



Toward Computing Extremal Quasiconformal Maps via Discrete Harmonic Measured Foliations

Citation

Palmer, David R. 2016. Toward Computing Extremal Quasiconformal Maps via Discrete Harmonic Measured Foliations. Bachelor's thesis, Harvard College.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:39060036>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Acknowledgements

I am very lucky to have had two excellent advisors and collaborators in developing this thesis—Professors Steven J. Gortler and Dylan P. Thurston. I am supremely grateful to them for all the time they generously gave me, and for their insight and coaching. This thesis forms only a small part of a larger project we have been working on, and I hope we will be able to continue to work toward our overall objective of a mathematically rigorous treatment of discrete Teichmüller mapping.

I am also grateful to my family for supporting me at every stage of my education, and especially to my grandfather for his inspiring and infectious optimism.

Finally, I want to thank Li Murphy and Jen Guidera for encouraging me to finish this thesis despite my doubts.

Contents

Abstract	i
Acknowledgements	ii
Contents	iii
1 Introduction	1
1.1 Our Approach	2
1.2 Related Work	3
1.2.1 Measured Foliations	3
1.2.2 Discrete Exterior Calculus	4
1.2.3 Extremal Quasiconformal Parameterization	5
1.2.4 Other Work	6
2 Smooth Theory	7
2.1 Quasiconformal Maps	7
2.1.1 Conformal Maps	7
2.1.2 Dilatation	8
2.2 Quadratic Differentials and Teichmüller's Theorem	10
2.2.1 Teichmüller's Theorem	11
2.3 Measured Foliations	12
2.3.1 Whitehead Classes	15
2.3.2 Dirichlet Energy and the Hubbard-Masur Theorem	16
3 Discrete Theory	18
3.1 The Index Theorem	19
3.2 Paths and Whitehead Equivalence	22
3.3 Harmonicity	25
3.4 Discrete Dirichlet Energy	26
4 Implementation	31
4.1 Discrete Whitehead Moves	31
4.2 Main Algorithm	33
4.2.1 Initialization	34
4.3 Convergence	36

4.4	Experimental Validation	36
4.5	Conclusion	38
A	Rendering Measured Foliations	40
	Bibliography	42

Dedicated to the memory of my grandmother, Hilda Gage.

Chapter 1

Introduction

Computing mappings between surfaces is one of the most fundamental problems in computer graphics and computational geometry. When one of the surfaces is a planar domain, this is called parameterization. Most prominently, it is used to apply textures to surfaces. Given a parameterized surface, a texture can be represented as a rectangular image covering the parameterization domain. The variation in the texture value (e.g., color) is captured by a uniform grid of square pixels (“texels”). The parameterization then determines which texels to look up when rendering any given point on the surface. More generally, maps between pairs of surfaces are important in texture transfer, deformation transfer, and other applications.

There are several important characteristics one might try to optimize in computing a map for applications such as texturing. One of the most important is injectivity. If a parameterization maps distinct patches of surface to overlapping regions of the plane, then it is impossible to texture these patches independently. Parameterization techniques including conformal and extremal quasiconformal mapping have been designed in part to ensure that generated maps will be injective regardless of the geometry of the input surface.

In choosing between injective maps, finer-grained considerations like geometric distortion become important. Distortion is defined as the degree to which distances and angles change under a map. Suppose a square image is used to texture a curved surface. If a large region of the surface is compressed onto a small region of the square domain, the quality of the texture in that region will suffer as it will appear “pixelated.” A higher-resolution image might be required to compensate for the distortion even though the remainder of the surface might be textured at an acceptable resolution. In order to avoid such inefficiencies, it pays to minimize overall distortion.

It is a fact of differential geometry that a curved surface cannot be mapped to the plane without distortion (isometrically) [1]. Thus there are tradeoffs between different types of distortion one might try to minimize in generating a map. Conformal maps, for example, preserve angles (e.g., between curves), while equiareal maps preserve areas of regions. Exact preservation of both angles and areas is impossible, as it implies isometry [2], but conformal maps and equiareal maps are available between large classes of surfaces. In particular, the Riemann Mapping Theorem guarantees the existence of conformal maps between pairs of simply connected surfaces (topological disks and spheres). Unfortunately, conformal maps between surfaces of higher genus are typically unavailable. Even when mapping disks, imposing extra constraints on the map quickly makes conformal mapping impossible.

Partly as a response to these problems, a new approach to parameterization using quasiconformal maps has gained some currency. Quasiconformal maps preserve neither angles nor areas exactly. Rather, quasiconformal parameterization techniques seek to minimize distortion in the form of *dilatation*, which measures the degree to which small circles are stretched into ellipses under a map. Extremal quasiconformal maps, which minimize dilatation, have many favorable features, including inejectivity and smoothness [3]. Moreover, Teichmüller's Theorem guarantees the existence of extremal maps between surfaces, including those of higher genus, and describes them in terms of objects called holomorphic quadratic differentials. Rather than optimizing dilatation directly, approaches to quasiconformal parameterization generally arrive at a map indirectly by optimizing quadratic differentials or closely related objects called Beltrami differentials. Previous approaches to extremal quasiconformal parameterization have discretized quadratic differentials as discrete 1-forms (comprising a signed real number per edge) [3] or Beltrami differentials as functions defined per-face [4].

1.1 Our Approach

We propose to represent and compute holomorphic quadratic differentials indirectly via measured foliations. Thus our strategy is, broadly speaking

$$\text{Measured Foliations} \rightarrow \text{Holomorphic QDs} \rightarrow \text{Extremal Quasiconformal Maps}$$

In this thesis, we focus on the first stage of this process, namely, going from measured foliations to holomorphic quadratic differentials.

Intuitively, one can think of a measured foliation as a smooth pattern of stripes on a surface. A form of homotopy equivalence (Whitehead equivalence) on foliations divides

them into so-called Whitehead classes. Every holomorphic quadratic differential gives rise to a measured foliation, and this special (“harmonic”) foliation is unique in its Whitehead class. Moreover, Gardiner and Lakic [5, 6] have characterized this unique representative as the minimum of a Dirichlet-like energy in its Whitehead class.

We define a new object—a discrete measured foliation—comprising a nonnegative number on each mesh edge. By analogy to the technique used to discretize 1-forms, the edge values of our discrete foliation represent path integrals of a smooth foliation. Since foliations lack orientation, we do not store signs. Fortunately, the nonnegative values preserve enough information to define discrete notions of closedness, harmonicity, and Whitehead class faithful to their smooth counterparts.

Further, we have developed a framework for manipulating discrete foliations within their Whitehead classes. The Dirichlet energy of Gardiner and Lakic is easily translatable to the discrete setting. As in the smooth setting, the energy has a unique, harmonic minimum in each Whitehead class. This forms the basis of our main algorithm, which evolves a given foliation into the harmonic representative in its class. The algorithm converges to the same harmonic foliation regardless of how the initial representative is chosen within its Whitehead class, demonstrating that our discretization preserves the essential structure of the continuous theory.

In conclusion, the distinguishing features of our approach are:

- It preserves the essential features of the smooth theory of measured foliations and quadratic differentials, especially the natural structure of Whitehead equivalence classes.
- It maintains the data type distinctions between measured foliations, quadratic differentials, and 1-forms.
- Our algorithm converges to a globally unique element.

1.2 Related Work

1.2.1 Measured Foliations

Measured foliations are familiar in the mathematical study of Riemann surfaces (2-dimensional complex manifolds). They figured especially prominently in the work of W. Thurston on low-dimensional geometry [7].




FIGURE 1.1: A harmonic foliation on a double torus computed using our techniques.

In the continuous setting, the first statement and proof of the Hodge-like theorem for measured foliations that we use is due to Hubbard and Masur [8]. They actually prove a stronger result, namely that the map taking a holomorphic quadratic differential to the Whitehead class of its vertical foliation is a homeomorphism.

Wolf [9] gives another proof of Hubbard-Masur using global maps to \mathbb{R} -trees. Thinking of foliations as maps to \mathbb{R} -trees has inspired some of our thinking about the appropriate way to discretize foliations.

Gardiner and Lakic [5, 6] define the Dirichlet problem for measured foliations in the continuous setting and show that its solution, the minimum-energy foliation, corresponds to the unique holomorphic quadratic differential of Hubbard and Masur. This is the continuous analog of our Theorem 3.12 and has inspired our overall approach.

Fock and Goncharov [10] describe assigning edge-weights on a triangulation corresponding to a measured lamination (foliation) the same way we do in this thesis. However, they are interested in defining coordinates on the space of laminations, rather than discretizing and computing them.

1.2.2 Discrete Exterior Calculus

Our approach to discretizing measured foliations is also inspired by discrete exterior calculus (DEC), pioneered by Hirani [11] and also explicated by Crane et al. [12]. The program of discrete exterior calculus is to replace smooth objects with their integrals over simplices. Because de Rham cohomology and simplicial cohomology are equivalent, many theorems about smooth differential forms can be easily transferred to the discrete

setting. In particular, de Rham cochains (forms) can be approximated by simplicial cochains, preserving cohomology exactly.

Gu and Yau [13] use discrete 1-forms as an intermediate in computing parameterizations. Their method first solves a set of linear equations for a basis of holomorphic 1-forms on a surface of arbitrary topology. Then they integrate these 1-forms to arrive at conformal parameterizations. Because their maps are conformal, their method provides relatively little control over the generated result—in particular, there is no way to specify the image of the map or to impose specific point correspondences.

Gortler, Gotsman, and Thurston [14] use discrete 1-forms and index-counting arguments to prove that certain parameterizations of disks and tori are embeddings. With discrete measured foliations, it should be possible to prove an analogous result for surfaces of higher genus.

1.2.3 Extremal Quasiconformal Parameterization

Parameterization is a mature subfield of computer graphics, and we will not attempt to enumerate parameterization techniques exhaustively. We refer the reader to Floater and Hormann's [2] survey of various approaches to parameterization including those based on discrete conformal maps.

Extremal quasiconformal parameterization is a recent development. Weber, Myles, and Zorin [3] propose extremal quasiconformal maps as a favorable alternative to harmonic maps in fixed-boundary parameterization problems for which conformal maps are unavailable. They compute extremal maps via Teichmüller's Theorem, using alternating gradient descent to solve the Beltrami Equation. However, they only consider maps between punctured discs and annuli in the plane, and they discretize quadratic differentials as real 1-forms, eliding type distinctions that become more important on curved surfaces. Moreover, their algorithm comes without guarantees on convergence or uniqueness.

Ng, Gu, and Lui [4] take a different approach, optimizing a Beltrami coefficient and then generating a quasiconformal map from it. They use an iterative process of damping and projection onto the space of Teichmüller-type Beltrami coefficients to minimize the norm of the Beltrami coefficient, simultaneously updating the map itself to match the Beltrami coefficient by iteratively solving a linear system. Lui et al. [15] replace the flow used to update the map with a single linear system solve. Their algorithm, which they call Quasiconformal Iteration, is proven to converge on smooth Riemann surfaces [16].

1.2.4 Other Work

Another thread of study that bears comparison with ours is the literature on line fields and polyvector fields more generally. These are fields of multiple vectors invariant under a local group action (an involution in the case of line fields). Line fields in particular are closely related to measured foliations. Knöppel et al. [17] and Diamanti et al. [18] compute optimal fields using a Dirichlet-like energy, much as we do with measured foliations. However, they do not keep track of Whitehead classes as they are not aiming to compute holomorphic quadratic differentials.

Knöppel et al. [19] describe an algorithm for rendering vector fields or line fields as stripe patterns by computing consistent texture coordinates. Their method produces high-quality results even for fields which are not closed by optimizing a global energy. For our needs, this global optimization is unnecessary because we start with closed foliations. However, their method of rendering singularities within triangles proved useful.

Chapter 2

Smooth Theory

2.1 Quasiconformal Maps

Quasiconformal maps generalize conformal maps. Conformal maps are often characterized as angle-preserving. But in order to more naturally motivate quasiconformal maps, we begin with another characterization in terms of complex geometry.

2.1.1 Conformal Maps

Let D be an open subset of \mathbb{C} and consider a smooth map $f : D \rightarrow \mathbb{C}$. If f maps $z = x + iy \mapsto f(z)$, then its differential can be expressed as

$$df = \frac{\partial f}{\partial x} dx + \frac{\partial f}{\partial y} dy = \frac{\partial f}{\partial z} dz + i \frac{\partial f}{\partial \bar{z}} d\bar{z},$$

where

$$\frac{\partial}{\partial z} := \frac{\partial}{\partial x} - i \frac{\partial}{\partial y}, \quad \frac{\partial}{\partial \bar{z}} := \frac{\partial}{\partial x} + i \frac{\partial}{\partial y}$$

are the holomorphic and anti-holomorphic derivative operators, respectively. (We will write f_z and $f_{\bar{z}}$ for $\partial f / \partial z$ and $\partial f / \partial \bar{z}$.) f is holomorphic (equivalently, conformal) if it satisfies the Cauchy-Riemann equation:

$$\frac{\partial f}{\partial \bar{z}} = 0.$$

It is a fact of complex analysis that if f is holomorphic, then so is f_z , and by extension all higher derivatives of f .

The Riemann Mapping Theorem guarantees the existence of conformal maps between suitably nice domains. In particular, given any simply connected, open sets $C, D \subset \mathbb{C}$,

there exists a holomorphic bijection $f : C \rightarrow D$ [20]. This result explains the popularity and success of conformal mapping in computer graphics. Unfortunately, the theorem does not extend to multiply-connected domains. More importantly for applications, constraining f —such as by imposing specific correspondences

$$f(z_i) = w_i, \quad i = 1, \dots, n$$

—can make it impossible to find a conformal map, even between simply-connected domains. For example, Kahn [21] proves that if f maps the rectangle $[0, a] \times [0, 1]$ onto the rectangle $[0, b] \times [0, 1]$ conformally, preserving corresponding boundary segments, then $a = b$.

Let $f : R \rightarrow S$ be a map between Riemann surfaces. For any $p \in R$, let z be a coordinate around p and w be a coordinate around $f(p)$. f is locally conformal at p if $w \circ f \circ z^{-1}$ is holomorphic. f is conformal on R if it is locally conformal at every point. The Riemann mapping theorem implies the existence of a conformal map between any two simply-connected Riemann surfaces. But just as in the plane, it does not extend to multiply-connected domains.

2.1.2 Dilatation

Following Kahn [21], we define the dilatation of a diffeomorphism as a measurement of its failure to be conformal. Rather than measuring angle distortion, dilatation measures the distortion of infinitesimal circles into infinitesimal ellipses. A conformal map has a complex derivative, so is locally a similarity transformation. So we should expect conformal maps to have dilatation identically zero.

Formally, let $f : R \rightarrow S$ be an orientation-preserving diffeomorphism between closed Riemann surfaces, and let $z \in R$. The differential $df_z : T_z R \rightarrow T_{f(z)} S$ is linear and invertible. Think of $T_z R$ as a real 2-dimensional vector space, with basis $\partial/\partial x, \partial/\partial y$. Then df_z is represented by a 2×2 real matrix. By singular value decomposition, we may write

$$df_z = UAV,$$

where $U, V \in SO(2, \mathbb{R})$ and A is diagonal. Factoring out a scalar from A and allowing U and V to be similarity transformations, we can express the derivative in the form

$$df_z = UA_k V, \quad A_k = \begin{pmatrix} 1+k & 0 \\ 0 & 1-k \end{pmatrix}$$

for some positive real number k . Since U and V are similarity transformations, they map circles to circles. A_k maps the unit circle to an ellipse with semimajor axis $1+k$ and semiminor axis $1-k$. The **local dilatation** $\text{Dil}_z(f)$ at z is defined to be k . Equivalently,

Proposition 2.1.

$$\text{Dil}_z(f) = \left| \frac{f_{\bar{z}}}{f_z} \right|.$$

Proof. To see this, note that the linear transformation A_k can be expressed as

$$z = x + iy \mapsto (1+k)x + i(1-k)y = z + k\bar{z}.$$

As maps, U and V act as multiplication by complex scalars (say u and v), which is conformal. In a neighborhood of z , we can form the composed map $g : z \mapsto f(z/v)/u$. $dg_z = A_k$, so

$$\text{Dil}_z g = k = \left| \frac{g_{\bar{z}}}{g_z} \right|.$$

Now $f(z) = ug(vz)$, so by the chain rule

$$f_{\bar{z}} = ug_{\bar{z}}\bar{v} \quad f_z = ug_z v,$$

whence

$$\left| \frac{f_{\bar{z}}}{f_z} \right| = k \left| \frac{\bar{v}}{v} \right| = k = \text{Dil}_z f$$

□

Figure 2.1 illustrates the local dilatation of a map between planar regions.




FIGURE 2.1: The mapping illustrated above is quasiconformal. Lighter colors indicate greater local dilatation. Note that the dilatation is greatest where squares are mapped to elongated rectangles.

f is **quasiconformal** if $\text{Dil}_z(f)$ is bounded over M . In this case, the **dilatation** of f is

$$\text{Dil}(f) = \sup_{z \in R} \text{Dil}_z(f).$$

An extremal quasiconformal map is one with minimum dilatation within a homotopy class of maps. That is, given some diffeomorphism $h : R \rightarrow S$, we seek another diffeomorphism $f : R \rightarrow S$ homotopic to h , and such that for any other $g \sim h$,

$$\text{Dil}(f) \leq \text{Dil}(g).$$

It turns out that given some constraints on h , there is a unique such f , i.e., the inequality is strict. This is the content of Teichmüller's Theorem and the next section.

2.2 Quadratic Differentials and Teichmüller's Theorem

Teichmüller's Theorem not only guarantees the existence and uniqueness of extremal quasiconformal maps, but also characterizes them as *Teichmüller maps*, which have a specific form—in particular, constant local dilatation. Describing this form requires the notion of a quadratic differential. One can think of a quadratic differential as “locally the square of a 1-form”. More precisely,

Definition 2.2 (*cf.* Kahn, [21]). Let R be a Riemann surface. A **quadratic differential** is a map $\varphi : TR \rightarrow \mathbb{C}$ such that

$$\varphi(\lambda v) = \lambda^2 \varphi(v)$$

for any tangent vector $v \in TR$ and scalar $\lambda \in \mathbb{C}$.

Let $z : U \subset R \rightarrow \mathbb{C}$ be a coordinate patch. Then dz sends vectors in $TR|_U$ to complex numbers in $T\mathbb{C} \simeq \mathbb{C}$. We define the symbol $dz^2 : TR \rightarrow \mathbb{C}$ by

$$dz^2(v) = (dz(v))^2$$

for all $v \in TR|_U$. Note that dz^2 is a quadratic differential on U . Moreover, let φ be any quadratic differential on R , and let $v \in T_z R$ for some $z \in U$. There is a unique vector $1_z \in T_p R$ such that $dz(1_z) = 1$. Since $T_p R$ is a 1-dimensional complex vector space, we can write $v = \lambda 1_z$ for some $\lambda \in \mathbb{C}$. Then

$$\varphi(v) = \varphi(\lambda 1_z) = \lambda^2 \varphi(1_z) = dz^2(v) \varphi(1_z) = \varphi(z) dz^2(v),$$

so that

$$\varphi = \varphi(z)dz^2$$

on U , where the coefficient function is defined as $\varphi(z) = \varphi(1_z)$. We say φ is a **holomorphic quadratic differential** if its coefficient function $\varphi(z)$ is holomorphic on every coordinate patch.

A holomorphic quadratic differential φ has isolated zeros. Let $z_0 \in R$ be in the support of φ . Let (U, z) be a simply connected coordinate neighborhood of z_0 disjoint from the zeros of φ , and such that $\varphi = \varphi(z)dz^2$ on U . Then φ has a well-defined holomorphic square root

$$\sqrt{\varphi(z)}dz$$

on U . Since it is a holomorphic 1-form, we can integrate it path-independently to obtain a holomorphic function

$$\begin{aligned} \zeta : U &\rightarrow \mathbb{C} \\ \zeta(z) &= \int_{z_0}^z \sqrt{\varphi(\xi)}d\xi \end{aligned}$$

known as a **natural coordinate** on U . By definition, $\varphi = d\zeta^2$ on U . In other words, the coefficient of φ in the coordinate patch (U, ζ) is just unity [21].

In a neighborhood of a zero w of φ , the behavior of the quadratic differential is determined by its lowest residue, i.e., φ is locally of the form $\lambda z^n dz^2$ for some positive integer n and complex scalar λ .

2.2.1 Teichmüller's Theorem

Teichmüller's Theorem guarantees the existence and uniqueness of extremal quasiconformal maps. It also characterizes the form of an extremal map, which is of great importance for computing these maps:

Definition 2.3 ([21]). Let R and S be Riemann surfaces, and let $\varphi : TR \rightarrow \mathbb{C}$ and $\psi : TS \rightarrow \mathbb{C}$ be quadratic differentials. A **Teichmüller map** of dilatation k from R to S is a homeomorphism

$$f : R \rightarrow S$$

which maps the zeros of φ to the zeros of ψ , and such that for any $p \in R \setminus \{\text{zeros of } \varphi\}$, there is a natural coordinate ζ in a neighborhood of p and a natural coordinate ω in a neighborhood of $f(p)$ such that

$$\omega \circ f \circ \zeta^{-1}(z) = z + k\bar{z}. \tag{2.1}$$

In particular, since the maps ζ and ω are holomorphic, hence conformal, f has constant dilatation k . Moreover, f is smooth away from the zeros of φ .

Theorem 2.4 (Teichmüller’s Theorem [3, 21]). *Let R and S be Riemann surfaces and $h : S \rightarrow T$ an orientation-preserving homeomorphism. Then*

- (i) *there exists a unique Teichmüller map $f : S \rightarrow T$ homotopic to h ; and*
- (ii) *f is extremal, i.e., for any other map $g : S \rightarrow T$ homotopic to h and smooth almost everywhere, $\text{Dil}(g) \geq \text{Dil}(f)$.*

Teichmüller’s theorem suggests a path toward computing extremal quasiconformal maps: given surfaces R and S and a homeomorphism $h : R \rightarrow S$, find a pair of holomorphic quadratic differentials ψ on S and φ on R such that φ is the pullback of ψ under h . Then update h , φ , and ψ until they satisfy (2.1), with h in place of f .

2.3 Measured Foliations

Every smooth quadratic differential φ on a surface S defines a foliation of S as follows [21]. A vector $v \in T_z S$ is horizontal if $\varphi(v)$ is real and positive. When z is not a zero of φ , φ is of the form $\varphi(z)dz^2$ at z , and so the pullback of the positive real axis under φ is a line. Thus φ defines a smooth subbundle of the tangent bundle away from its zeros. The horizontal trajectories or *leaves* of the foliation are curves whose tangent vectors are everywhere horizontal.

For a holomorphic quadratic differential φ , let Z be the set of zeros of φ . $S \setminus Z$ has an open covering $\mathcal{U} = \{U_i\}$ and a natural coordinate $\zeta_i = x_i + iy_i$ for φ on each U_i . On U_i ,

$$\varphi = d\zeta_i^2.$$

The horizontal vectors are precisely those that map onto the real axis under $d\zeta_i$, i.e., multiples of $\partial/\partial x_i$, and so the horizontal trajectories are the pullbacks under ζ_i of horizontal lines $y_i = a$.

Moreover, on each U_i , dy_i defines a *transverse measure*, denoted by $|dy_i| = |\text{Im} \sqrt{\varphi}|$, which allows us to “count” how many leaves cross a curve, or, in other words, its total distance traveled in the vertical direction. We call this the “height” of the curve relative to the foliation, and denote it

$$\text{ht}(c, |\text{Im} \sqrt{\varphi}|) = \int_c |\text{Im} \sqrt{\varphi}|$$




FIGURE 2.2: Natural coordinates map the leaves of the vertical foliation of a holomorphic quadratic differential to horizontal lines in the plane.

The quantity $|dy_i|$, which will be described more formally below, is known as the *vertical foliation* of the holomorphic quadratic differential φ .

Note that on $U_i \cap U_j$,

$$d\zeta_i^2 = d\zeta_j^2,$$

i.e., $d\zeta_i = \pm d\zeta_j$. So $dy_i = \pm dy_j$, whence $y_i = \pm y_j + c_{ij}$ for some $c_{ij} \in \mathbb{R}$. This motivates the following definitions [5]:

Definition 2.5. A **partial measured foliation** $[dv]$ of a Riemann surface R is a family of open subsets $\mathcal{U} = \{U_i\}$ together with Lipschitz-continuous functions $v_i : U_i \rightarrow \mathbb{R}$ such that on $U_i \cap U_j$,

$$v_i = \pm v_j + c_{ij}$$

for some constant c_{ij} . Note that the U_i do not have to cover R , whence the term *partial* measured foliation.

Definition 2.6. The **leaves** or **level sets** of $|dv|$ are sets of points L such that for each i ,

$$L \cap U_i = v_i^{-1}(a_i)$$

for some constant $a_i \in \mathbb{R}$. Note that on $U_i \cap U_j$,

$$v_i^{-1}(a_i) = v_j^{-1}(\pm(a_i - c_{ij})),$$

so the level sets are well-defined on overlaps.

Definition 2.7. A partial measured foliation $|dv|$ is a **measured foliation** if every point $p \in R$ has a neighborhood with a homeomorphism h to a neighborhood of 0 in the complex plane and such that h carries the leaves of $|dv|$ onto the leaves of $z^{n/2}|dz|$ for some $n \geq 0$.

If $n = 0$, p is called a **regular point**. If $n > 0$, it is called an $(n + 2)$ -pronged **singular point** or a **zero** of degree n . If $n < 0$ (not permitted for a measured foliation), p is a **pole** (Figure 2.3).




FIGURE 2.3: Singular and regular points of a measured foliation.

Definition 2.8. Let $|dv|$ be a partial measured foliation on a surface R , and let $c : [0, 1] \rightarrow R$ be a curve on R . The **height** of c relative to $|dv|$,

$$\text{ht}(c, |dv|) = \int_c |dv|$$

is defined as follows. Let $\mathcal{P} = \{0 = t_0 < t_1 < \dots < t_n = 1\}$ be a partition of the unit interval. Let $I_k = [t_{k-1}, t_k]$ for each k . The integral of $|dv|$ over the partition is defined to be

$$\int_{\mathcal{P}} |dv| = \sum_i \left[\sum_{c(I_k) \subset U_i} |v(c(t_k)) - v(c(t_{k-1}))| \right]$$

Note that the quantity $|v(c(t_k)) - v(c(t_{k-1}))|$ is well-defined on overlaps, since v_i and v_j differ by a sign and a shift.

Next, we define the integral of $|dv|$ over c as the Riemann integral

$$\int_c |dv| = \limsup_{\Delta t \rightarrow 0} \int_{\mathcal{P}} |dv|,$$

where $\Delta t = \max_k (t_k - t_{k-1})$.

The **height of a homotopy class** $[c]$ is the infimum over all homotopic curves

$$\text{ht}([c], |dv|) = \inf_{c' \sim c} \text{ht}(c', |dv|).$$

It is difficult to calculate the infimum in this definition directly. However, it becomes easier if we can find a minimum-height representative. Luckily, there is a sufficient condition for minimality:

Definition 2.9 ([7]). A curve c is **quasitransverse** to a foliation $|dv|$ if each connected component of $c \setminus \{\text{singularities of } |dv|\}$ is either an arc of a leaf of $|dv|$ or transverse to the leaves of $|dv|$. In other words, c is always either transverse or parallel to the leaves, switching only at singularities. Moreover, c must leave a singularity in a different sector than the one it entered in (sectors are regions divided by prongs of the singularity—see Figure 2.3).

Theorem 2.10 ([7]). *If c is quasitransverse to $|dv|$, then*

$$\text{ht}(c, |dv|) = \text{ht}([c], |dv|).$$

2.3.1 Whitehead Classes

There is a natural equivalence relation on measured foliation which extends the notion of homotopy of curves.

Definition 2.11 ([7]). A **Whitehead operation** (or Whitehead move) is an operation that either joins two singularities of a foliation, or splits one singularity into a pair of singularities connected by a leaf (Figure 2.4).




FIGURE 2.4: A Whitehead move.

Definition 2.12 ([7]). Two foliations $|dv|$ and $|dw|$ are **Whitehead equivalent** if they are related by a sequence of isotopies interrupted by Whitehead operations. Since Whitehead operations are reversible, this is an equivalence relation, and it divides the space of measured foliations into **Whitehead classes**. We will denote the class of $|dv|$ by $W(|dv|)$. We will denote the set of Whitehead classes of foliations on a surface S by $\mathcal{MF}(S)$.

A Whitehead class can also be defined by its heights on simple closed curves:

Theorem 2.13 ([7]). $|dv| \in W(|dw|) \iff \text{ht}([c], |dv|) = \text{ht}([c], |dw|)$ for every homotopy class of simple closed curves $[c]$.

2.3.2 Dirichlet Energy and the Hubbard-Masur Theorem

It is clear that a holomorphic quadratic differential defines a unique measured foliation, and by extension, a unique Whitehead class. Is the converse true? Does each Whitehead class have a unique representative which is the vertical foliation of a holomorphic quadratic differential? This question is inspired by Hodge's theorem, which answers the analogous question for 1-forms. It states that each cohomology class of 1-forms is represented by a unique harmonic 1-form.

The question for foliations was first answered in the affirmative by Hubbard and Masur [8], but we shall follow the treatment of Gardiner and Lakic [6]. They define a Dirichlet-type energy as follows:

$$D(|dv|) = \int_R ((v_x)^2 + (v_y)^2) dx dy,$$

where v_x and v_y are the partial derivatives of v (defined as v_i on each covering set U_i). Note that D is invariant under change of coordinates and on overlaps $U_i \cap U_j$.

If we were considering not a foliation, but a 1-form ω equal to dv_i on each U_i , with $v_i = v_j + c_{ij}$, then D would be the usual Dirichlet energy. A minimum ω^0 for D would satisfy the Laplace equation

$$0 = \Delta v_i^0 = (d\delta + \delta d)v_i^0 = \delta(dv_i^0) = \delta\omega^0.$$

Therefore, ω_0 would be coclosed, hence harmonic. By analogy, we call a foliation which minimizes D in its Whitehead class **harmonic**. The following is the analog of the Hodge theorem for measured foliations:

Theorem 2.14 (Gardiner-Lakic [5, 6]). Let $|d\tilde{v}|$ be a partial measured foliation on a Riemann surface R .

(i) *There is a unique measured foliation $|dv_0|$ such that*

$$D(|dv_0|) = \inf \{D(|dv|) : \forall \text{ curves } c, \text{ht}([c], |dv|) \geq \text{ht}([c], |d\tilde{v}|\}\}.$$

(ii) *$\text{ht}([c], |dv_0|) = \text{ht}([c], |d\tilde{v}|)$ for all curves c . In particular, if $|d\tilde{v}|$ is a measured foliation, and not merely a partial foliation, then $|dv_0|$ is the minimum-energy foliation in $W(|d\tilde{v}|)$.*

(iii) *$|dv_0|$ is the vertical measured foliation of a holomorphic quadratic differential, uniquely determined by the Whitehead class of $|d\tilde{v}|$.*

Proof. [6]

□

Gardiner and Lakic prove a further result about the energy as a function on Teichmüller Space, but that is beyond the scope of this thesis.

Chapter 3

Discrete Theory

To discretize the smooth theory of measured foliations, we use the same technique employed to discretize 1-forms in [11]. To wit, for a measured foliation $|dv|$ and a mesh edge $e = (i, j)$, we assign a value

$$F_e = \int_e |dv| = |v_j - v_i|$$

Unlike 1-forms, measured foliations are not orientable. That is, there is in general no way to assign consistent signs to the values F_e . In view of this issue, we declare the edge values F_e to be positive. Then we restrict to the smaller class of closed foliations to impose consistency.

Let M be a triangulated mesh in \mathbb{R}^3 . Let V be the set of vertices, E the set of edges, and T the set of triangles. We will identify edges by unordered pairs of vertices $e = [i, j] \in E$ and triangles by ordered triples $t = (i, j, k) \in T$. M also has a set H of *halfedges*, or directed edges, which we denote by ordered pairs of vertices. For example, the edge $[i, j]$ comprises the halfedges (i, j) and (j, i) , and the boundary of the triangle (i, j, k) consists of the halfedges (i, j) , (j, k) , and (k, i) . By convention, $(i, j) = -(j, i)$.

Definition 3.1. A **discrete measured foliation** is a map $F : E \rightarrow \mathbb{R}^{\geq 0}$. We will denote the value $F([i, j])$ by F_{ij} . (Note that $F_{ij} = F_{ji}$.)

In the theory of discrete exterior calculus, the differential or coboundary operator d plays a major role. By Stokes' theorem, a smooth 1-form is closed if and only if its integral over the boundary of any disk is zero. In the discrete setting, this is used to define the differential of a 1-form ω on each triangle

$$(d\omega)_{ijk} = \omega_{ij} + \omega_{jk} + \omega_{ki}$$

so that $d\omega = 0$ if and only if ω is closed.

A measured foliation is locally (away from singularities) an exact (hence closed) 1-form. Thus, on any regular triangle $t = (i, j, k)$ there should be an assignment of signs σ_η to the halfedges such that

$$\sigma_{ij}F_{ij} + \sigma_{jk}F_{jk} + \sigma_{ki}F_{ki} = 0.$$

This implies that one of the signs must differ from the other two. Without loss of generality, assume that $\sigma_{ij} = -\sigma_{jk} = -\sigma_{ki}$. Then we can write the above equation more simply as

$$F_{ij} = F_{jk} + F_{ki}.$$

In this case, we say the foliation is closed on t at the corner k :

Definition 3.2. Let $t \in T$ with vertices i, j, k in cyclic order. We say a discrete measured foliation F is **closed** at the corner k if

$$F_{jk} + F_{ki} = F_{ij},$$

and similarly for the other corners. To represent this structure visually, we will mark closed corners of a foliation with an \circ and the rest with an \times (see Figure 3.1). The \circ s divide the edge-neighborhood of each vertex into **sectors**.

F is **regular** on t if it is closed at at least one of the corners of t . Otherwise, it is **singular** on t . If one of the edge values dominates the others ($F_{ij} > F_{jk} + F_{ki}$), we say F has a simple pole on t . If the three edge values obey the strict triangle inequalities

$$F_{ij} < F_{jk} + F_{ik}, \quad F_{jk} < F_{ij} + F_{ik}, \quad F_{ik} < F_{ij} + F_{jk}$$

we say F has a simple zero on t . Intuitively, the triangle inequalities say that every leaf that enters one edge must leave through one of the other edges. In a regular triangle, one of the inequalities is an equality. The edge opposite the closed corner saturates the other two edges with leaves.

If $F_e = 0$ for every edge e bounding a triangle t , we say t is an empty triangle of F . If F has empty triangles, it is a **partial foliation**.

3.1 The Index Theorem

Note that the markings \circ and \times are completely determined by the edge values. Moreover, the signs σ_η defined above are recoverable up to a global sign per triangle. Given




FIGURE 3.1: The foliation values in (a) determine the markings and leaves in (b). The top two triangles are singular, each containing a simple (degree-1) zero, while the other triangles are regular. Note that the central vertex has two sectors comprising the “northern” and “southern” outgoing edges, respectively.

two adjacent regular triangles sharing an edge, we may choose the signs σ consistently. In fact, this should be possible throughout any simply-connected region on which F is regular. This will be the case provided that the sign does not change when walking around any vertex. When walking around a vertex, every \circ subtends a sign change, since

$$\sigma_{jk} = \sigma_{ki} = -\sigma_{ik}$$

when k is a closed corner of the triangle (i, j, k) . This motivates the following definition:

Definition 3.3. Let F be a discrete measured foliation defined on a triangulated surface M , possibly with boundary. For a vertex $v \in V$, let $O(v)$ and $X(v)$ be the number of \circ s and \times s, respectively, at corners abutting v . In particular,

$$X(v) + O(v) = \deg(v).$$

Let $Z(v)$ be the number of halfedges $\eta = (u, v)$ for which $F_\eta = 0$.

The **index** of F at a vertex v is

$$\iota_v(F) = \begin{cases} 1 - O(v) + Z(v) & v \in \partial M \\ 2 - O(v) + Z(v) & \text{otherwise.} \end{cases}$$

F is regular at v if $\iota_v(F) = 0$. It has a zero of degree d at v if $\iota_v(F) = +d$, and it has a pole of degree d at v if $\iota_v(F) = -d$.

Similarly, the index of F on a triangle t is

$$\iota_t(F) = 2 - X(t) - Z(t).$$

where $X(t)$ is the number of \times s in corners of t and $Z(t)$ is the number of zero halfedges in ∂t .

We discount the number of zero halfedges from the number of \circ s at a vertex because, though a zero halfedge has \circ s at both its head and its tail, a slight perturbation of the foliation would collapse these into only one \circ .

The indices of vertices and faces are closely related to the topology of the mesh via a discretized version of the Poincare-Hopf Index Theorem similar to the one proven for 1-forms in [14]:

Theorem 3.4 (Index Theorem). *Let F be a discrete measured foliation on a mesh $M = (V, E, T)$ whose boundary ∂M is a collection of zero or more simple edge-loops. Then*

$$\sum_{v \in V} \iota_v(F) + \sum_{t \in T} \iota_t(F) = 2\chi(M),$$

where $\chi(M)$ is the Euler characteristic of M .

Proof. Let $\partial V = V \cap \partial M$ and $V^\circ = V \cap M^\circ$.

$$\begin{aligned} \sum_{v \in V} \iota_v(F) + \sum_{t \in T} \iota_t(F) &= \sum_{v \in \partial V} [1 - O(v) + Z(v)] + \sum_{v \in V^\circ} [2 - O(v) + Z(v)] + \sum_{t \in T} [2 - X(t) - Z(t)] \\ &= |\partial V| + 2|V^\circ| + 2|T| - \sum_{v \in V} O(v) - \sum_{t \in T} X(t) \\ &= |\partial V| + 2|V^\circ| + 2|T| - |H| \\ &= |\partial V| + 2|V^\circ| + 2|T| - [2(|E| - |\partial V|) + |\partial V|] \\ &= 2|V| - 2|E| + 2|T| \\ &= 2\chi(M) \end{aligned}$$

The third equality is due to the fact that there is a one-to-one correspondence between corners and halfedges (each halfedge lies in a unique triangle; take the corner at its tip).

The fourth equality enumerates the halfedges: there are two halfedges for each interior edge, but only one for each boundary edge. Moreover, since the boundary consists of simple closed curves, the number of boundary edges is the same as the number of boundary vertices. \square

Definition 3.5. A discrete measured foliation is **closed** if it has no poles.

Unless otherwise noted, we will assume that all of our foliations are closed. Note that Theorem 3.4 constrains the number of singularities a closed foliation may have. For

example, on a torus ($\chi = 0$), a closed foliation must be regular at all vertices and on all triangles.

3.2 Paths and Whitehead Equivalence

Let M be a mesh and F be a discrete measured foliation on M . The edge data of F defines a unique continuous, piecewise-linear foliation on M , as follows. First consider a regular triangle $t = (o, i, j)$, where o is the closed corner. Define a linear function f on t , with $f(o) = 0$, $f(i) = F_{oi}$, and $f(j) = -F_{oj}$. So for a point $p \in t$ with barycentric coordinates $(\beta_o, \beta_i, \beta_j)$,

$$f(p) = \beta_i F_{oi} - \beta_j F_{oj}.$$

For any arc $c : [0, 1] \rightarrow t$, we can now define the line integral

$$\int_c |df| = \int_0^1 \left| F_{oi} \frac{\partial \beta_i}{\partial t} - F_{oj} \frac{\partial \beta_j}{\partial t} \right| dt.$$

If c is a line segment, then $\dot{\beta}_i$ and $\dot{\beta}_j$ are constant, and the integral is equal to

$$\left| \int_c df \right| = |f(c(1)) - f(c(0))|.$$

In particular, for each edge e in the boundary of t ,

$$\int_e |df| = F_e.$$

Now suppose $t = (i, j, k)$ is singular. Let m be its barycenter, and divide t into three sub-triangles

$$(i, j, m), (j, k, m), (k, i, m).$$

If F satisfies the triangle inequalities on t , we can define values F_{im} , F_{jm} , and F_{km} so that each sub-triangle is closed at m . Indeed, this is equivalent to solving the following system of equations

$$\begin{aligned} F_{im} + F_{jm} &= F_{ij} \\ F_{im} + F_{km} &= F_{ik} \\ F_{jm} + F_{km} &= F_{jk} \end{aligned}$$

which has solutions

$$F_{im} = \frac{1}{2}(F_{ij} + F_{ik} - F_{jk}) \quad F_{jm} = \frac{1}{2}(F_{ij} + F_{jk} - F_{ik}) \quad F_{km} = \frac{1}{2}(F_{ik} + F_{jk} - F_{ij}).$$

These are nonnegative precisely when the original foliation values satisfy the triangle inequalities on t , i.e., when t is a zero of F . So we can divide any singular triangle of a closed foliation into three regular triangles sharing a singular vertex. Thus, the line integral on a regular triangle described above suffices to define the line integral of an arbitrary curve in M . In particular, if c is an edge-path consisting of vertices c_1, \dots, c_n , its height (intersection number) with respect to the foliation F is simply the sum of the foliation values on its edges:

$$\text{ht}(c, F) = \sum_{i=1}^{n-1} F_{c_i c_{i+1}}$$

For any closed curve c , we can define the height of its homotopy class as in the smooth setting:

$$\text{ht}([c], F) = \inf_{c' \sim c} \text{ht}(c', F)$$

We can then use this to define Whitehead equivalence. Whereas in Section 2.3.1 Whitehead equivalence was defined as the transitive closure of isotopy and Whitehead moves and claimed to be equivalent to height-equivalence, here we use height-equivalence as the primary definition:

Definition 3.6. Two discrete measured foliations F and F' are **Whitehead equivalent** if

$$\text{ht}([c], F) = \text{ht}([c], F')$$

for any class of simple closed paths $[c]$. We denote the Whitehead equivalence class of F by $W(F)$, and the space of Whitehead classes of discrete foliations by \mathcal{MF} .

Definition 3.7. Let F be a discrete measured foliation. An edge-path

$$c = c_1 - c_2 - \dots - c_n$$

is **quasitransverse** if it passes between two \circ s at each vertex c_i , i.e., if $[c_{i-1}, c_i]$ and $[c_i, c_{i+1}]$ are in different sectors at c_i for each i .

Lemma 3.8. *Let F be a closed discrete measured foliation, and let c be a quasitransverse edge-path. Then there is no leaf of F that, together with an arc of c , encloses a disk.*

Proof. We ignore zero-edges for simplicity.

Suppose the lemma is false. Let L be an arc of a leaf of F intersecting c at points p and q and enclosing a disk D_0 . Assume without loss of generality that $p \in [c_1, c_2]$ and $q \in [c_{n-1}, c_n]$. Let t_1, \dots, t_m be the triangles L traverses between p and q . Then

$$D = D_0 \cup t_1 \cup \dots \cup t_m$$

is a triangulated disk whose boundary consists of c together with another arc

$$l = \{c_1 = l_1, l_2, \dots, l_r = c_n\}$$

consisting of the outward-facing edges of the t_i .

$\chi(D) = 1$, so Theorem 3.4 implies that the indices of vertices and faces in D add up to 2. Since F is closed, $\iota_t \leq 0$ for every triangle, and $\iota_v \leq 0$ for every interior vertex of D . Moreover, since c is quasitransverse, $\iota_{c_i} \leq 0$ for each $i = 2, 3, \dots, n-1$ (where we compute the indices of these vertices relative to D , not to M). By possibly removing t_1 and t_m from D , we can ensure that ι_{c_0} and ι_{c_n} are also nonpositive. As such, we must have

$$2 = \sum_{j=2}^{r-1} \iota_{l_j} + \sum_{i=1}^m \iota_{t_i} \leq r - 2 - \#\{\circ\text{s along } l\} - \#\{\text{singular triangles along } l\}. \quad (3.1)$$

Consider one of the triangles t_i containing an edge $[l_j, l_{j+1}]$ of l . Call its third vertex v . Since the leaf L crosses between $[v, l_j]$ and $[v, l_{j+1}]$, v is not closed. Thus, either t_i is singular, or it has an \circ at l_j or l_{j+1} . Since there are $r-3$ edges $[l_2, l_3], \dots, [l_{r-2}, l_{r-1}]$, the number of \circ s on l is

$$r - 3 - \#\{\text{singular triangles along } l\}.$$

Combining this with 3.1, we obtain a contradiction. \square

Lemma 3.9. *Let F be a discrete measured foliation without poles, and let c be a closed quasitransverse edge-path. Then*

$$\text{ht}(c, F) = \text{ht}([c], F).$$

Similarly, let c be a quasitransverse edge-path between vertices s and t . Then $\text{ht}(c, F)$ is minimal in the homotopy class of paths between s and t .

Proof. Let c' be a closed path homotopic to c (resp., a path between s and t homotopic to c). In particular, c and c' are homologous, so $c' - c = \partial K$ for some region K . By 3.8, any leaf that enters K through c must exit through c' . Therefore,

$$\text{ht}(c', F) \geq \text{ht}(c, F). \quad \square$$

3.3 Harmonicity

We have defined closedness for discrete foliations, but we have yet to define harmonicity. We want this to formalize the property of being the vertical foliation of a holomorphic quadratic differential. For a holomorphic quadratic differential φ , the vertical foliation $|\operatorname{Im} \sqrt{\varphi}|$ is locally (away from zeros) equal to $|dy|$, where $\zeta = x+iy$ is a natural coordinate for φ . Since ζ is holomorphic, y is harmonic, i.e.,

$$\Delta y = 0,$$

where Δ is the Laplacian.

Let us now translate this to the discrete setting. In a simply-connected, regular region, we have $F_{ij} = |y_i - y_j|$ for some function y on the vertices. We want y to be harmonic. In other words, y should satisfy the discrete Laplace's equation $\Delta y = 0$. The discrete Laplacian takes the form

$$(\Delta y)_i = \frac{1}{2} \sum_{[i,j] \in E} \alpha_{ij} (y_j - y_i),$$

where $\alpha_{ij} = \cot \theta_{ij} + \cot \phi_{ij}$ are the standard cotangent weights [12]. So Laplace's equation becomes

$$0 = (\Delta y)_i = \frac{1}{2} \sum_{[i,j] \in E} \alpha_{ij} F_{ij} \sigma_{ij},$$

for each regular vertex i , where σ_{ij} are the halfedge signs defined earlier in this chapter.

We do not know the signs σ_η , but we do know that there are two sign changes as we walk around a regular vertex, so that the two sectors have different signs. Thus, we can rewrite the harmonicity condition as

$$\sum_{[i,j] \in S_1} \alpha_{ij} F_{ij} = \sum_{[i,j] \in S_2} \alpha_{ij} F_{ij},$$

where S_1 and S_2 are the two sectors of edges at i . This motivates the more general definition of harmonicity, as follows:

Definition 3.10. A discrete measured foliation F is **coclosed** at a vertex v if for each sector S at v ,

$$\sum_{e \in S} \alpha_e F_e \leq \frac{1}{2} \sum_{e \in \mathcal{E}(v)} \alpha_e F_e,$$

where $\mathcal{E}(v)$ is the set of edges abutting v . If F is regular at v , there are two sectors (say S_1 and S_2), and the condition becomes

$$\sum_{e \in S_1} \alpha_e F_e = \sum_{e \in S_2} \alpha_e F_e.$$

A foliation which is both closed and coclosed is **harmonic**.

3.4 Discrete Dirichlet Energy

Definition 3.11. The Dirichlet Energy of a discrete measured foliation F is

$$D(F) = \sum_{e \in E} \alpha_e F_e^2$$

Theorem 3.12 (Discrete Hubbard-Masur Uniqueness [22]). *Let M be a mesh such that the cotangent weights α_e are positive everywhere. Let F be a discrete harmonic foliation on M . Let F' be another foliation such that for any closed path c ,*

$$\text{ht}(c, F') \geq \text{ht}([c], F)$$

Then

$$D(F) < D(F').$$

In particular, F is the unique harmonic foliation in its Whitehead class.

The proof will require some preliminary work.

Definition 3.13. Let F be a discrete foliation on M . A **train track** τ subordinate to F is an assignment of nonnegative real values τ_e to edges $e \in E(M)$ such that

- (i) If $F_e = 0$, then $\tau_e = 0$.
- (ii) For any vertex $v \in V(M)$ and any sector S at v (as defined by F),

$$\sum_{e \in S} \tau_e \leq \frac{1}{2} \sum_{e \in \mathcal{E}(v)} \tau_e.$$

That is, no sector of the train track “dominates” the vertex.

τ is a rational (respectively, integral) train track if $\tau_e \in \mathbb{Q}$ (respectively, \mathbb{Z}) for each e . We can compute the height of a train track as a generalization of the height of an

edge-path:

$$\text{ht}(\tau, F) = \sum_{e \in E} \tau_e F_e.$$

The terminology “train track” is inspired by an object W. Thurston defined on smooth surfaces.

Lemma 3.14. *The rational train tracks are dense in the set of all train tracks subordinate to a foliation F . In particular, for any train track τ and $\epsilon > 0$, there is a rational train track $\bar{\tau}$ such that $|\tau_e - \bar{\tau}_e| < \epsilon$ for all $e \in E$.*

Proof. Rational points are dense in \mathbb{R}^E , and the train tracks are defined as a subset of \mathbb{R}^E by linear inequalities with rational coefficients. \square

Lemma 3.15. *Let τ be an integral train track subordinate to a foliation F , and such that τ_e is even for every edge e . Then there is a closed, quasitransverse edge-path c such that the multiplicity of c on each edge e is equal to τ_e . In particular,*

$$\text{ht}(c, F) = \text{ht}(\tau, F)$$

Moreover, by perturbing c slightly off the edges of M , it can be made simple.

Proof. Start with τ_e arcs on each edge e . Constructing c is simply a matter of ensuring that every arc entering a vertex v connects up with another arc leaving v through a different sector. We show such a matching can be found by induction on the number of sectors at v .

If v has two sectors, then simply match each arc entering one sector to an arc leaving the other sector—this is possible since τ is a valid train track.

Suppose v has three sectors, S_1, S_2, S_3 . For each $i = 1, 2, 3$, let

$$N_i = \sum_{e \in S_i} \tau_e.$$

Since the τ_e are even, so are the N_i . Moreover, since τ is a valid train track, they satisfy the triangle inequality

$$N_i \leq N_j + N_k$$

for all $i \neq j \neq k$. Denote by n_{ij} the number of arcs entering through S_i and leaving through S_j . Then we must find a solution in nonnegative integers to the system

$$\begin{aligned} n_{12} + n_{23} &= N_2 \\ n_{12} + n_{13} &= N_1 \\ n_{23} + n_{13} &= N_3 \end{aligned}$$

Indeed, the solution is

$$n_{ij} = \frac{1}{2}(N_i + N_j - N_k), \quad \forall i \neq j \neq k.$$

Since the N_i are even, and by the triangle inequalities, we have found the (nonnegative integral) solution we seek.

Now suppose there are r sectors S_1, \dots, S_r , and suppose by way of induction that the matching can be done for $r - 1$ sectors. Pick two adjacent sectors, say S_1 and S_2 , and define $\bar{S} = S_1 \cup S_2$ and

$$\bar{N} = \min \left\{ N_1 + N_2, \sum_{i=3}^r N_i \right\}.$$

Consider \bar{S} as though it were a single sector with total weight \bar{N} . By definition,

$$\bar{N} \leq \sum_{i=3}^r N_i,$$

and for $j = 3, \dots, r$,

$$N_j \leq \bar{N} + \sum_{i>2, i \neq j} N_i.$$

So the induction hypothesis guarantees that there is a matching of arcs between the sectors \bar{S}, S_3, \dots, S_r . Within sector \bar{S} , we have


$$\bar{N} \leq N_1 + N_2,$$

and

$$N_1 \leq \sum_{i=2}^r N_i = \bar{N} + N_2$$

and similarly $N_2 \leq N_1 + \bar{N}$. So by the three-sector case, we can find a matching *within* \bar{S} , so that each arc that enters S_1 (respectively, S_2) either leaves through S_2 (respectively, S_1) or leaves \bar{S} and is subsequently matched to one of the other sectors.

To build a *simple* curve, use τ_e parallel arcs in a small strip around edge e , and connect arcs “outermost-first” at vertices (see Figure 3.2). \square



(A) An even integral train track

(B) A corresponding simple curve

FIGURE 3.2: The curve in (b) is uniquely determined by the edge values and markings in (a).

Proof of Theorem 3.12. We will prove the contrapositive—namely, that if F' is a distinct foliation such that $D(F') \leq D(F)$, then there is some simple closed curve c such that

$$\text{ht}(c, F') < \text{ht}([c], F).$$

If $F' = \lambda F$ for some $\lambda < 1$, the conclusion is true for any simple closed curve. So assume otherwise.

For arbitrary maps $G, G' : E \rightarrow \mathbb{R}$, let

$$\langle G, G' \rangle = \sum_{e \in E} \alpha_e G_e G'_e,$$

and observe that if $\alpha_e > 0$ for all e , this is an inner product.

Let τ be a train track defined by $\tau_e = \alpha_e F_e$ for each $e \in E$. Since F is harmonic, this is a valid train track. Observe that

$$\text{ht}(\tau, F) = \sum_{e \in E} \tau_e F_e = \sum_{e \in E} \alpha_e F_e^2 = D(F)$$

and

$$\text{ht}(\tau, F') = \sum_{e \in E} \tau_e F'_e = \sum_{e \in E} \alpha_e F_e F'_e = \langle F, F' \rangle$$

By the Cauchy-Schwarz inequality,

$$\text{ht}(\tau, F') = \langle F, F' \rangle < \sqrt{D(F)D(F')} \leq D(F) = \text{ht}(\tau, F),$$

where the first (Cauchy-Schwarz) inequality is strict since F' and F are linearly independent as elements of $\mathbb{R}^{|E|}$.

Let $\delta > 0$ be small enough that $\text{ht}(\tau, F') < \text{ht}(\tau, F) - \delta$. Define

$$\epsilon = \frac{\delta}{\sum_{e \in E} |F_e - F'_e|}$$

Using Lemma 3.14, let $\bar{\tau}$ be a rational train track subordinate to F and such that $|\bar{\tau}_e - \tau_e| < \epsilon$ for each e . Then

$$\begin{aligned} \text{ht}(\bar{\tau}, F) - \text{ht}(\bar{\tau}, F') &= \sum_{e \in E} \bar{\tau}_e (F_e - F'_e) \\ &= \sum_{e \in E} \tau_e (F_e - F'_e) + \sum_{e \in E} (\bar{\tau}_e - \tau_e) (F_e - F'_e) \\ &> \delta - \epsilon \sum_{e \in E} |F_e - F'_e| \\ &> 0. \end{aligned}$$

Let N be such that $N\bar{\tau}_e$ is an even integer for each e . Then $N\bar{\tau}$ is an even integral train track subordinate to F . By Lemma 3.15, there is a closed curve c , quasitransverse to F , whose multiplicity on edge e is $N\bar{\tau}_e$. So

$$\text{ht}([c], F) = \text{ht}(c, F) = \text{ht}(N\bar{\tau}, F) > \text{ht}(N\bar{\tau}, F') = \text{ht}(c, F').$$

Note that while c is not simple, we can find a simple closed curve arbitrarily close to c and with height arbitrarily close to $\text{ht}(c, F)$. This justifies writing $\text{ht}([c], F) = \text{ht}(c, F)$. \square

Chapter 4

Implementation

4.1 Discrete Whitehead Moves

Definition 4.1. Let $v \in V$. Divide its edge ring $\mathcal{E}(v)$ into sectors E_1, \dots, E_k separated by \circ markings. Assume $k \geq 2$. A discrete Whitehead move on sector 1 at v takes $F \mapsto W_s F$ with

$$(W_s F)_e = \begin{cases} F_e - s & e \in E_1 \\ F_e + s & e \in \mathcal{E}(v) \setminus E_1 \end{cases}$$

for $s \in \mathbb{R}^+$ small enough that the following constraints are satisfied:

- (i) $(W_s F)_e > 0$ for all $e \in \mathcal{E}(v)$.
- (ii) If $t = (u, v, w)$ is a triangle abutting v , then for all s , $W_s F$ obeys the triangle inequalities on t :

$$\begin{aligned} (W_s F)_{uv} + (W_s F)_{vw} &\geq F_{uw} & (W_s F)_{vw} + F_{uw} &\geq (W_s F)_{uv} \\ F_{uw} + (W_s F)_{uv} &\geq (W_s F)_{vw}. \end{aligned}$$

We will need the following lemma from [10]:

Lemma 4.2. *Consider a quadrilateral on a surface with vertices a, b, c, d in clockwise order and diagonal bd . Let F be a foliation on the surface. If we flip the diagonal (erase bd and replace it with ac) then the minimum integral of the foliation over the new edge is given by*

$$F_{ac} = \max\{F_{ab} + F_{cd}, F_{ad} + F_{bc}\} - F_{bd}.$$

Lemma 4.3. *If F' is reachable from F by a sequence of discrete Whitehead moves, then for any quasitransverse closed edge-path c ,*

$$\text{ht}([c], F') = \text{ht}(c, F) = \text{ht}([c], F).$$

Proof. Each discrete Whitehead move only changes foliation values on edges in the 1-ring of a single vertex. So it suffices to consider vertices c traverses. Let v be such a vertex, and let W_s be a Whitehead move on sector 1 at v . Since c is quasitransverse, it enters and leaves v on edges in different sectors, say $[u, v] \in E_i$ and $[v, w] \in E_j$. If $i = 1$ or $j = 1$, then one of F_{uv} and F_{vw} increases by s , while the other decreases by s , so the height of c remains unchanged, i.e.,

$$\text{ht}(c, W_s F) = \text{ht}(c, F) = \text{ht}([c], F).$$

The \circ s surrounding E_1 stay there after the Whitehead move, so c remains quasitransverse, and so $\text{ht}(c, W_s F) = \text{ht}([c], W_s F)$.

Otherwise, if neither i nor j is 1, then both F_{uv} and F_{vw} get larger, so the height of c increases. Walk around v from u to w in the direction that does not intersect E_1 . Label the encountered vertices $u = v_0, v_1, \dots, v_r = w$. Flip each edge $[v, v_k]$ in turn. By Lemma 4.2, we have

$$F_{uv_k} = \max\{F_{uv_{k-1}} + F_{vv_k}, F_{uv} + F_{v_{k-1}v_k}\} - F_{vv_{k-1}}. \quad (4.1)$$

We know that $(W_s F)_{uv_1} = F_{uv_1}$. Assume by way of induction that $(W_s F)_{uv_{k-1}} = F_{uv_{k-1}}$. We know further that

$$\begin{aligned} (W_s F)_{vv_k} &= F_{vv_k} + s & (W_s F)_{uv} &= F_{uv} + s & (W_s F)_{vv_{k-1}} &= F_{vv_{k-1}} + s \\ (W_s F)_{v_{k-1}v_k} &= F_{v_{k-1}v_k}. \end{aligned}$$

Plugging these into (4.1) yields


$$(W_s F)_{uv_k} = F_{uv_k}.$$

By induction,

$$(W_s F)_{uw} = F_{uw} = F_{uv} + F_{vw}.$$

Let c' be the path which is identical to c outside the 1-ring of v , and which takes the path of minimal height from u to w with respect to $W_s F$. Then c' is homotopic to c , and

$$\text{ht}([c], W_s F) = \text{ht}(c', W_s F) = \text{ht}(c, F) = \text{ht}([c], F).$$



(A) Initial foliation and path

(B) After a Whitehead move

FIGURE 4.1: The blue arc in (a) would no longer be quasitransverse to the foliation depicted in (b), but it can be deformed as shown to preserve its height. The singular triangles on the right are divided into regular sub-triangles by dotted lines.

Thus, for any Whitehead move W_s ,

$$\text{ht}([c], W_s F) = \text{ht}(c, F).$$

The lemma follows by induction. □

4.2 Main Algorithm

The idea of our algorithm is to gradually decrease the Dirichlet energy of a foliation while staying in its original Whitehead class. To that end, we apply gradient descent one vertex at a time via a Whitehead move at that vertex.

Suppose we modify a discrete foliation F by a whitehead move W_s on sector E_1 at a vertex v , and we allow s to vary. Then

$$\begin{aligned} \frac{d}{ds}[D(F)] &= \frac{d}{ds} \left[\sum_{e \in E} \alpha_e (W_s F)_e^2 \right] \\ &= 2 \sum_{e \in \mathcal{E}(v)} \alpha_e (W_s F)_e \frac{d}{ds} [(W_s F)_e] \\ &= 2 \sum_{e \in \bar{E}_1} \alpha_e (W_s F)_e - 2 \sum_{e \in E_1} \alpha_e (W_s F)_e, \end{aligned}$$

where $\bar{E}_1 = \mathcal{E}(v) \setminus E_1$. Setting $d/ds[D(F)] = 0$, we see that the Dirichlet energy is minimized when the foliation values in sector E_1 are reduced until

$$\sum_{e \in E_1} \alpha_e(W_s F)_e = \sum_{e \in \bar{E}_1} \alpha_e(W_s F)_e.$$

Substituting for the values of $W_s F$ and solving for s , we find that the optimal s is

$$s_0 = \frac{\sum_{e \in E_1} \alpha_e F_e - \sum_{e \in \bar{E}_1} \alpha_e F_e}{\sum_{e \in \mathcal{E}(v)} \alpha_e}$$

This is only a valid Whitehead move if $s_0 > 0$, which means E_1 must be such that

$$\sum_{e \in E_1} \alpha_e F_e > \sum_{e \in \bar{E}_1} \alpha_e F_e.$$

There can be only one such “dominant” sector at any given vertex. By applying W_{s_0} we make the vertex coclosed.

Algorithm 1 implements these ideas, adjusting one vertex at a time toward being coclosed via a Whitehead move. We use a heap to prioritize vertices with the largest energy gradients. Finally, we avoid creating strong poles:

Definition 4.4. A **strong pole** of a foliation F is either a pole in a triangle or a vertex v with only one sector. A **weak pole** is a pole at a vertex with at least two sectors and at least one zero halfedge.

4.2.1 Initialization

Our system allows the user to initialize a foliation from a set of disjoint, nontrivial dual loops. A dual loop is a cycle of triangles, each sharing one edge with the next. We define

$$F_e^0 = \begin{cases} 1 & \text{if } e \text{ is between two selected triangles} \\ 0 & \text{otherwise.} \end{cases}$$

This defines a closed partial foliation, since each selected triangle has exactly two interior edges. The height of any class of curves with respect to this initial foliation is the minimum number of times it crosses the dual loops. Since our algorithm maintains Whitehead class, the height of any class of curves with respect to the computed harmonic representative will still be an integer.

Algorithm 1 Relax Foliation

```

procedure RELAX( $M = (V, E, T), F^0, \epsilon$ )
  MAX-HEAP  $H \leftarrow \{\}$ 
  for  $v \in V$  do UPDATEHEAP( $H, F^0, v$ ) ▷ Add  $v$  to the heap.
  end for
  while  $\max_{v \in V} H[v] > \epsilon$  do
     $v \leftarrow H.\text{MAX}$ 
    WHITEHEADMOVE( $F, v$ )
    UPDATEHEAP( $H, F, v$ )
    for  $v \in \mathcal{N}(v)$  do
      UPDATEHEAP( $H, F, v$ )
    end for
  end while
end procedure

procedure WHITEHEADMOVE( $F, v$ )
   $(\nabla D, A, i) \leftarrow \text{COMPUTEGRADIENT}(F, v)$ 
   $s \leftarrow (\nabla D)/A$ 
  for  $e \in E_i$  do
     $s \leftarrow \min\{s, F_e\}$  ▷ Make sure foliation values stay positive.
  end for
  for  $t \in T$  abutting  $v$  do
    Reduce  $s$  until  $W_s F$  satisfies all 3 triangle inequalities on  $t$ 
    Reduce  $s$  until no strong poles are created at neighboring vertices
  end for
  for  $e \in E_i$  do
     $F_e \leftarrow F_e - s$ 
  end for
  for  $e \in \bar{E}_i$  do
     $F_e \leftarrow F_e + s$ 
  end for
end procedure

function COMPUTEGRADIENT( $F, v$ )
  Divide  $\mathcal{E}(v)$  into sectors  $E_1, \dots, E_n$ 
  for  $i = 1, \dots, n$  do
     $M_i \leftarrow \sum_{e \in E_i} \alpha_e F_e$ 
  end for
  if  $M_i > \sum_{j \neq i} M_j$  for some  $i$  then
     $\nabla D \leftarrow \sum_{e \in E_i} \alpha_e F_e - \sum_{e \in \bar{E}_i} \alpha_e F_e$ 
     $A \leftarrow \sum_{e \in \mathcal{E}(v)} \alpha_e$ 
    return  $(\nabla D, A, i)$ 
  else
    return  $(0, 0, \text{null})$ 
  end if
end function

procedure UPDATEHEAP( $H, F, v$ )
   $(\nabla D, A, i) \leftarrow \text{COMPUTEGRADIENT}(F_0, v)$ 
   $H[v] \leftarrow \nabla D$ 
end procedure

```

4.3 Convergence

We conjecture the following:

Conjecture 4.5. *If F has no strong poles, then every homotopy class of curves has a quasitransverse representative.*

If this is the case, then by Lemma 4.3, our algorithm preserves the Whitehead class of the foliation.

Conjecture 4.6. *Suppose M is such that the cotangent weights α_e are positive everywhere. If F has no poles, then its Whitehead class $W(F)$ contains a minimum of the Dirichlet energy D .*

Conjecture 4.7. *Suppose*

- (i) *M is such that the cotangent weights α_e are positive everywhere.*
- (ii) *F^0 contains no poles.*
- (iii) *Group the empty triangles of F^0 into connected components, then remove those empty components that are topological disks. Suppose that after this procedure, the sum of the indices of vertices on each boundary component is at most -2 .*


Then Algorithm 1 converges to the unique discrete harmonic foliation F in the Whitehead class of F^0 .

The idea of the last condition in Conjecture 4.7 is that an empty disk could be “hiding” a pole, which would only be revealed when the algorithm reduces the empty disk to a single vertex or triangle. However, by counting indices on the boundary of the empty disk, we can ensure that this will not happen.


4.4 Experimental Validation

We implemented the system as a plug-in to OpenFlipper [23], which provides basic infrastructure like a mesh library and a rendering system. All of our tests were performed on a 15” MacBook Pro (mid 2015) with a 2.8 GHz Intel Core i7 and 16 GB of RAM.

Figure 4.2 illustrates the convergence behavior of our algorithm. The algorithm was initialized with a foliation derived from three distinct dual loops on a genus-2 mesh. The Dirichlet energy (Figure 4.2c) declines steeply and then levels out.



(A) Initialization

(B) Harmonic foliation ($\epsilon = 10^{-5}$)

(c) Dirichlet Energy


FIGURE 4.2: A foliation initialized with the three dual loops illustrated in (a) converges to the harmonic representative in (b). (c) shows convergence of the Dirichlet energy over approximately 6 million Whitehead moves.

Figure 4.3 illustrates the uniqueness of the harmonic foliation within a given Whitehead class. Two foliations were initialized using two different, but isotopic, dual loops on the same genus-2 mesh. After running Algorithm 1 with $\epsilon = 10^{-5}$, they converge to nearly identical foliations, as we should expect from Theorem 3.12.



4.5 Conclusion

We have described a discrete representation for measured foliations on triangulated meshes, and we have developed an algorithm for generating harmonic measured foliations from arbitrary measured foliations. Future directions for research include

- proving the remaining conjectures in Section 4.3 to show that our algorithm is guaranteed to converge to the minimum of D
- developing an alternative, more efficient algorithm that is not limited to one Whitehead move at a time
- completing the research program set out in this paper by developing an algorithm to generate quasiconformal maps from harmonic measured foliations.



(A) Initialization 1.

(B) After 3.8 million Whitehead moves ($\epsilon = 10^{-5}$)

(C) Initialization 2


(D) After 4.3 million Whitehead moves ($\epsilon = 10^{-5}$)

FIGURE 4.3: Two different initializations (a), (c) in the same Whitehead class converge to the same harmonic representative (b), (d). Convergence took 18.3 s for (a) \rightarrow (b) and 19.3 s for (c) \rightarrow (d).

Appendix A

Rendering Measured Foliations

The goal of our rendering algorithm is to accurately display the leaves of a foliation and their transverse density by painting a texture on the surface. In Chapter 3, we discussed how, in a simply-connected region U not containing any zeros of a foliation F , the edge values can be assigned consistent signs so that F is locally an exact 1-form. In particular, there is a function $f : V \rightarrow \mathbb{R}$ such that

$$F_{vw} = |f_v - f_w|.$$

for all $[v, w] \in E$. If we take f to be our (1-dimensional) texture coordinate, we will be able to see the level sets of f (i.e., the leaves of F) as bands of constant color within U .

What if we try to expand the support U ? Either the region will close up (and no longer be simply connected) or we will reach a zero. Assume that zero triangles have been subdivided so that all zeros are at vertices. Walking around a simple zero introduces a sign change. Say v_0 is the zero vertex. If the texture coordinate of v_0 is f_{v_0} , then there is some edge $[v_0, w]$ such that on one side of the edge,

$$f_w = f_{v_0} + F_{v_0w}$$

and on the other side,

$$f_w = f_{v_0} - F_{v_0w}.$$

So the rendering will appear seamless if the texture is symmetric with respect to reflection about f_{v_0} .

Next suppose we walk around a nontrivial dual loop (strip of triangles). This will introduce a phase shift; if F is locally exact everywhere along the strip, the phase shift will be equal to the height of the homotopy class of the dual loop with respect to the

foliation. As such, the rendering will be seamless if the texture period is divisible by the height of the class.

So we have two conditions on our texture: it must be periodic with period divisible by the heights of curve classes and symmetric about the coordinates of zeros. Since the texture coordinates are only defined up to a global phase and sign, we can choose to set $f_{v_0} = 0$ for some singular vertex v_0 . If v'_0 is another singularity, and $f_{v'_0} - f_{v_0} = c$, the period of our texture must also be divisible by c , so that it is symmetric about $f_{v'_0}$ as well.

For foliations defined based on intersection numbers with one or more closed curves, heights of classes and “distances” between singularities will always be integers. So we choose our texture to be of the form

$$\text{color}(p) = g(\cos(2\pi f(p))),$$

where $f(p)$ is defined linearly on each triangle with $f(v) = f_v$ for each corner v , and g is an arbitrary function. Since the color function is periodic with period 1, and every curve class has integer height, the period condition is satisfied.

We compute the per-face texture coordinates by doing a breadth-first traversal of the dual graph. First, we pick one regular triangle t_0 and choose its texture coordinates. On each tree edge of the traversal, we assign coordinates to the target triangle so as to match the sign on the source triangle. At the end of the traversal, we apply a phase shift so that a zero (if one exists) has coordinate 0. This satisfies the symmetry condition.

Bibliography

- [1] Manfredo Perdigão do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice-Hall, Englewood Cliffs, NJ, 1976.
- [2] Michael S. Floater and Kai Hormann. Surface parameterization: a tutorial and survey. *Advances in multiresolution for geometric modelling*, 1(1), 2005.
- [3] Ofir Weber, Ashish Myles, and Denis Zorin. Computing extremal quasiconformal maps. In *Computer Graphics Forum*, volume 31, pages 1679–1689. Wiley Online Library, 2012.
- [4] Tsz Ching Ng, Xianfeng Gu, and Lok Ming Lui. Computing extremal Teichmüller map of multiply-connected domains via Beltrami holomorphic flow. *Journal of Scientific Computing*, 60(2):249–275, 2014.
- [5] Frederick P. Gardiner and Nikola Lakic. A synopsis of the Dirichlet principle for measured foliations. *preprint*, 2010.
- [6] Frederick P. Gardiner and Nikola Lakic. The Dirichlet principle for measured foliations. *preprint*, 2010.
- [7] Albert Fathi, François Laudenbach, and Valentin Poénaru. *Thurston’s Work on Surfaces*, volume 48. Princeton University Press, second edition, 2012.
- [8] John Hubbard and Howard Masur. Quadratic differentials and foliations. *Acta Mathematica*, 142(1):221–274, 1979.
- [9] Michael Wolf. On realizing measured foliations via quadratic differentials of harmonic maps to R-trees. *Journal D’Analyse Mathématique*, 68(1):107–120, 1996.
- [10] Vladimir V Fock and Alexander B Goncharov. Dual teichmuller and lamination spaces. *arXiv preprint math/0510312*, 2005.
- [11] Anil N Hirani. *Discrete exterior calculus*. PhD thesis, California Institute of Technology, 2003.

-
- [12] Keenan Crane, Fernando de Goes, Mathieu Desbrun, and Peter Schröder. Digital geometry processing with discrete exterior calculus. In *ACM SIGGRAPH 2013 Courses*, page 7. ACM, 2013.
- [13] Xianfeng Gu and Shing-Tung Yau. Global conformal surface parameterization. In *Proceedings of the 2003 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pages 127–137. Eurographics Association, 2003.
- [14] Steven J. Gortler, Craig Gotsman, and Dylan Thurston. Discrete one-forms on meshes and applications to 3D mesh parameterization. *Computer Aided Geometric Design*, 23(2):83–112, 2006.
- [15] Lok Ming Lui, Ka Chun Lam, Shing-Tung Yau, and Xianfeng Gu. Teichmüller extremal mapping and its applications to landmark matching registration. *arXiv preprint arXiv:1211.2569*, 2012.
- [16] Lok Ming Lui, Xianfeng Gu, and Shing-Tung Yau. Convergence of an iterative algorithm for Teichmüller maps via harmonic energy optimization. *Mathematics of Computation*, 2015.
- [17] Felix Knöppel, Keenan Crane, Ulrich Pinkall, and Peter Schröder. Globally optimal direction fields. *ACM Transactions on Graphics (TOG)*, 32(4):59, 2013.
- [18] Olga Diamanti, Amir Vaxman, Daniele Panozzo, and Olga Sorkine-Hornung. Integrable PolyVector fields. *ACM Trans. Graph.*, 34(4):38:1–38:12, July 2015. ISSN 0730-0301. doi: 10.1145/2766906. URL <http://doi.acm.org/10.1145/2766906>.
- [19] Felix Knöppel, Keenan Crane, Ulrich Pinkall, and Peter Schröder. Stripe patterns on surfaces. *ACM Trans. Graph.*, 34(4):39:1–39:11, July 2015. ISSN 0730-0301. doi: 10.1145/2767000. URL <http://doi.acm.org/10.1145/2767000>.
- [20] Lars V. Ahlfors. *Complex Analysis*. McGraw-Hill, New York, NY, third edition, 1979.
- [21] Jeremy Kahn and Maxime Fortier Bourque. Teichmüller theory and dynamical systems. https://wfs.gc.cuny.edu/mfortier_bourque/www/notes/teich2013/TTDS_Fall2013.pdf, Fall 2013. Course notes.
- [22] Dylan P. Thurston. Personal communication, 2015.
- [23] Jan Möbius and Leif Kobbelt. OpenFlipper: An open source geometry processing and rendering framework. In *Proceedings of the 7th International Conference on Curves and Surfaces*, pages 488–500, Berlin, Heidelberg, 2012. Springer-Verlag. ISBN 978-3-642-27412-1. doi: 10.1007/978-3-642-27413-8_31. URL http://dx.doi.org/10.1007/978-3-642-27413-8_31.