6

Teichmüller spaces

Now we introduce the main actors of this book: Teichmüller spaces. Although the applications we will consider mainly use finite-dimensional Teichmüller spaces, we discuss them in full generality. This makes our job somewhat harder; we do it because we hope that theorems involving finite-dimensional Teichmüller spaces will have analogs for infinite-dimensional Teichmüller spaces. In particular, we hope that Thurston's theorem on the topological characterization of rational functions might be extended to mappings that are not postcritically finite. Indeed, in one case, David Brown [22] has proved such a result.

Thus we will treat finite-dimensional Teichmüller spaces, associated to Riemann surfaces of finite type, as a special case of infinite-dimensional Teichmüller spaces, associated to general Riemann surfaces.

REMARK This view, mainly represented by the work of Ahlfors and Bers, is quite analytical. The alternative would be to see finite-dimensional Teichmüller spaces as moduli spaces of compact complex curves, generalizing to moduli spaces of higher-dimensional compact complex manifolds, for instance surfaces of general type. This view was championed by Grothendieck [51], who used techniques from complex analytic geometry and algebraic geometry, and also by Earle and Eells; their paper [37] is still probably the best place to start learning the theory.

The situation is like that of $SL_2(\mathbb{Z})$, which can be viewed as either the genus one case of Teichmüller modular groups or as the first of the sequence $SL_2(\mathbb{Z}), SL_3(\mathbb{Z}), \ldots$. These two views diverge rapidly and lead to quite different descriptions of $SL_2(\mathbb{Z})$. Similarly, the two views of Teichmüller theory lead to quite different treatments of finite-dimensional Teichmüller spaces, reflecting which constructions one wants to be able to carry over to the more general setting.

I used to favor the Grothendieck-Earle-Eells approach. In [59] I gave a construction inspired by this view, using smooth, almost complex structures and the Serre duality theorem, and never mentioning quasiconformal mappings. The Bers simultaneous uniformization theorem, key to the Bers approach, seemed to me unnatural and even unpalatable; I could not see why anyone would ever want this result. Sullivan's no wandering domains theorem showed me that I was wrong; I have come to see that the simultaneous uniformization theorem is essential in proving Thurston's hyperbolization theorem for 3-manifolds that fiber over the circle, presented in volume 2. Bers's theorem still seems unnatural to me, just as the paintings of Hieronymus Bosch seem unnatural. But I have come to see beauty as well as utility in an approach that first seemed to me simply horrible.

6.1 Quasiconformal surfaces

A Teichmüller space is the set of Riemann surfaces of a given quasiconformal type. There is one Teichmüller space for every quasiconformal surface: we speak of the "Teichmüller space modeled on S", where S is a quasiconformal surface. This requires knowing what a quasiconformal surface is.

A quasiconformal surface S is a topological surface with a Riemann-surface structure; two Riemann surface structures on S define the same quasiconformal structure if the identity map between them is quasiconformal. If S_1, S_2 are two quasiconformal surfaces, a map $f : S_1 \to S_2$ is quasiconformal if it is a quasiconformal homeomorphism for one, hence all, analytic structures on each of S_1 and S_2. In particular, by definition all quasiconformal maps are isomorphisms.

If X is a Riemann surface, we denote by $\mathrm{qc}(X)$ its equivalence class. By Rado's theorem, all connected quasiconformal surfaces are σ-compact.

For compact surfaces, a quasiconformal structure carries little information.

Proposition 6.1.1 *If two compact quasiconformal surfaces S_1 and S_2 are homeomorphic, then they are isomorphic as quasiconformal surfaces.*

Proof We may take $S_1 = \mathrm{qc}(X_1)$ and $S_2 = \mathrm{qc}(X_2)$. In dimension 2, homeomorphic differentiable surfaces are diffeomorphic (for compact surfaces, this follows from the classification of surfaces), so X_1 and X_2 are diffeomorphic, and on a compact surface a diffeomorphism is quasiconformal. □

Proposition 6.1.1 is wildly wrong for noncompact surfaces. Already $\mathbb{C}$ and $\mathbf{D}$ are homeomorphic, but not isomorphic as quasiconformal surfaces (see Exercise 4.3.7). More generally, the quasiconformal surface gotten by removing a point from a compact Riemann surface and the quasiconformal surface gotten by removing a disc from the same surface are homeomorphic, but they are not isomorphic as quasiconformal surfaces. But the situation can get much wilder: when the fundamental group of a surface is infinitely generated, there are *uncountably many* distinct quasiconformal surfaces that are homeomorphic.

Example 6.1.2 Let Z be $\{0, 1, 2, 3, \dots\}$. Then there are uncountably many different quasiconformal surfaces all homeomorphic to $\mathbb{C} - Z$. Figure

6.1.1 shows how to construct one such surface. Since (Theorem 3.5.8) we can make the lengths $l_1, l_2, \dots$ anything we like, there are uncountably many such surfaces.

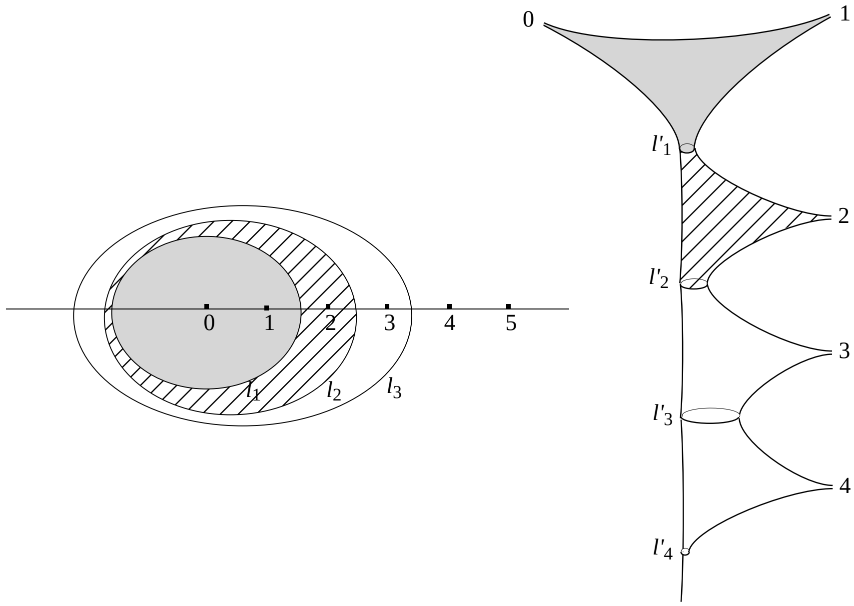

FIGURE 6.1.1 How to construct one quasiconformal surface homeomorphic to $\mathbb{C}-Z$. The shaded oval at left is homeomorphic to the trouser at top right, with cuffs of length 0 at 0 and 1 and a waist of length l'_1. The surface drawn with slanted lines has boundary l_1 and l_2 and a puncture point at 2; it is homeomorphic to the second trouser at right, which we may think of as having cuffs of lengths l'_1 and l'_2 and a waist (at the puncture point 2) of length 0 .... (Note that although at left we draw the lengths $l_1, l_2, \dots$ using the Euclidean metric, so that $l_1 < l_2 < l_3 \dots$, these geodesics are really with respect to the hyperbolic metric of $\mathbb{C}-Z$; they are all more or less the same length.) This gives a recipe for creating a quasiconformal surface topologically identical to $\mathbb{C}-Z$. Since we can make the lengths $l'_1, l'_2, \dots$ whatever we like, we can create uncountably many such surfaces. $\triangle$

Beltrami forms on quasiconformal surfaces

We now need to define the *space of Beltrami forms* on a quasiconformal surface S. It is tempting to define this as the unit ball in $L^\infty_*(TS, TS)$ (see Definition 4.8.11 and equation 4.8.18). But this does not work in any natural way. The problem is that S is not naturally a C^1 manifold, so it doesn't have a tangent bundle TS. It does have a "tangent bundle almost everywhere", and since Beltrami forms are only defined almost everywhere, this is good enough. However, setting up the machinery to make this precise takes more effort than the dodge we will adopt.

Definition 6.1.3 (Beltrami form on a quasiconformal surface) A *Beltrami form on a quasiconformal surface* S is represented by a pair $\big((\varphi : S \to X), \mu\big)$ where X is a Riemann surface, φ is an isomorphism $S \to \mathrm{qc}(X)$, and μ is a Beltrami form on X, i.e., $\mu \in \mathcal{M}(X)$. Two pairs $\big((\varphi_1 : S \to X_1), \mu_1\big)$ and $\big((\varphi_2 : S \to X_2), \mu_2\big)$ represent the same element of the space $\mathcal{M}(S)$ of Beltrami forms on S if

$$\mu_2 = (\varphi_1 \circ \varphi_2^{-1})^* \mu_1. \tag{6.1.1}$$

This is just a disguised way of "identifying" $\mathcal{M}(S)$ with $\mathcal{M}(X)$, as the following statement makes clear.

Proposition and Definition 6.1.4 (Analytic structure on the space of Beltrami forms)

1. *Let S be a quasiconformal surface, X a Riemann surface, and $\varphi : S \to \mathrm{qc}(X)$ an isomorphism of quasiconformal surfaces. Then the mapping $\mathcal{M}(X) \to \mathcal{M}(S)$ given by*

$$\mu \mapsto \big((\varphi : S \to X), \mu\big) \tag{6.1.2}$$

is bijective.

2. *If we make $\mathcal{M}(S)$ into a Banach analytic manifold by requiring that the identification 6.1.2 be an isomorphism, then this structure is independent of the choice of $\varphi : S \to \mathrm{qc}(X)$.*

Proof 1. By Definition 6.1.3, we know that μ_1 and μ_2 map to the same point if $(\varphi \circ \varphi^{-1})^* \mu_1 = \mu_2$, which evidently means $\mu_1 = \mu_2$. This shows injectivity. For surjectivity, suppose $m \in \mathcal{M}(S)$ is represented by $\big((\varphi_1 : S \to X_1), \mu_1\big)$ for some φ_1, X_1, μ_1. Then it is also represented by $\big((\varphi : S \to X), (\varphi \circ \varphi_1^{-1})^* \mu_1\big)$.

2. Again using Definition 6.1.3, we need to know that

$$(\varphi \circ \varphi_1^{-1})^* : \mathcal{M}(X_1) \to \mathcal{M}(X) \tag{6.1.3}$$

is an analytic isomorphism. That is the content of Proposition 4.8.17. □

Remark If $\mathcal{M}(S)$ is just $\mathcal{M}(X)$ in light disguise, why bring it in at all? The reason is that $\mathcal{M}(X)$ has a distinguished point (the point 0); $\mathcal{M}(S)$ does not. When we work in $\mathcal{M}(X)$, we are studying complex structures where a particular *background* complex structure has been chosen, namely, that of X. When we work in $\mathcal{M}(S)$, we are working with the same complex structures but there is now no distinguished background structure. Whenever we make an argument about $\mathcal{M}(X)$ that is really about $\mathcal{M}(S)$, we need to show that the argument does not depend on the choice of base point, i.e., the background complex structure.

However, this advantage of $\mathcal{M}(S)$ over $\mathcal{M}(X)$ is phony: we won't be able to avoid making these arguments. For instance, part 2 of Proposition and Definition 6.1.4 asserts that $\mathcal{M}(S)$ has a complex structure that does not depend on any background structure, but we had to go back to Proposition 4.8.17 to prove it. But we will be able to avoid referring to the base point or background structure in the *statements*; I hope this results in conceptual clarification. △

Ideal boundaries of quasiconformal surfaces

Recall (Proposition and Definition 3.7.1) the definition of the ideal boundary of a hyperbolic Riemann surface. Quasiconformal surfaces also have ideal boundaries. It follows from Proposition 6.1.5 that the ideal boundary of a Riemann surface X depends only on the underlying quasiconformal surface: every quasiconformal surface $S = \mathrm{qc}(X)$ has ideal boundary $I(S) = I(X)$.

Proposition and Definition 6.1.5 (Ideal boundary of a quasiconformal surface)

1. *If X and Y are Riemann surfaces and $f : X \to Y$ is quasiconformal, then f extends to a homeomorphism $\overline{f} : \overline{X} \to \overline{Y}$.*
2. *If S is a quasiconformal surface and X is a Riemann surface such that $S = \mathrm{qc}(X)$, then the ideal boundary of S is $I(S) = I(X)$. If Y is another Riemann surface such that $S = \mathrm{qc}(Y)$, then there is a quasiconformal mapping $X \to Y$, which by part 1 induces a homeomorphism $I(X) \to I(Y)$, so that $I(S) = I(Y)$, and the ideal boundary is well defined.*

Proof 1. This follows from Proposition 4.9.1. Let $\widetilde{X}$ and $\widetilde{Y}$ be the universal covering spaces of X and Y. Choose isomorphisms $\varphi_X : \widetilde{X} \to \mathbf{D}$ and $\varphi_Y : \widetilde{Y} \to \mathbf{D}$. There are then Fuchsian groups Γ_X, Γ_Y such that φ_X, φ_Y induce isomorphisms $X \to \mathbf{D}/\Gamma_X$ and $Y \to \mathbf{D}/\Gamma_Y$. The homeomorphism f lifts to a quasiconformal homeomorphism $\tilde{f} : \mathbf{D} \to \mathbf{D}$ with the property that $\Gamma_Y \tilde{f} = \tilde{f}\Gamma_Y$. By Proposition 4.9.1, $\tilde{f}$ extends to a homeomorphism $\overline{\tilde{f}} : \overline{\mathbf{D}} \to \overline{\mathbf{D}}$. Moreover, $\tilde{f}$ maps the limit set of Γ_X to the limit set of Γ_Y, and induces a homeomorphism

$$\overline{f} : \left(\overline{\mathbf{D}} - \Lambda_{\Gamma_X}\right)/\Gamma_X \to \left(\overline{\mathbf{D}} - \Lambda_{\Gamma_Y}\right)/\Gamma_Y. \qquad 6.1.4$$