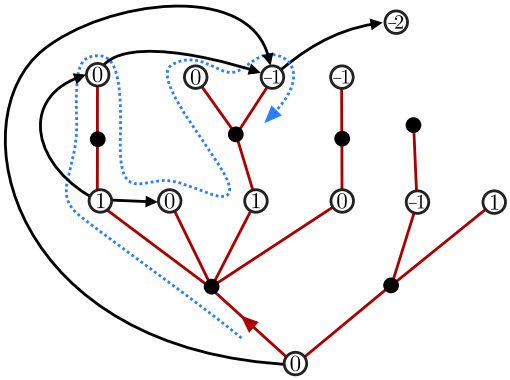

Figure 3: Illustration of the inverse map Φ^{-1}. The contour is illustrated in dashed blue, and the new black edges associated to five of the corners have been drawn.

is the generating function of the maps in $\vec{M}_0^\bullet$. By Theorem 4 it therefore also is the generating function of mobiles,

$$R(t,q) - t = \sum_{\mathfrak{t} \in \mathrm{Mob}} t^{|V_{\mathrm{white}}(\mathfrak{t})|} \prod_{v \in V_{\mathrm{black}}(\mathfrak{t})} q_{\deg v}. \tag{26}$$

Or if you wish, $R(t,q)$ is the generating function of mobiles which includes the degenerate mobile consisting of a single white vertex (with label 0).

Now we can understand that R satisfies the equation $g_q(R) = t$, because mobiles admit a decomposition analogous to the rooted plane trees before. Let us denote the root vertex by v_0. If the degree of the black vertex at the end of the root edge is k, then it has $k-1$ white children $v_1, \dots, v_{k-1}$. Each vertex v_i together with its offspring, excluding the branch of the root edge in case of v_0, determines a mobile, once the labels have been shifted such that the root vertex v_i receives label 0. Noting that these mobiles may take the form of a single white vertex with no children, this leads immediately to the equation

$$R = t + \sum_{k=1}^{\infty} q_{2k} \sum_{\substack{\text{labels on} \\ v_1,\dots,v_{k-1}}} \quad = \sum_{k=1}^{\infty} q_{2k} \binom{2k-1}{k} R^k, \tag{27}$$

because there are precisely $\binom{2k-1}{k}$ choices for the labels $\ell_1, \cdots, \ell_{k-1}$ on $v_1, \dots, v_{k-1}$ satisfying the requirements $\ell_1 \geq -1$, $\ell_k \leq 1$, and $\ell_{i+1} \geq \ell_i - 1$ for $i = 1, \dots, k-1$. This is equivalent to $g_q(R) = t$.

Since the summand gives the contribution of maps with root face of degree $2k$, we also reproduce

$$W_\bullet^{(2k)} = 2\binom{2k-1}{k} R^k = \binom{2k}{k} R^k, \tag{28}$$

where the factor of 2 takes care of lifting the restriction on the root edge to point away from the origin. This finishes the bijective proof of Theorem 3.

3.5 Graph distance statistics in Boltzmann maps

The tree bijection not only provides insights into the structure of the generating functions, but also provides an effective way of studying refined enumeration where we have control on graph distances.

This insight has been crucial in the study of scaling limits of metric structure of large random planar maps, culminating in the proofs of Gromov–Hausdorff convergence towards the Brownian sphere [10, 14] by Le Gall and Miermont.

Suppose $t, q_1, q_2, \ldots \geq 0$ are chosen small enough such that $R(t,q) < \infty$, then according to (25),

$$\mathbb{P}(\mathfrak{m}) = \frac{1}{R} t^{|V(\mathfrak{m})|-1} \prod_{f \in F(\mathfrak{m})} q_{\deg f/2} \tag{29}$$

determines a probability measure on the set $\vec{M}_0^\bullet \cup \{\bullet\}$ of rooted pointed bipartite maps (including for convenience a degenerate map $\bullet$ with weight t). We call this random map the *pointed* (t,q)*-Boltzmann planar map*. According to Theorem 4, the associated random mobile $\mathfrak{t} = \Phi(\mathfrak{m})$ is distributed as

$$\mathbb{P}(\mathfrak{t}) = \frac{1}{R} t^{|V_{\text{white}}(\mathfrak{t})|} \prod_{v \in V_{\text{black}}(\mathfrak{t})} q_{\deg v}. \tag{30}$$

Let $\mathfrak{t}^\circ$ denote the bicolored tree obtained from the mobile $\mathfrak{t}$ by forgetting the labels on the white vertices. By the same reasoning as in (27), this mapping is

$$\prod_{v \in V_{\text{black}}(\mathfrak{t})} \binom{2\deg v - 1}{\deg v} \quad \text{to} \quad 1. \tag{31}$$

Hence $\mathfrak{t}^\circ$ is distributed among the all (properly bicolored) plane trees as

$$\mathbb{P}(\mathfrak{t}^\circ) = \frac{1}{R} t^{|V_{\text{white}}(\mathfrak{t}^\circ)|} \prod_{v \in V_{\text{black}}(\mathfrak{t}^\circ)} \binom{2\deg v - 1}{\deg v} q_{\deg v}. \tag{32}$$

This may be recognized as the law of a *two-type Bienayme–Galton–Watson tree* with appropriately chosen offspring distributions for the white and black vertices.

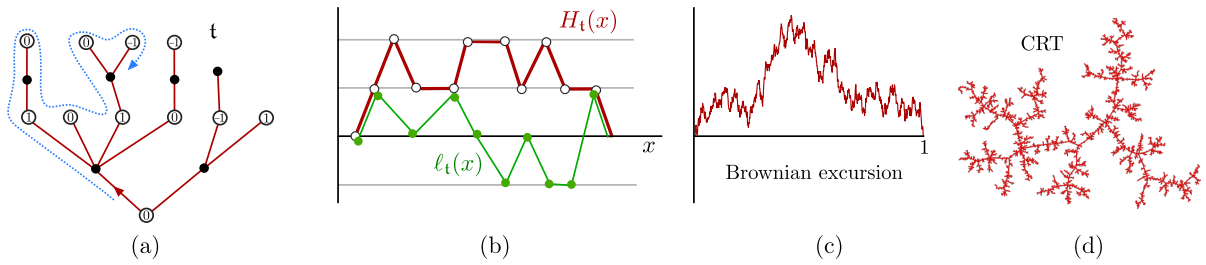

Figure 4: Illustration of the height and label function (b) associated to the mobile (a). The scaling limit of the height function is the Brownian excursion (c), which is the height function of the CRT (d).

It is natural to ask what this random tree looks like when conditioned to have n white vertices as $n \to \infty$. Under mild conditions on the weights q such random trees admit a scaling limit described by the *Continuous Random Tree* (CRT) of Aldous [1]. To be more precise, one may associate to a tree $\mathfrak{t}^\circ$ with n white vertices a *height function* $H_{\mathfrak{t}^\circ} : [0,n] \to \mathbb{R}$ by setting $H_{\mathfrak{t}^\circ}(i)$ for $i = 0, \ldots, n$ to be the height of the ith white vertices (in the order encountered in the contour), where the height is (half) the graph distance within the tree to the root vertex. Then by the invariance principle of Marckert & Miermont

[12], there exists a $c > 0$ such that we have the convergence in distribution (with respect to the uniform topology on $[0,1]$)

$$\left(\frac{1}{c\sqrt{n}} H_{\mathfrak{t}^\circ}(x/n)\right)_{0\le x\le 1} \xrightarrow[n\to\infty]{\text{(d)}} \text{Brownian excursion on } [0,1]. \tag{33}$$

The right-hand side is the height function of the CRT. In particular, $\mathfrak{t}^\circ$ is typically of height $\sqrt{n}$.

Conditionally on $\mathfrak{t}^\circ$ the labels of $\mathfrak{t}$ can be understood as a type of random walk along the branches of the tree. We should therefore expect the typical range of the labels to be of order $\sqrt{\sqrt{n}}$, since the height is of order $\sqrt{n}$. Indeed, the aforementioned invariance principle [12] implies that the *label function* $\ell_{\mathfrak{t}} : [0,n] \to \mathbb{R}$, defined similarly as the height function but recording the label at each white vertex, converges when normalized by $n^{-1/4}$, together with the height function to a continuous process known as the *Brownian snake* [8],

$$\left(\frac{1}{c\sqrt{n}} H_{\mathfrak{t}}(x/n), \frac{1}{c' n^{1/4}} \ell_{\mathfrak{t}}(x/n)\right)_{0\le x\le 1} \xrightarrow[n\to\infty]{\text{(d)}} \text{Brownian snake on } [0,1]. \tag{34}$$

Even though the label function only records distances to the origin, for large random maps it contains sufficient information to approximate distances between arbitrary pairs of points. This was utilized by Le Gall and Miermont [10, 14] to prove that the discrete metric space arising from the graph distance on the vertices of a (t,q)-Boltzmann planar map with a large number n of vertices converges in distribution (with respect to the Gromov-Hausdorff topology)

$$\left(V(\mathfrak{m}), n^{-1/4} d_{\text{graph}}\right) \xrightarrow[n\to\infty]{\text{(d)}} \text{Brownian sphere}. \tag{35}$$

The *Brownian sphere* or *Brownian map* [13] is a random metric space that can be naturally constructed from the Brownian snake. It has the topology of the 2-sphere but a geometry that is truly fractal [9], as exemplified by having a Hausdorff dimension of 4.

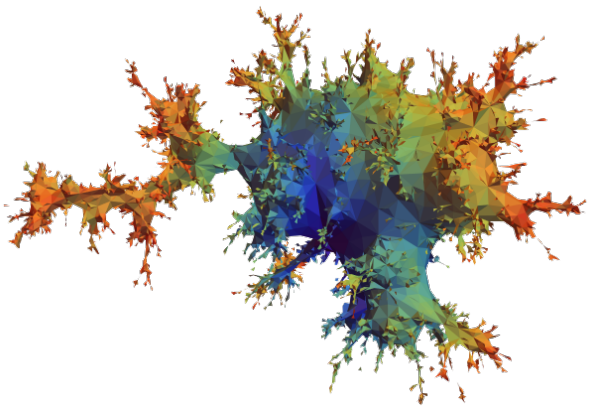

Figure 5: Simulation of the Brownian sphere.

4 Tree bijection for hyperbolic surfaces

Based on the similarity between the generating functions of Weil-Petersson volumes (Theorem 2) and planar maps (Theorem 3), as well as the interpretation of the latter in terms of trees, the goal of this

section should be clear. Can we encode hyperbolic surfaces with two distinguished cusps, whose generating function is given by $R = \frac{\delta^2 F_0^{\mathrm{WP}}}{\delta q(0)^2}$, in terms of a family of trees such that $Z^{\mathrm{WP}}(R) = 0$ is the characteristic equation associated to these trees?

In analogy with the planar map case, we should expect that one of the cusps plays the role of the origin, from which we measure distances, and the other will be the root of the tree. It turns out that such a tree can be identified in a hyperbolic surface with the help of a variant of the spine construction of Bowditch & Epstein [4] and Penner [17].

4.1 Spine of a hyperbolic surface

Let $X \in \mathcal{M}_{0,1+n}(0, L)$ be a genus-0 hyperbolic surface with a cusp, that we call the *origin*, and $n \geq 2$ additional boundaries of length $L_1, \ldots, L_n$ (or cusps whenever $L_i = 0$). The universal cover of X can be represented as a convex domain P in the Poincaré upper-half plane $\mathbb{H}$. Denote by Γ the Fuchsian group $\Gamma \subset \mathrm{PSL}(2, \mathbb{R})$ such that $X = P/\Gamma$. It is convenient to consider the extended surface $\check{X} = \mathbb{H}/\Gamma \supset X$, corresponding to the surface X in which to each geodesic boundary of positive length we glue a *funnel* with a geodesic boundary of the same length.

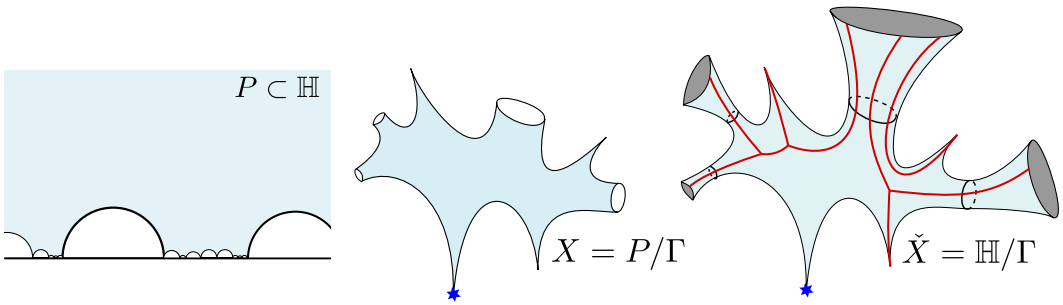

Figure 6: Illustration of the universal cover P of the surface X. Note that a boundary geodesic of X lifts to a geodesic side of P. The surface $\check{X} = \mathbb{H}/\Gamma$ is extended with funnels. The spine of $\check{X}$ is illustrated in red.

For any point $x \in X$ we can make sense of the shortest geodesics from x to the origin as follows. Let c be a horocycle around the origin that separates x from the cusp and $d(x, c)$ the hyperbolic distance between x and c. Let $w(x)$ be the number of distinct geodesics from x to c of length $d(x, c)$. Since each of these geodesics meets c perpendicularly and thus continues into the cusp, the extended geodesics and their number $w(x)$ is independent of the choice of horocycle c.

The *spine* Σ of $\check{X}$ is defined as the subset of points with at least two shortest geodesics to the root, i.e.

$$\Sigma = \{x \in \check{X} : w(x) \geq 2\} \subset \check{X}. \tag{36}$$

The points with three or more geodesics will be called *internal vertices*

$$V = \{x \in \check{X} : w(x) \geq 3\} \subset \Sigma. \tag{37}$$

The following (informally formulated) lemma is a variant of [4, Lemma 2.2.1].

Lemma 5. *The spine of* Σ *satisfies the following properties.*

1. *V is a finite set.*
2. *$\Sigma \setminus V = \{x \in \check{X} : w(x) = 2\}$ consists of a finite union of open geodesic arcs. Tracking each arc in each direction it either ends after finite distance at a point in V or extends indefinitely into a cusp or funnel.*
3. *Each point in $x \in V$ is the endpoint of exactly $w(x) \geq 3$ arcs.*
4. *Each boundary or cusp (except the origin) is the endpoint of at least one arc.*
5. *The compactified spine $\overline{\Sigma}$, obtained by joining all arcs running into the same funnel or cusp with a point at infinity (a* boundary vertex*), is a tree.*

Sketch of the proof. Let us fix a horocycle c around the origin that is short enough not to self-intersect or touch any of the funnels (by the collar lemma, the unit-length horocycle satisfies this). Then c does not intersect the spine Σ and it lifts to a countable collection $\mathcal{C}$ of disjoint horocycles in $\mathbb{H}$ (or even disjoint horodisks, because the horocycles cannot nest). Let $x \in \check{X}$ be a point separated from the origin by c and $y \in \mathbb{H}$ a lift of x. Then $w(x)$ counts the number of horocycles in $\mathcal{C}$ that are closest to x, say at distance r. The number $w(x)$ is finite, since every hyperbolic disk, in particular the ball of radius r around y, can meet only finitely many disjoint horodisks.

Let $\varepsilon > 0$ be such that there are still only $w(x)$ horocycles in $\mathcal{C}$ within distance $r + 2\varepsilon$ of y. Then the shape of the spine can be established within an ε-neighborhood of x, i.e. $\Sigma \cap \mathrm{Ball}(x, \varepsilon)$. This is particularly conveniently seen in the Poincaré disk $\mathbb{D}$ when y is put at the origin (Fig. 7). A compactness argument then rather straightforwardly leads to the first three stated properties.

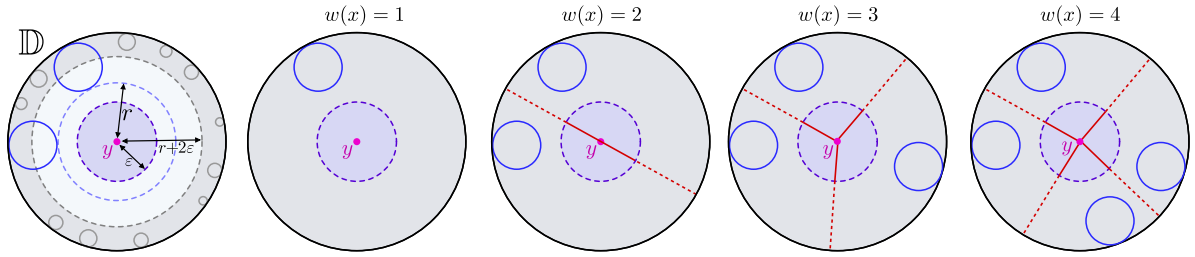

Figure 7: The first figure illustrates the collection $\mathcal{C}$ of disjoint horocycles. The blue ones are at closest distance r from y, while the gray ones are at distance at least $r + 2\varepsilon$. The other figures illustrate the spine neighbourhoods $\Sigma \cap \mathrm{Ball}(x, \varepsilon)$ for different number $w(x)$.

For the complement of the spine $\check{X} \setminus \Sigma = \{x \in \check{x} : w(x) = 1\}$ we obtain a retraction onto the horodisk around the origin, by tracing the unique shortest geodesics to the origin. Hence, $\check{X} \setminus \Sigma$ is topologically an open punctured disk. This in turn implies that the compactified spine $\overline{\Sigma}$ is connected and simply-connected, so must be a tree. □

We denote the combinatorial plane tree encoding the incidence relations of the compactified spine $\overline{\Sigma}$ by t. It has n white vertices of arbitrary degree and zero or more red (internal) vertices of degree at least three.

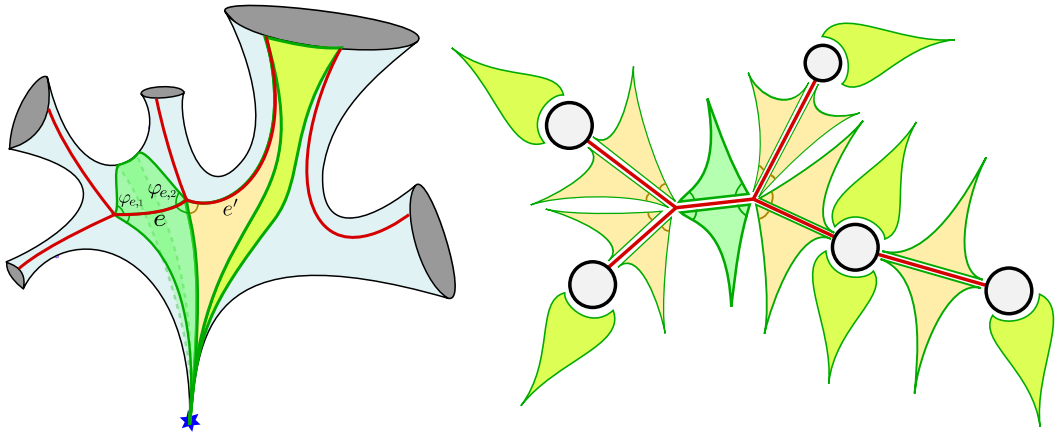

Figure 8: Example of two triangles (in green) associated to the finite-length arc e. One of the triangles (in orange) for the half-infinite arc e' is shown. The green-yellow region indicates a hyperbolic wedge. The right figure shows schematically how the triangles and wedges are arranged in the contour of the tree $\mathfrak{t}$.

4.2 A canonical triangulation of X

With the help of the spine Σ we can now canonically triangulate $\check{X}$. To each arc e we associate two triangles: each $x \in e$ has two shortest geodesics to the origin, one on the left and one on the right (assuming we have chosen an orientation of e). The union of all such geodesics on the left when x ranges over the arc determines a triangle, which has one side along the arc and two sides made by geodesics to the origin. Similarly we get a triangle on the right. Since the geodesics to the left and right are equally long, it should be clear that the two triangles are mirror copies of each other. We denote the angles of the triangles at the ends of e by $\varphi_{e,1}, \varphi_{e,2} \in [0, \pi)$, which can vanish (corresponding to an ideal vertex) in case the respective end of e runs into a cusp or funnel. Note that the angle at the origin is necessarily 0, so it is a hyperbolic triangle of area

$$\pi - \varphi_{e,1} - \varphi_{e,2} \in (0, \pi). \tag{38}$$

The pairs of triangles associated to the arcs of the spine together nearly triangulate the whole $\check{X}$. What is left is a collection of *ideal wedges*, i.e. regions isometric to $\{x + iy : 0 < x < 1, y > 0\} \subset \mathbb{H}$, extending to the boundary at infinity in the hyperbolic cylinders (see the green-yellow region in Figure 8). An ideal wedge is

Knowledge of the angles $\varphi_{e,i}$ specifies uniquely the shapes of constituent triangles and the tree structure $\mathfrak{t}$ indicates how the sides of the triangles and wedges are to be glued pairwise. The geodesic shared between two triangles is half-infinite and can thus be glued unambiguously. The sides shared by a triangle and a wedge, however, is infinite so there is a shift degree of freedom that we would like to capture.

To this end let us zoom in on the neighbourhood of a funnel, see Figure 9 for an example where two arcs enter the funnel. Viewed in the universal cover, we can position the lift of the closed geodesic b of length L along the imaginary axis, such that the Möbius transformation associated with circling around the boundary becomes $z \longrightarrow e^L z$. In case there are $k \geq 1$ arcs running into the funnel, we can identify k pairs of triangles (in orange) and k ideal wedges (in green-yellow). Note that the geodesics that

separate a wedge from a triangle always intersect the imaginary axis and form a zigzagging pattern of semicircles (right-hand side of Figure 9). This pattern is uniquely described by recording the Euclidean radii of the semicircles, or equivalently the log-ratios of consecutive radii, which can be interpreted as measuring hyperbolic distances along b.

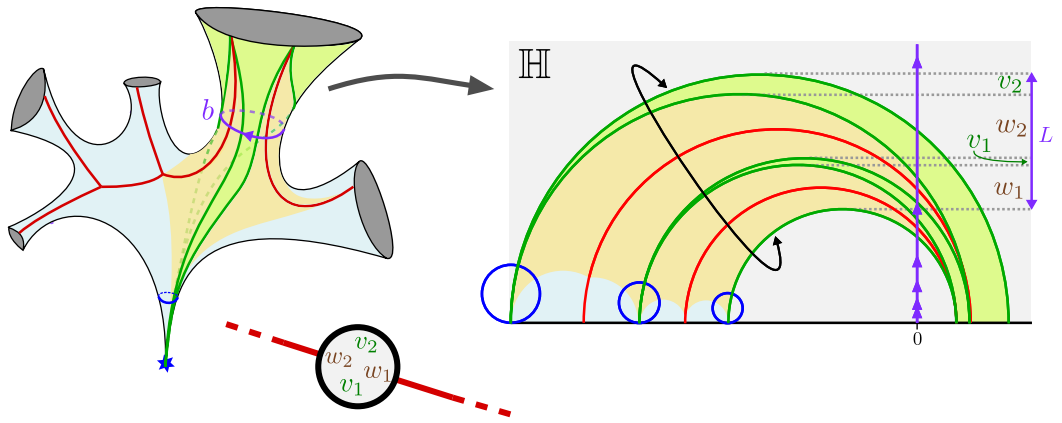

Figure 9

Since the geodesic b is partitioned into intervals of length $v_1, w_1, \ldots, v_k, w_k$ we thus find that

$$\sum_{i=1}^{k} v_i + w_i = L. \tag{39}$$

It is an exercise to check that a consistent choice of horocycle around the origin imposes the further condition

$$\sum_{i=1}^{k} v_i = \sum_{i=1}^{k} w_i. \tag{40}$$

4.3 Tree bijection

We are now in a position to formulate the tree bijection. Denote by T_n the set of plane trees with n white vertices (carrying indices $1, \ldots, n$) and any number of red vertices of degree at least three. And for $\mathfrak{t} \in T_n$ we introduce the set of allowed label assignments

$$\mathcal{A}_{\mathfrak{t}}(L) = \{(\varphi_{e,1}, \varphi_{e,2})_{e \in E(\mathfrak{t})}, (v_{i,1}, w_{i,1}, \ldots, v_{i,k_i}, w_{i,k_i})_{i=1}^{n} : \text{constraints}\} \subset \mathbb{R}^{6n-6} \tag{41}$$

These constraints are

- $\varphi_{e,1}, \varphi_{e,2} \in [0, \pi)$ with $\varphi_{e,i} = 0$ iff it is incident to a white vertex;
- $\varphi_{e,1} + \varphi_{e,2} < \pi$;
- $\sum_e \varphi_{e,1} = \pi$ where the sum runs over edges meeting at an internal vertex;
- $\sum_{j=1}^{k_i} v_{i,j} = \sum_{j=1}^{k_i} w_{i,j} = L_i/2$ for $i = 1, \ldots, n$.

Some book keeping shows that $\mathcal{A}_{\mathrm{t}}(L)$ is a convex polytope of dimension

$$\dim \mathcal{A}_{\mathrm{t}}(L) = 2n - 4 - \sum_{v \in V_{\mathrm{internal}}(\mathrm{t})} (\deg v - 3). \tag{42}$$

Theorem 6 (Tree bijection). *For $n \geq 2$ and $L_1, \ldots, L_n > 0$ the spine construction determines a bijection*

$$\mathcal{M}_{0,n+1}(0, L) \to \bigsqcup_{\mathrm{t} \in T_n} \mathcal{A}_{\mathrm{t}}(L). \tag{43}$$

4.4 Weil-Petersson volume form

To be continued.

References

[1] D. Aldous, *The continuum random tree. I*, Ann. Probab., 19 (1991), pp. 1–28.

[2] J. Ambjørn, B. Durhuus, and T. Jonsson, *Quantum Geometry: A Statistical Field Theory Approach*, Cambridge Monographs on Mathematical Physics, Cambridge University Press, 1997.

[3] J. Bouttier, P. Di Francesco, and E. Guitter, *Planar maps as labeled mobiles*, Electron. J. Combin., 11 (2004), pp. Research Paper 69, 27 pp. (electronic).

[4] B. Bowditch and D. Epstein, *Natural triangulations associated to a surface*, Topology, 27 (1988), pp. 91–117.

[5] B. Eynard, *Counting surfaces*, Progress in Mathematical Physics, 70 (2016).

[6] C. Itzykson and J.-B. Zuber, *Combinatorics of the modular group. II. The Kontsevich integrals*, Internat. J. Modern Phys. A, 7 (1992), pp. 5661–5705.

[7] M. Kontsevich, *Intersection theory on the moduli space of curves and the matrix Airy function*, Comm. Math. Phys., 147 (1992), pp. 1–23.

[8] J.-F. Le Gall, *Spatial branching processes, random snakes and partial differential equations*, Lectures in Mathematics ETH Zürich, Birkhäuser Verlag, Basel, 1999.

[9] ——, *The topological structure of scaling limits of large planar maps*, Invent. Math., 169 (2007), pp. 621–670.

[10] ——, *Uniqueness and universality of the Brownian map*, Ann. Probab., 41 (2013), pp. 2880–2960.

[11] Y. I. Manin and P. Zograf, *Invertible cohomological field theories and Weil-Petersson volumes*, Ann. Inst. Fourier (Grenoble), 50 (2000), pp. 519–535.

[12] J.-F. Marckert and G. Miermont, *Invariance principles for random bipartite planar maps*, Ann. Probab., 35 (2007), pp. 1642–1705.

[13] J.-F. Marckert and A. Mokkadem, *Limit of normalized quadrangulations: the Brownian map*, Ann. Probab., 34 (2006), pp. 2144–2202.

[14] G. Miermont, *The Brownian map is the scaling limit of uniform random plane quadrangulations*, Acta Math., 210 (2013), pp. 319–401.

[15] M. Mirzakhani, *Simple geodesics and Weil-Petersson volumes of moduli spaces of bordered Riemann surfaces*, Invent. Math., 167 (2007), pp. 179–222.

[16] ——, *Weil-Petersson volumes and intersection theory on the moduli space of curves*, J. Amer. Math. Soc., 20 (2007), pp. 1–23.

[17] R. Penner, *The decorated Teichmüller space of punctured surfaces*, Comm. Math. Phys., 113 (1987), pp. 299–339.

[18] W. T. Tutte, *On the enumeration of planar maps*, Bull. Amer. Math. Soc., 74 (1968), pp. 64–74.

[19] E. Witten, *Two-dimensional gravity and intersection theory on moduli space*, in Surveys in differential geometry (Cambridge, MA, 1990), Lehigh Univ., Bethlehem, PA, 1991, pp. 243–310.