

Optimal Constrained Hyperbolic Surface Mappings

By

ALEXANDER TSUI
B.S. (San Jose State University) 2010

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

COMPUTER SCIENCE

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

Nina Amenta, Chair

Owen Carmichael

Patrice Koehl

Committee in Charge

2016

To whom it may concern

Contents

1	Introduction	1
2	Mathematical background	10
2.1	Geometry	10
2.1.1	Distortion	12
2.1.2	Surface maps	14
2.1.3	Minimal distortion maps	18
2.1.4	Conformal geometry	24
2.1.5	Hyperbolic geometry	26
2.1.6	Orbifolds	30
2.1.7	Harmonic Map – Existence and Uniqueness	33
2.2	Related work and applications	34
3	Hyperbolic Orbifold Maps	37
3.1	Introduction	37
3.2	Prior work	39
3.3	Method	40
3.3.1	Orbifold Mapping	40
3.3.2	Marking selection	43
3.3.3	Initial Mapping	43
3.3.4	Energy minimization	44

3.4	Experiments	45
3.4.1	Data	45
3.4.2	Competing methods	45
3.4.3	Performance measures	46
3.4.4	Experimental settings	47
3.5	Results	47
4	Curve Extension For Hyperbolic Orbifold Maps	53
4.1	Introduction	53
4.2	Prior work	56
4.3	Method	58
4.3.1	Theory and Curve Landmarks Extension	58
4.3.2	Algorithm	59
4.4	Results	65
4.4.1	Datasets	65
4.4.2	Comparison methods	66
4.4.3	Energy minimum robustness to differences in initial conditions	67
4.4.4	Comparing distortion and landmark matching across methods	67
4.4.5	Generalization of maps to unseen landmark mismatch across methods	70
4.4.6	Biological relevance of map distortion measures	70
4.4.7	Mapping in the presence of anatomical uncertainty	73
4.5	Discussion	74
5	Towards Angle Distortion-Minimizing Maps	88
5.1	Theory and Prior Work	89
5.1.1	Dilatation	89
5.1.2	Conformal maps	91
5.1.3	Prior approaches	92

5.2	Methods	97
5.2.1	OrbifoldPoints	97
5.2.2	OrbifoldBeltrami	98
5.2.3	OrbifoldAngleShare	102
5.2.4	OrbifoldMIPS	103
5.3	Results	105
5.3.1	Characteristic behavior of angle optimization on 2D grid	105
5.3.2	Characteristic behavior of angle optimization on 3D surface	106
5.3.3	Relative behavior of angle optimization across real-world datasets	111
5.3.4	OrbifoldAngleShare is robust to length-only deformations	118
5.3.5	Discussion	118
A	Technical Overview Of Mapping Software	125
A.1	Prerequisites	127
A.2	Conventions	127
A.3	Usage	127
B	List of Mapping Software	128
6	References	133

Optimal Constrained Hyperbolic Surface Mappings

Abstract

Constrained surface mapping is an important component in applications ranging from biomedical shape analysis to shape deformation. There traditionally has been a tradeoff between optimality and accuracy of landmarks that has been addressed in an ad-hoc manner. Orbifolds are topological surfaces that have nice properties that we find useful and applicable to constrained surface mapping. Particularly, the identification of special cone points on orbifolds with surface landmarks allows us to construct a map with exact landmark matching, and the introduction of a hyperbolic metric allows us to apply a theorem that guarantees that such a map is unique. The following work is about how to set up the orbifold structure needed to match either point or curve landmarks, and how to approximate such unique surface maps that reduce various forms of distortion. We show some empirical results that demonstrate the technical and biological relevance of the method with respect to other methods that use similar approaches.

Acknowledgments and Thanks

First, I want to thank the committee, specifically the core biogeometry group formed by the committee and Joel Hass, for their lively discussions, insightful observations and advice, and unending patience. Being a part of this cross-disciplinary project, I had the perfect opportunity to put everything I learned in math and computer science into use, and working with such an enthusiastic group of people was no doubt one of the reasons I was able to make it this far. I couldn't have asked for a better group of mentors.

My work was made possible with the help of many others, and it is a privilege to have the opportunity to speak with some of the giants who boost me up on their shoulders through their work. I want to thank Lok Ming Lui, Ying He, Yalin Wang, and Feng Luo for the time they had taken to speak to me in regards to technical details on methods and research prospects. I want to acknowledge the open source community, particularly the developers with Computational Geometry Algorithms Library, which has been vital for my own work and undoubtedly so for the community of computational geometry researchers at large. I had the privilege of personally working with some of the developers and intend to use my capacities to contribute to this and other projects as I move forward. I also want to acknowledge the students who I had the pleasure to work with at Davis to utilize early versions of my software to collect data – thank you to Casey Wilson, Melissa Chenok, Devin Fenton, and Jessica Liu for your efforts.

I want to acknowledge all of the moral support that I received in my time at Davis. There are points during graduate school where I wonder if I would be able to continue, if it weren't for this dependable foundation that kept me going. I am glad that you were part of my experience here. Whether it was for the insightful discussions about the industry or academia, or whether it was simply for the time spent sharing meals or playing games, I am

grateful to have been in your company. In particular, I wanted to thank the colleagues who I have all worked with directly or indirectly, who helped me navigate graduate school and collectively made me feel like I was valued part of a research community. I want to thank Sean Gilpin and Kate Isaacs as fellow alumni from San Jose State University for offering their perspectives and insights before and during my time as a graduate student. Thanks to Chris Schwarz, Phong Vuong, and Kris Kalish with the Carmichael Lab, and Sam Lockhart, Alexandra Roach, Baljeet Singh, Evan Fletcher, Charlie DeCarli, and everyone with the IDeA Lab for getting me started and getting my feet wet. Thanks to Stewart He, Lance Simons, Shengren Li, and Carlos Rojas with the Geometry Lab for a true lab experience complete with chasing deadlines, poster sessions, and Uno and Magic games. Thanks to Armen Khodaverdian, Gabriel Maganis, Martin Velez, Wencin Poh, Minh Tran, and the Iwamoto family for being great roommates and even greater friends.

Finally, I would like to recognize the unconditional support from my family. It means a lot to me to have someone to rely on, without having to worry about justifying myself. I want to thank them for their kindness and patience through the past several years, which has remained a steady constant throughout the highs and lows of graduate school.

Chapter 1

Introduction

Maps are useful tools that gives the reader a bird's-eye view of a particular region of interest and, in doing so, orients and gives a sense of direction through the simplified, alternative model of the actual area. Maps have existed since ancient times, with the earliest wall paintings dating back to the 7th millenium BCE. The makings of a rigorous, mathematical approach to mapmaking probably began with Ptolemy when he proposed a world map projection (Figure 1.0.1) in his famous book *Geographia*. In a time where areas of the world had been scaled by political importance, the fact that Ptolemy provided a projection that preserves lengths and angles relatively well (albeit only near the central latitude) and made use of latitude and longitude lines to construct a global coordinate system for locating places of interest precisely, marks a point in history where mapmaking became more systematic. Techniques of physically creating flat maps of the known world also became more sophisticated as time went by, for instance, with the Mercator projection introduced by Flemish geographer Gerardus Mercator in 1569. At the same time, the means of measuring the physical world also became more sophisticated. In the past, we relied on optical technology such as the telescope and the sextant to survey the land, and with the development of printing technology, we could put to paper finely detailed maps that are more accurate in the sense that they physically don't distort and wear out as easily. Nowadays, we not only have access



Figure 1.0.1: A 15th century reconstruction of the map described by Ptolemy in Geographia (circa 150).

to aerial photography, satellite imaging, and remote sensing, which gives us a high degree of precision, but we also have access to hardware that allows us to perform mapping on-site, and computers to visualize and overlay analytics dynamically. Furthermore, we are not restricted to what we can gather with our naked eyes from the visible portion of the electromagnetic spectrum. Laser technology allows us to generate depth information with great accuracy, for instance by scanning physical objects of interest, or perhaps surveying wide swaths of terrain with a airplane-mounted laser, to produce point clouds that can be reconstructed into manifolds. Techniques such as magnetic image resonance imaging allow us to inspect internal structures of living organisms in a noninvasive manner. Suddenly, it is possible with the aid of computers to reconstruct anatomical manifolds of domains that are otherwise inaccessible to the naked eye, consequently opening up the possibility of producing maps across such surfaces, for the purpose of performing high-level tasks such as visualization, statistical analysis, and other applications.

Whether a map is a simple caricature of what it is meant to represent (such as in a hand-drawn map) or whether it tries to be as close to reality by precisely and accurately capturing information, a map essentially serves the function of indicating a *correspondence* – specifying that one point in one region is identified with another point in another region. The utility of solving this *correspondence problem* is clear considering the purpose of maps to aid in navigating between different locations, but how to answer the question of specifying which point goes to where is also an important problem in computer science. For instance, in archaeological morphometrics, we have the problem of having in hand some similar fossil specimens for which we do not know the correspondence. Consider that the skulls of two different species of monkeys may have drastic differences in size yet still exhibit common features that the field experts may decide to annotate as corresponding labeled points. Given a sparse set of points that correspond between two shapes, the *correspondence problem* asks the question: "Where does the rest of the surface correspond to?" We would like to be able to examine some unspecified region on one skull and see what the matching region on the other

skull is. Constructing a map is how we explicitly answer this question, and the algorithm for doing so should produce a map that is correct and intuitive in the context of the given surfaces.

Aside from navigation, there are a plethora of applications and computational geometric techniques that open up when we have access to a surface map. We can impose a correspondence across a set of surfaces that we are interested in studying and compare where individuals features differ and coincide. We can do interpolation of functions in order to overlay or transfer data from one surface onto another, whether that is color or motion data. This has a lot of practical uses in graphics and animation, allowing for such tasks as keyframe interpolation, transfer of attributes such as color and pose, and shape deformation. As a corollary, we can perform statistical analyses of surface-defined functions, examining how certain examples behave in certain regions of the surface, with respect to the average value. A subfield has emerged in biomedical image analysis where anatomical models are extracted from patient data and analyzed for biological relevance.

The quality of said applications depends on how “good” our maps are, so we will try to motivate the discussion with an anecdotal example of mapping a region of the Earth to the plane. Suppose we were given the task of mapping a large region, such as the state of California, and consider for the sake of discussion, the familiar-looking map in Figure 1.0.2 as a possible solution to this task.

Note we have a difference in scale, as the entire state of California is visible from on a single sheet of paper. Here, the scale is $1 : 2,000,000$, which means that an inch on the paper corresponds to two million inches (just over 31 miles, or just under 51 kilometers). Measuring the straight-line distance between any two points on the map, and then multiplying by the scale of the map, tells the reader how to retrieve the true distance between the corresponding two points in California. In the ideal case, the difference between this computed distance and the actual physical distance will be the same for any pair of points on the map. This is an example of *metric distortion*, whereby a map changes the distance between two points. In



Figure 1.0.2: USGS relief map of California, 2001.

some sense, if the only thing that differs between the map and the reality that it represents is the scale, then the map is ideal. However, this is clearly not the case if we do the mental exercise of “reversing” the transformation and scaling the flat map to the size of California. In this case, the sense of elevation has been lost when transferring actual terrain onto a flat medium, so a straight line as measured on the flat map cannot possibly account for the vertical increases and decreases in elevation along the corresponding path in physical space. A further issue arises when one considers that the physical region encompassed in the map does not represent a flat sheet but rather a curved portion of a spherical surface. This means that even if elevation was uniform across the neighborhood of California, the distance between two points in California cannot be exactly computed by scaling up the straight-line distance on a flat map of California. Maps distort reality in order to give alternate representations of the region that they are meant to capture, and as such, the maps hold intrinsic geometric information about the surfaces. By studying these intrinsic deformations, we can understand to what degree the mapped shapes differ from what is actually there. Conversely, the map distills reality down to certain crucial pieces of information, and it is important to understand how exactly this information remains in line with what is actually there. Whether the task is to match a spherical patch to a flat patch, or it is something more complex such as the fossils of two different species of monkeys, the idea of metric distortion is always part of the discussion of maps.

To make the idea of metric distortion more precise, we need to fix our data representation. Let us consider the map of California. In the context of geographical information system (GIS) data, there are a number of ways to represent the world, and depending on the context and application, each of the many representations can lend itself to the task at hand; however, the project that we will discuss in this thesis concerns the mapping and analysis of three-dimensional surfaces, and as such, the most pertinent of these data structures is known in the field of GIS as the triangulated irregular network (TIN). In general applications, this is known as a triangulated surface mesh, or simply triangulation, and basically, it is a collection

of points in three-dimensional space, together with local topological information expressed as a list of ordered triplets of said points. The points of the mesh are called vertices, and the ordered triplets, known as the faces of the mesh, specify which points are locally connected. The points are a discrete sampling of a surface, and assuming a sufficient sampling rate, any surface of interest can be represented faithfully in a computer as a list of points and triangles.

Now that we have fixed a data representation, we can begin to talk about the concept of a metric on the surface. Locally, distances between neighboring points can be measured by simply taking the length of the connecting line between them. If we suppose we have a discrete representation of the state of California as a network of points connected into triangles, then we can visualize a physical process that deforms and compresses this network of points so that it fits in the planar bounds of a sheet of paper. We have introduced a distortion to the metric in the sense that all pairwise distances have changed. From our motivating example, we understand that there is a scale that is necessary to transform between the map and physical space; however, the differences in the metric that remains after such scaling is what we are interested in and refer to as metric distortion introduced by the map.

Maps can be seen as deformations, which can range from transformations parameterizable as linear transformations determined by a matrix of parameters, to an arbitrary specification as a non-rigid deformation. See [3] for more details. If maps are deformed, then the algorithm that generates such a map is responsible for any deformation that is added. The intuition from data modeling that a parsimonious model should be favored when fitting to data, and similarly, it is desirable to be conservative when introducing deformation, warping the surface only when it is necessary. We can factor in geometric data to guide in matching, which could be given explicitly if we use landmark data to specify that certain areas must be in correspondence. Alternatively, we can define some matching function over the surface and perform an energy optimization with respect to that function. Whatever data is ultimately introduced in a non-rigid registration attempts to guide map construction by strictly

requiring or suggesting certain areas be kept in alignment.

This thesis presents a surface mapping framework that takes into account the considerations of metric distortion minimization and landmark matching and provides a canonical solution that is applicable to general surfaces. A brief summary of the contributions of this body of research is as follows. This thesis explores a particular strategy of constrained surface mapping that utilizes a novel mathematical construction called the *orbifold*. At the start of the project, this was the first known application of orbifolds to surface mapping. We applied hyperbolic orbifolds to constrained surface mapping with landmark points. We extended hyperbolic orbifold construction to allow for application to constrained surface mapping with landmark curves. We performed some validation experiments to demonstrate the discrete approximation follows closely to the behavior that we expect from the theory. We experimented with various quantities of metric distortion, exploring length distortion in a least-squares sense and angle distortion in a least-squares and maximum sense. We implemented these experiments in a framework that allows us to compare it to similar approaches that are based on conformal maps, and demonstrates that our proposed method performs comparably, and sometimes favorably.

The thesis is organized as follows. In Chapter 2, we will discuss some preliminaries and general mathematical background to provide some context to understand the rationale and approach in the studies described in this thesis. In Chapter 3, we introduce the concept of orbifolds and our novel application of orbifolds to surface mapping with point-based landmark constraints. The mapping framework utilizes hyperbolic orbifolds to produce surface maps that are optimal in the sense that the process minimizes an energy functional that is uniminimal, resulting in a map in which overall length distortion is minimized in a least-squares sense. We present some experiments performed on a cortical surface dataset that demonstrates the technical and biological relevance of the method in surface-based biomedical image analysis, as well as compare the method to similar state of the art mapping techniques. Then in Chapter 4, we extend the orbifold mapping framework described in

the preceding chapter to overcome the apparent limitation of orbifold mapping to be able to match only point landmarks. We present a novel construction that allows for surfaces with curve landmarks to be mapped such that the curves matched exactly. We perform experiments on an extensive collection of real-world datasets, including cortical surfaces, human faces, and monkey skulls. The experiments demonstrate the practical behavior of the curve-based landmark matching method, a modest improvement of landmark matching over similar prior methods, and an ability to detect differences between different groups within each dataset. Finally, in Chapter 5, we present an alternative approach to orbifold map optimization that minimizes overall angle distortion. We show empirical examples where the angle-optimizing method produces characteristically different results from the method described in preceding chapters. The experiments show that angle- and length-based optimizations can form the basis of for shape difference measures that give complementary information. In Appendix A, we give some supplementary technical information about the software implementation of the surface mapping framework described in this thesis.

Chapter 2

Mathematical background

2.1 Geometry

In the context of surface mapping, we have shapes and objects that are distinguished by the individual lengths, angles, and areas of their distinctive parts, so our problem is one that is fundamentally concerned with *geometry*. This branch of mathematics is concerned with the questions of shape, size, and properties of space, and it focuses on the information gained by examining the metric, a quantitative concept of distance between two points in the underlying space. It arose from practical considerations of measuring lengths, angles, and areas and forms a basis for many other branches of mathematics. Historically, Euclid in the sixth century BC put geometry in rigorous form, giving the subject a comprehensive treatment built on a number of axiomatic statements:

1. Two points determine a line segment
2. A line segment can be extended indefinitely along a straight line
3. A circle is determined by a line segment by taking it as its radius
4. All right angles are congruent
5. There is a unique line that is parallel to another line and passes through a given point.

Consider in a discrete computational setting, we can represent finite regions of Euclidean space as a *triangulation*, a special type of graph with triangular faces. To fix notation, consider a triangulation T to be a collection of $|V|$ vertices, $|F|$ triangles, and $|E|$ edges. A triangulation allows us to conveniently and simply define notions of length and angle by considering the measurements between locally adjacent mesh elements. The distance between two adjacent vertices v_i and v_j is computed by measuring the length ℓ_{ij} of the edge e_{ij} connecting them. This assignment of distance is called a *Euclidean distance metric* ℓ precisely if for a face f_{ijk} that has vertices i, j , and k , the following equations hold

$$\begin{aligned}\ell_{ij} &\leq \ell_{jk} + \ell_{ki}, \\ \ell_{jk} &\leq \ell_{ki} + \ell_{ij}, \\ \ell_{ki} &\leq \ell_{ij} + \ell_{jk}.\end{aligned}\tag{2.1.1}$$

A Euclidean triangulation is a triangulation T together with a Euclidean *distance metric* ℓ . A particular Euclidean triangulation is denoted (T, ℓ) , and note that as long as the triangle inequality is satisfied, ℓ can be replaced with another metric $\tilde{\ell}$ to give a new Euclidean triangulation. Two Euclidean triangulations are said to be *combinatorially equivalent* if they differ only by the distance metric. If we know all the lengths of a triangle, then the angles are fully determined, and we can denote an angle by α_{jk}^i , which indicates the angle α in triangle ijk at vertex i . Note that the order of the vertices that specify a triangle are important. By convention, the vertices are always oriented in a counterclockwise fashion with respect to the “top” side of the triangle – that is, using the right hand rule by circulating the vertices i, j , and k , we get a normal vector pointing out of the side of the triangle facing “up”, or “out” if the triangle is part of a closed, orientable surface. Following this convention gives a consistent orientation of all the triangles in the triangulated region of Euclidean space. This representation is applicable not just 2D planar regions but general 2D surfaces embedded in 3D space as well. A general surface admits a triangulation by the following argument:

every surface patch is associated with a flat parameterization, in which every triangle in the triangulation is contained – see [1] for more details. For a details about a software implementation of a data structure that expresses this surface representation, please refer to the technical overview in Appendix A.

2.1.1 Distortion

From our earlier motivation in the introduction, we understand that in the context of surface mapping, aside from functioning as a correspondence, a map is also a deformation of an original shape into a different shape, and distortion quantifies the degree to which geometry is changed locally. There are many ways to distort a pliable shape. Consider by analogy, changing the shape of a pliable piece of clay by moving, stretching, twisting, and physically manipulating it. Now suppose a fine wire mesh of triangles is embedded into the clay and examine that the local changes in this mesh capture the deformation of the clay from its initial configuration to its final deformed configuration.

A deformation in the most general sense is a combination of any number of these shape-changing actions, and by restricting ourselves to only certain types of actions, we get subclasses of deformations that preserve particular properties and thus are given special names. First, consider maps restricted to the Euclidean plane in Figure 2.1.4. To introduce one special subclass of deformations, an *isometry* is a map that does not introduce length distortion. Translation and rotation combine to form a subset of isometries, and they also form an *equivalence class* of shapes that differ only by their rotation and translation in space.

A *similarity* encompasses a broader class of transformations that allows uniform scaling and reflection in addition to the rigid motions. The lengths may be uniformly changed, but the angles remain congruent. Any map that preserves angles is called *conformal*. Similarly to rigid motion, the set of conformal transformations of a particular shape also forms an equivalence class. For any pair of surfaces in this equivalence class, the angles are preserved and are considered equivalent if we are comparing the surfaces only by their conformal

structure i.e. their angles. However, as it is a much larger parameter space compared to the isometries, there is much more variation in the types of shapes when considering them from the point of view of lengths.

In general, an *affine transformation* need not preserve lengths nor angles, as the only requirement is that straight lines are mapped to straight lines. Certainly, affine transformations do not characterize the set of all possible planar maps, but considering a discrete representation of space in a computational setting is often given as a triangulation, it is useful to consider general maps as approximately represented as a collection of affine transformations restricted to corresponding pairs of triangles. That is, the two surfaces are endowed with a combinatorically equivalent triangulation, where each triangle on one surface corresponds to one triangle on the second surface, which fixes an affine transformation for that restricted region. Thus, by finer subdivision of space, we can express the mapping with greater precision and accuracy. As we define the map by piecewise linear maps, we can quantify distortion in a piecewise manner over the triangulation. Then length distortion is given by how much the original length l_{ij} between vertices v_i and v_j deviates under the mapping. While this can be considered between any pair of vertices, typically, length distortion considered over the set of edges E between adjacent vertices serves as a representative measure of how the map preserves lengths.

Similarly, we can quantify angular distortion by considering the collection of corresponding angles in similar fashion. Therefore, an angle preserving map will seek to minimize the change between corresponding angles. A more general way to state this is, an angle preserving map tries to preserve the distribution of angles around each vertex [25]. This particular view is useful for when we consider maps between surfaces embedded in space, where the total angle around each vertex varies. Yet there are other ways to quantify angle distortion that measure it across each triangle rather than vertex. Consider that the singular value decomposition of a linear map between two triangles yields two singular values Γ and γ , representing the major axes of scaling. Their ratio Γ/γ is called *dilatation* and expresses angle

distortion by how much an infinitesimal circle is deformed into an ellipse. In particular, an orthogonal transformation will have exactly a dilatation of 1.

Additionally, it is possible to consider area distortion by measuring change of area between triangles. It is known that a triangulation is fully determined by its angles together with its lengths, or its angles together with its areas. We take the approach of examining length and angle distortions only.

Homotopy

In the above discussion, we mention that the class of deformations on a surface forms an equivalence class. From a particular representative surface, a collection of new surfaces can be obtained by applying any number of deformations from the determined class of deformations to the surface. Consider now fixing the target surface and, instead, deforming the image of the source surface on the target surface, and in addition, suppose that some points remain fixed. This forms the basis of another equivalence class called *homotopy*. Two maps are called *homotopic* relative to the fixed points if there exists a continuous deformation from one to the other. As an example and non-example of homotopic curves, consider Figures 2.1.1 and 2.1.2.

This concept is important in the mapping relaxation when we have a surface with a *marking* – a set of curves with distinguished endpoints that are fixed in the mapping. From a given marking, consider continuously deforming the curves while fixing the endpoints. Then the resulting marking and the original marking are said to be homotopic relative to their endpoints. In our surface mapping approach, we are interested in minimizing the distortion of maps with respect to a given homotopy.

2.1.2 Surface maps

In order to discuss distortion on surfaces embedded in space, we need to consider *surface maps* that connect the simple 2D parameter space on which we can readily compute distortion,

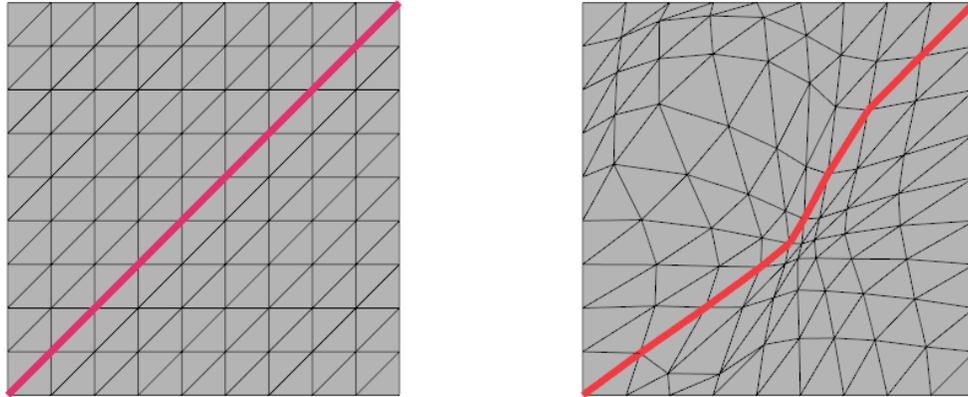


Figure 2.1.1: A homotopy between two squares in \mathbb{R}^2 that have been marked with a single curve whose endpoints are fixed. The curves are homotopic relative to their endpoints because there exists a continuous deformation from one to the other.

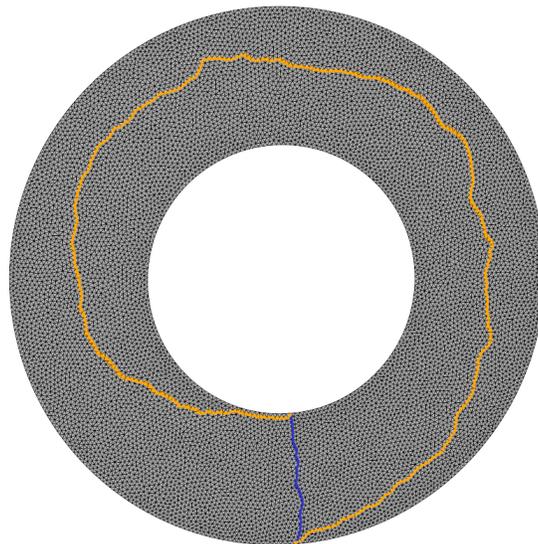


Figure 2.1.2: A non-example of a homotopy. Note that there is no way to continuously deform the blue curve to the orange curve while keeping the endpoints fixed on this annular mesh. Each curve specifies and belongs to a different homotopy class.

with the curved surface patch we are interested in. A surface map is a formal construction studied in *differential geometry*, a branch of math dedicated to answering questions of curves and shapes embedded in three-dimensional Euclidean space. Arising in the later 1800s, its early development is attributed to Leibniz [60], and its development is around this time was paralleled by the emergence of new and alternate theories of geometry. Informally, a manifold is a topological entity that locally resembles Euclidean space. We can picture a 2D sheet in a 3D space, and we can call that sheet a manifold because we can associate to every local patch on the sheet a flat patch of 2D Euclidean space. Generally speaking, the surface does not need to be developable from a sheet and oftentimes is very different from a sheet, but the intuitive requirement is that the surface is a collection of patches, each of which can be associated with a sheet locally. Now, given two manifolds M and N , a surface map $f : M \rightarrow N$ is a bijection from one manifold to the other. If f and its inverse are differentiable, then f is called a *diffeomorphism*. Now, each point on the surface also has an associated tangent plane, which is the best fit linear approximation of the surface at that point, and we can define the dot product of vectors in this tangent space, called the *first fundamental form* as

$$I_p(w) = \langle w, w \rangle_p = |w|^2. \quad (2.1.2)$$

For instance, suppose we have a surface with regular parameterization $x(u, v)$ and a curve $\alpha(t) = x(u(t), v(t))$ on the surface, then we consider a tangent vector to the curve and evaluating the first fundamental form:

$$\begin{aligned} I_p(\alpha'(0)) &= \langle \alpha'(0), \alpha'(0) \rangle \\ &= \left\langle \frac{dx}{du} \frac{du}{dt} + \frac{dx}{dv} \frac{dv}{dt}, \frac{dx}{du} \frac{du}{dt} + \frac{dx}{dv} \frac{dv}{dt} \right\rangle \\ &= \langle x_u du + x_v dv, x_u du + x_v dv \rangle \\ &= \langle x_u, x_u \rangle du^2 + 2 \langle x_u, x_v \rangle dudv + \langle x_v, x_v \rangle dv^2 \\ &= Edu^2 + 2Fdudv + Gdv^2. \end{aligned}$$

The functions E , F , and G are the coefficients of the first fundamental form, and it describes how the infinitesimal patch $dudv$ bends in the principal directions. Having this quantity defined at each point on the surface allows us to compute lengths, angles, and areas intrinsically by integrating over the curves and the regions that they bound. Said another way, the intrinsic distortion between two manifolds can be understood by examining the first fundamental form of the map.

As an example, a map $f : M \rightarrow N$ is an *isometry* for any pairs of vectors w_1, w_2 in the tangent plane, we have

$$\langle w_1, w_2 \rangle_p = \langle df(w_1), df(w_2) \rangle_{f(p)} \quad (2.1.3)$$

This means that lengths are not changed when passing under the map f . By considering the first fundamental forms of surfaces, we are able to compare intrinsically how different surfaces are. The interested reader can find a more in-depth exposition on differential geometry in [14].

A more general way to convey this intrinsic surface information is by way of *Riemannian geometry*. A central feature of this theory is the notion of *Riemannian metric*, which specifies the same information as the first fundamental form by assigning to each point on the manifold a positive symmetric definite matrix g , where

$$g = \begin{pmatrix} E & F \\ F & G \end{pmatrix}. \quad (2.1.4)$$

This defines a dot product in the tangent space exactly as above, and as long as g varies smoothly along a given manifold M , then we say that (M, g) is a Riemann manifold. In general, the Riemann metric is generally known as a metric tensor, and the entries of the tensor captures all information about curvature at a given point. Gauss's Theorema Egregium formalizes this idea, allowing us to make such statements as

1. Wrapping a sheet of paper into a cylinder is an isometry.

2. No piece of the plane can be bent into a piece of a sphere without distorting lengths.

Incidentally, Bernard Riemann, the mathematician who developed *Riemannian geometry*, also developed elliptic geometry, a system of geometry in which parallel lines do not exist, and the idea of Riemannian geometry is to unify non-Euclidean geometries like this and others with the notion of the Riemannian metric. Hyperbolic geometry, to be described in Section 2.1.5, is another system of geometry where the space has constant negative curvature.

2.1.3 Minimal distortion maps

In the previous sections, we introduced how to quantify length and angle distortion and a couple of classes of maps in which there is no distortion for either length, namely isometries, or angles, as in conformal maps. While these do form a broad class of maps, they are relatively restrictive considering the full space of possible deformations. Consider a map of a rectangular region of the Euclidean plane \mathbb{R}^2 to a circular region in \mathbb{R}^2 . If we consider the Euclidean metric from the background space, then we see that an isometry is not possible. According to the Riemann mapping theorem, however, it is possible to construct a conformal map from the square to the circle. However, if we were to select the image of each of the four corners of the square arbitrarily along the boundary, then the map cannot be conformal. In general, we will be faced with a particular boundary condition for which we will not be able to achieve a distortion-free map, but we want to construct a map that minimizes distortion. In this situation, we would like to consider distortion as some integral sum over the domain and minimize that. It is natural to consider length and angle distortion, and in fact if we consider these distortions under the L_p norms for particular values of p , we encounter some more well-known surface maps that are summarized in the following table:

	L_2 norm	L_∞ norm
Length distortion	Harmonic map	Absolutely minimizing Lipschitz extension
Angle distortion	-	Teichmüller map

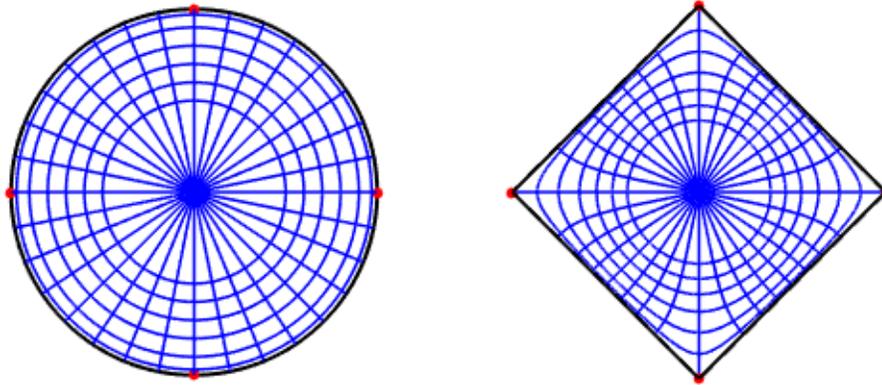


Figure 2.1.3: A conformal map from the square to the disk. As a conformal map is determined by the image of three points, supposing we fix two of the red points and the origin of the disk, if the remaining two red points are required to map elsewhere along the boundary, then the resulting map cannot remain conformal.

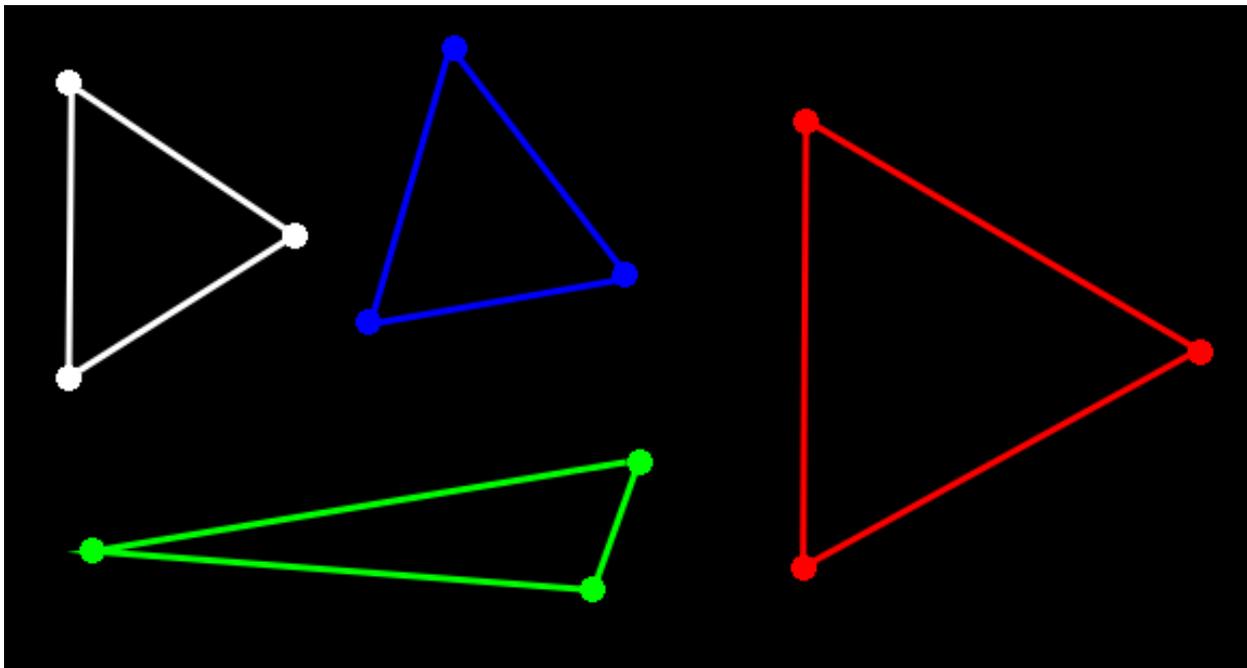


Figure 2.1.4: Basic examples of planar distortion. The combination of translation and rotation constitutes a rigid motion in which the lengths and angles of a triangle are preserved. A scaling map uniformly distorts all the lengths of a triangle, but the triangles remain congruent. In general, a map needs not preserve lengths nor angles.

Harmonic maps

Harmonic maps are defined to be those maps in a collection of maps between two given surfaces A and B that minimize overall length distortion in an L_2 , least-squares sense, and this distortion is quantified by the Dirichlet energy. The Dirichlet energy, also referred to as harmonic energy, is a natural measure of the amount of stretching induced by a surface map. In the discrete setting, a map f that deforms a Euclidean triangulation has Dirichlet energy

$$E(f) = \sum_{i,j} w_{ij} |f(v_j) - f(v_i)|^2 \quad (2.1.5)$$

where w_{ij} are the discrete cotangent weights

$$w_{ij} = \frac{1}{2}(\cot \alpha + \cot \beta) \quad (2.1.6)$$

with α and β being the two angles opposite of the edge between vertices v_i and v_j on the domain surface. The discrete cotangent weights actually make an appearance as early as 1949 [43] in the context of the solution of Laplace's equation on in asymmetric network of triangles, but in the discrete differential geometry literature, Pinkall and Polthier [48] formalized the discrete Dirichlet energy, showing how the cotangent terms arise from considering Dirichlet energy of a linear map between two triangles.

One key idea in surface mapping is that surface maps are connected by continuous deformation by a relation called *homotopy*. In the same principle that a taut mesh that is perturbed by an external force wants to restore itself to a relaxed configuration when the external force is released, we can perform energy optimization on a given surface map as a way of searching for a more optimal solution. The existence and uniqueness of such an optimal solution depends on the parameter space under consideration, but it is noteworthy that a hyperbolic parameter space ensures that the resulting map achieves the global minimum of the length distortion we have formulated. The Eells-Sampson theorem [15] states that there exists a harmonic map (i.e. a map that minimizes overall squared average length distortion)

between any pair of hyperbolic surfaces, and Hartman [22] proved that such a map is unique.

Teichmüller maps

As mentioned before, the extent to which conformal maps are able to satisfy arbitrary constraints and boundary conditions is limited. To be precise, in the continuous case the image of three points in the complex plane determines a conformal map. Therefore, additional points must then be fitted in a least squares sense, if the map must remain conformal. Alternatively, the natural extension beyond conformal maps are *quasiconformal maps*, which are maps that introduce a bounded amount of angular distortion. Formally, a planar map is called *quasiconformal* if it satisfies the *Beltrami equation*

$$\frac{\partial f}{\partial \bar{z}} = \mu(z) \frac{\partial f}{\partial z} \quad (2.1.7)$$

where the complex-valued function μ that satisfies $\|\mu\|_\infty < 1$ is called the *Beltrami coefficient*. Recall from the discussion about quantifying distortion from Section 1 that angular distortion can be quantified by *dilatation*, a description of how much a map distorts an infinitesimal circle into an ellipse (see Figure 2.1.5).

Dilatation K is related to the Beltrami coefficient by

$$K = \frac{1 + |\mu|}{1 - |\mu|}. \quad (2.1.8)$$

Note that a conformal map sends circles to circles, thus $|\mu| = 0$. Maps that are less conformal map circles to ellipses where the long axis becomes increasingly larger than the short axis, and it approaches sending a circle to a line as $|\mu|$ approaches to 1.

In theory, we know of the existence of the Teichmüller map, or the map that minimizes the L_∞ norm of $|\mu|$. This means that the maximal angular distortion throughout the map has been minimized. There have been several attempts to compute this extremal quasiconformal map, which we will discuss in some detail later on.

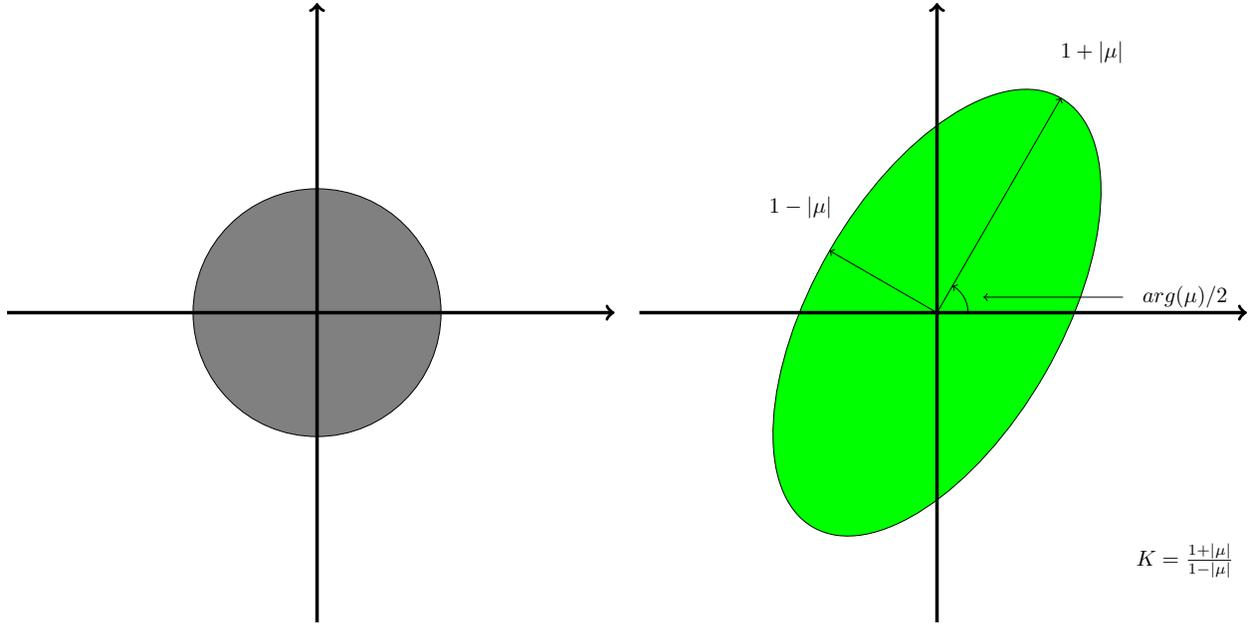


Figure 2.1.5: The dilatation is quantified as a ratio of the major over the minor axis of the image ellipse of an infinitesimal circle under a given mapping f .

To the best of our knowledge, there is no method that attempts to compute a map that minimizes the L_2 norm of $|\mu|$ rather than the L_∞ norm, which would be analogous to what the harmonic map does for isometric maps. We pursue this idea experimentally in a later section.

Absolutely minimizing Lipschitz extensions

Analogous to considering extremal maps that minimize the maximal angular distortion throughout the map in the previous section, it is natural to ask questions about maps that minimize the maximal length distortion. The harmonic map, which can be understood to minimize stretching of the domain under the mapping, minimizes an integral sum of length distortions averaged over the domain but does not necessarily account for distortion in an absolute, global sense. However, these two concepts are closely related by a particular family of functionals taken over surface \mathcal{B}_1

$$\mathbf{J}_p(\phi) = \left(\frac{1}{\mu(\mathcal{B}_1)} \int_{\mathcal{B}_1} \|D_{\mathcal{B}_1} \phi(x)\|_2^p \mu(dx) \right)^{1/p} \quad (2.1.9)$$

where $D_{\mathcal{B}_1}$ denotes intrinsic distances on the surface \mathcal{B}_1 . Note that for $p = 2$, $\mathbf{J}_p(\phi)$ is the Dirichlet functional and minimizers are harmonic maps. Taking p to infinity, the functional gives rise to a special notion of distortion known as the **Lipschitz constant**

$$\mathbf{J}_\infty(\phi) := \lim_{p \uparrow \infty} \mathbf{J}_p(\phi) = \operatorname{ess\,sup}_{x \in \mathcal{B}_1} \|D_{\mathcal{B}_1} \phi(x)\|_2 \quad (2.1.10)$$

In the equation above, the quantity on the right hand side is the Lipschitz constant of the map ϕ . This measure of distortion can be understood as the maximum of local metric distortion.

Note that we are also interested in maps that satisfy correspondence constraints between surfaces. This can be posed as follows: for two surfaces \mathcal{B}_1 and \mathcal{B}_2 we have landmarks $\Gamma_1 \subset \mathcal{B}_1$ that correspond to landmarks $\Gamma_2 \subset \mathcal{B}_2$. This determines a partial mapping $f' : \Gamma_1 \rightarrow \Gamma_2$. A continuous extension to the rest of \mathcal{B}_1 such that the distortion is minimized is known to exist [4]. This is an instance of a classical problem in real analysis, where a function defined on the boundary of a bounded, open subset of \mathbb{R}^n is extended to whole domain without increasing its Lipschitz constant. It turns out that the functional in Equation 2.1.10 is the Euler-Lagrange equation for this problem, and maps that are solutions to this problem are called *minimal Lipschitz extensions*. Theory tells us that these maps exist, but in general, they are not necessarily unique nor smooth. However, *viscosity solutions* provide a viable notion of weak solution, and a unique solution is guaranteed for this class of maps.¹ Maps of this type are called *absolutely minimizing Lipschitz extensions* (AMLE).

Memoli [45] proposed computational method for computing an AMLE directly between two surfaces, which is an extension of a computational method proposed by Oberman [46] to find the AMLE from a subset of \mathbb{R}^n to \mathbb{R} . However, the computation was formulated directly in the ambient Euclidean 3D space of the surface rather than in a 2D parameter space, so we did not pursue this method in this work, although future research could explore

¹The term “viscosity” is a historical name for the class of continuous but not necessarily differentiable functions which satisfy degenerate elliptic PDEs such as the infinity Laplace equation we are interested in ($\Delta_\infty u = 0$).

this direction for a comprehensive comparison of maps.

2.1.4 Conformal geometry

Recall from Figure 1.0.1 Ptolemy's map of the Western world as it was known to the ancient Greeks. Notice that in this conical projection of the globe, the parallel curves from the top to the bottom of the page are drawn as arcs while the meridians are straight lines that meet at a single point if extended. The property of this projection is that east-west distances are shown in correct proportionality to the north-south distances within a radius of the parallel line that crosses through the southern part of Greece on the map. As the map shifts further to the extremes in any of the four cardinal directions, the distortion of distances becomes more apparent, to the point where once the map extended beyond the equator, Ptolemy made the ad hoc choice to artificially resize the line below the equator to match the one directly above the equator. This is one of two map projections that was described in his *Geographia*, which attempted to address the problems that he had with world maps of his time that were using straight-line latitude and longitude lines, which considerably distorted lengths on top of not conveying that everything was part of a spherical world. As we understand now, this illustrates the concept that we cannot map a non-flat surface down to a flat space without inevitably introducing distortion to some extent and the constraint that we want to minimize distortion.

The same considerations must be taken into account for surfaces that have been triangulated: if the triangulation forms a planar graph, then it admits a straight-line embedding [16]; however, the embedding is not unique and geometric information can be lost when using an ad-hoc method. While we may not be able to preserve the metric structure while mapping to a flat space, conformal theory tells us that we can still preserve the *conformal structure* if we map to a space of constant curvature. Specifically, the *Riemann uniformization theorem* states that every simply connected Riemann surface can be conformally mapped to one of three canonical spaces of constant curvature: either the sphere, the Euclidean plane or the

hyperbolic plane. The map that does this is called the *Riemann map*. This is formalized by the notion of conformal equivalence as mentioned in the preliminary discussion on distortion. Two Riemannian metrics g and \tilde{g} on a smooth manifold M are said to be conformally equivalent if

$$\tilde{g} = e^{2u}g, \tag{2.1.11}$$

where u is a smooth function defined on M . From this, we gather two things. First, there is an equivalence class of a wide variety of surfaces with the same conformal structure that is to be had by varying the function u . Second, any representative in this equivalent class can be “uniformized” so that it has constant curvature i.e. embeddable in the sphere, plane, or hyperbolic plane.

Now while uniformization of a surface metric is unique and allows us to conformally map the surface into one of three two-dimensional parameter spaces of constant curvature, the particular map/parameterization is not unique. Depending on the space, it is unique up to an application of a particular class of operators, or self-mappings, called Möbius transformations. For the (complex) plane, the function $f : \mathbb{C} \rightarrow \mathbb{C}$ is a Möbius transformation if

$$f(z) = \frac{az + b}{cz + d}, \tag{2.1.12}$$

where a, b, c, d are complex numbers such that $ad - bc \neq 0$. By modifying the condition to be $ad - bc = 1$, we obtain a conformal map from the sphere onto itself. As for two-dimensional hyperbolic space, there are a number of equivalent models of this geometry, but the form of the Möbius transformation remains clear when we consider the upper half plane model, which is the upper half of the plane above the x-axis, and the metric on the space is given by

$$\frac{dx^2 + dy^2}{y^2}. \tag{2.1.13}$$

where $z = x + iy$, and x and y are the real and imaginary parts, respectively. Here, the equation remains the same as in Equation (2.1.12), except that now we have the restriction

that $a, b, c, d \in \mathbb{R}$ and $ad - bc > 0$ to achieve a smooth bijection into the upper half plane. It should be noted that Möbius transformations have the additional property of being isometric in that these transformations preserve hyperbolic distances.

An equivalent theory exists in the discrete case of triangulated surfaces and depends on a construction associated with triangulations called *circle packings*. For an introductory exposition on circle packings, see [58]. Thurston conjectured, and Rodin and Sullivan proved [49], that any triangulation of a sphere has an associated *circle packing*, a configuration of circles that are tangent to each other in the pattern of the triangulation, and that this circle packing is unique up to Möbius transformation. This discrete theory connects combinatorics to geometry and converges to the smooth theory with refinement. There are several related discrete formulations that were spawned from this central notion of circle packings, and we use the following definition of discrete conformal equivalence due to Springborn [9], which will be discussed in more detail later. Two Euclidean triangulations (T, ℓ) and $(T, \tilde{\ell})$ are discrete conformally equivalent if

$$\tilde{\ell}_{ij} = e^{\frac{1}{2}(u_i + u_j)} \ell_{ij} \tag{2.1.14}$$

for some scalar function $u : V \rightarrow \mathbb{R}$ defined on the vertices of T . Again, we can always tell from the combinatorics of a triangulation what type of surface it is, and this surface type determines to which surface of constant curvature the surface will be embedded by the Riemann map. In the latter case, such an assignment of u is called the *uniformization metric* and is unique up to scale for any particular surface.

2.1.5 Hyperbolic geometry

Hyperbolic geometry is a geometry that results by relaxing the parallel postulate of Euclidean geometry, which states that given a point and a line, there exists a unique line that passes through that point and is parallel to that line. The parallel postulate is actually independent

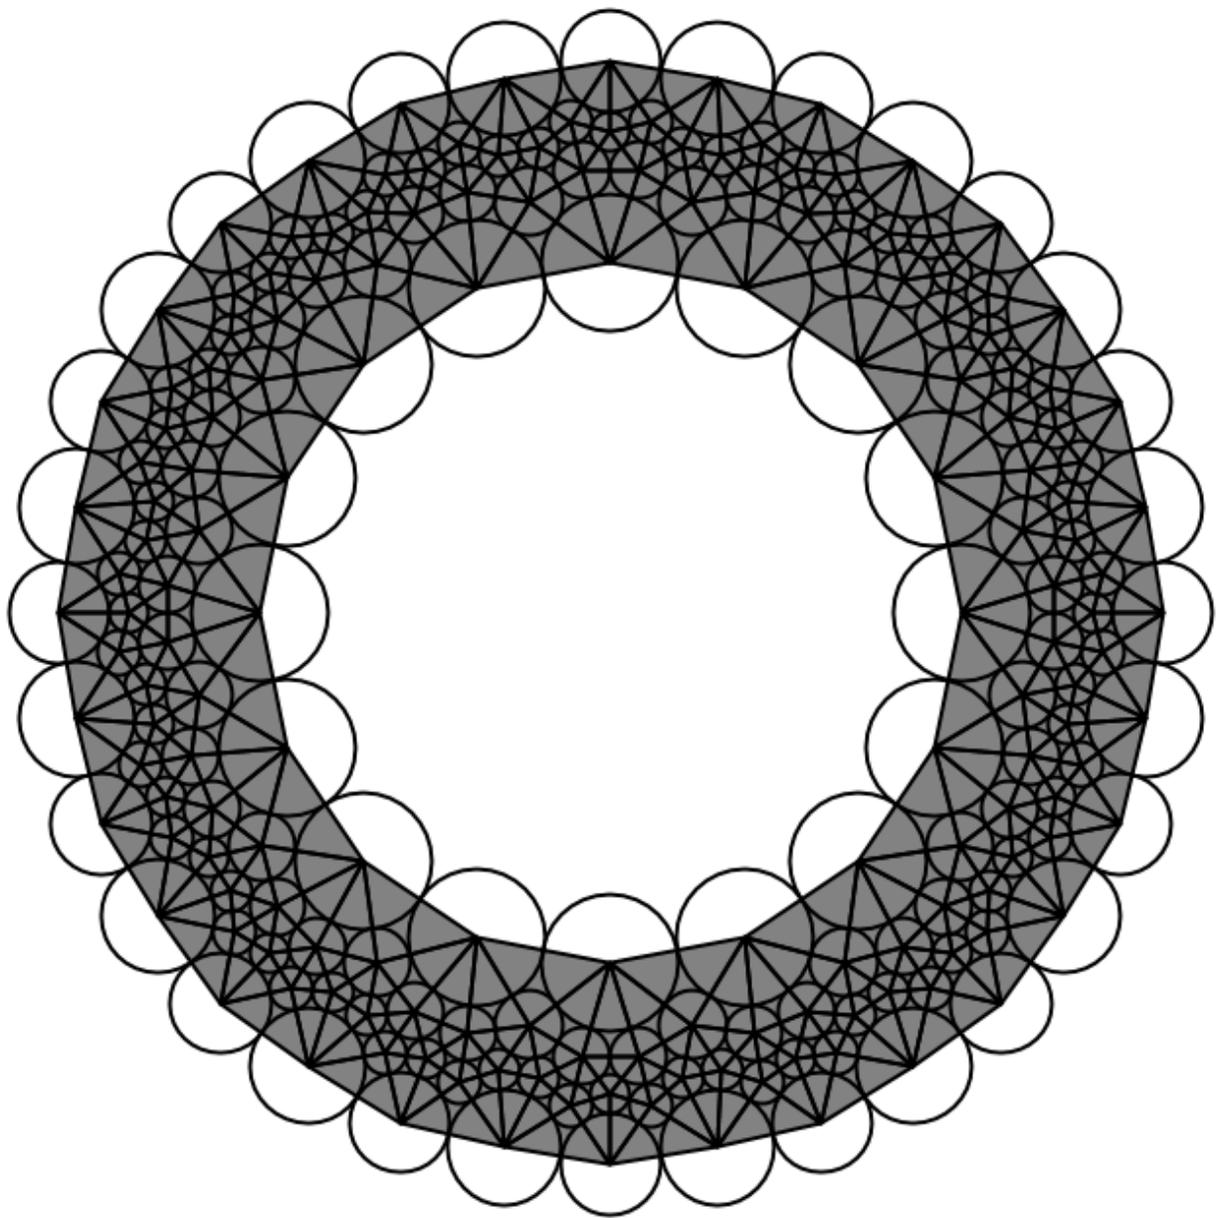


Figure 2.1.6: An example mesh parameterized in the Euclidean plane (in grey), together with an associated circle packing. Notice that the circles are associated with and centered on each vertex of the mesh, and circle tangencies encode adjacent vertices connected by edges.

of the other axioms of geometry, and it was shown in the 1800s that there are other models of geometry that are consistent without making use of the parallel postulate. There are a number of equivalent models of hyperbolic geometry, as alluded to in the preceding section. In particular, we will use the Poincaré disk model, which is the two-dimensional version of the sphere model that is described in Poincaré's words [59] as follows:

Suppose, for example, a world enclosed in a large sphere and subject to the following laws: the temperature is not uniform; it is greatest at the centre, and gradually decreases as we move towards the circumference of the sphere, where it is absolute zero. The law of this temperature is as follows: if R be the radius of the sphere, and r the distance from the centre, then the absolute temperature will be proportional to $R^2 - r^2$. Suppose that in this world the linear dilation of any body is proportional to its absolute temperature. A moving object will become smaller and smaller as it approaches the circumference of the sphere. Although from the point of view of our ordinary geometry this world is finite, to its inhabitants it will appear infinite. As they approach the surface of the sphere they will become colder, and at the same time smaller and smaller. The steps they take are also smaller and smaller, so that they can never reach the boundary of the sphere. If to us geometry is only the study of the laws according to which rigid solids move, to these imaginary beings it will be the study of the laws of motion of solids deformed by the differences in temperature alluded to.

In the disk model, the entirety of two-dimensional hyperbolic space is mapped to a Euclidean unit disk. The metric, given by

$$ds^2 = \frac{dzd\bar{z}}{(1 - \bar{z}z)^2} \tag{2.1.15}$$

is nonuniform throughout the interior of the disk, with lengths being compressed near the boundary of the circle. Geodesics in this space are given by either straight lines through the

origin or circular arcs that are perpendicular to the boundary circle.

The motivation for using hyperbolic geometry is the fact that it offers the possibility of a conformal global parameterization: any surface that admits a hyperbolic metric, regardless of its complexity, can be embedded in the Poincaré disk without angle distortion. Furthermore, the surface can be used to tile the disk. To understand this concept, consider the analogous example of a prototypical torus embedded in \mathbb{R}^3 . We can identify this surface with a square patch by cutting along a radial loop as well as the inner loop and identifying the edges. Now the patch can be embedded in \mathbb{R}^2 . No given parameterization is particularly special in the sense that an isometry applied in the plane, such as a translation or rotation, yields another valid parameterization of the torus on the plane. However, consider fixing one “copy” of the torus in the plane, and tiling the space by laying out additional copies adjacent to the fixed copy by corresponding edges, as in Figure 2.1.7. We can see that there is a subset of isometric transformations that maps from one tile to another – we can tile the entirety of the space by repeated application of these transformations to additional copies of the torus. The set of transformations that map from one tile to another is called *deck transformations*. Notice that the adjacent tiles form a lattice structure – this shows the finite subgroup structure of the set of deck transformations on the plane.

Analogously, the unique hyperbolic parameterization of a surface allows us to produce a hyperbolic tile that can be used to fill the entirety of hyperbolic space. Let us consider tiling hyperbolic space in the Poincaré disk model. The rigid motion in the space, known as a *gyrotranslation*, is determined by a few parameters: a point z_0 in the disk that is mapped to the origin and a rotation θ :

$$w = e^{i\theta} \frac{z - z_0}{1 - \bar{z}_0 z} \quad (2.1.16)$$

These gyrotranslations are equivalent to the Möbius transformations on the upper half plane model of hyperbolic space described in the preceding section. Now consider Figure 2.1.8, which exhibits a two-holed torus that has been uniformized and embedded in the Poincaré disk, producing a tile. Equivalently to the deck transformations of the flat torus in the

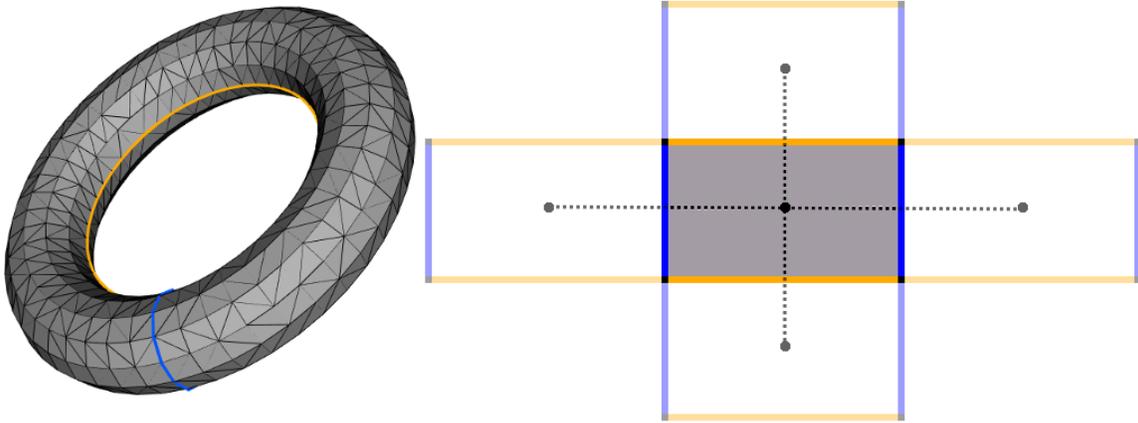


Figure 2.1.7: A torus (left) is mapped flat into the rectangular grey region (right). A copy of a flat torus in the Euclidean plane. The edges with matching colors are identified to show torus topology. The plane can be tiled by arranging multiple copies as shown. The set of transformations that generates the tiling is called deck transformations and forms a group whose relation forms a lattice as suggested by the dotted lines connecting adjacent tiles.

Euclidean plane, there is a finite subgroup of the conformal transformations in hyperbolic space called the Fuchsian group that takes a particular tile to a corresponding tile.

2.1.6 Orbifolds

Recall that manifolds can be thought of a topological structure in which every point can be associated with a patch of flat Euclidean space. An orbifold is a generalization of the manifold in which certain points are homeomorphic to the quotient of \mathbb{R}^2 by finite group actions. Intuitively, an orbifold can simply be visualized as a pillow-like surface, where most points on the surface can be identified with a flat patch of Euclidean space, and the corners of the pillow represent something like a the tip of a cone and hence are called *cone points*. The total angle around the cone point is less than 2π . We can visualize taking a circular patch of \mathbb{R}^2 with a wedge cut out of it and rolling it up as if to form a cone where the edges adjacent to the cone point are identified with each other. To illustrate the example, consider the 1-ring neighborhood of a vertex of an orbifold that has been parameterized in the hyperbolic plane in Figure 2.1.9. In particular, the point is an orbifold point of order 2,

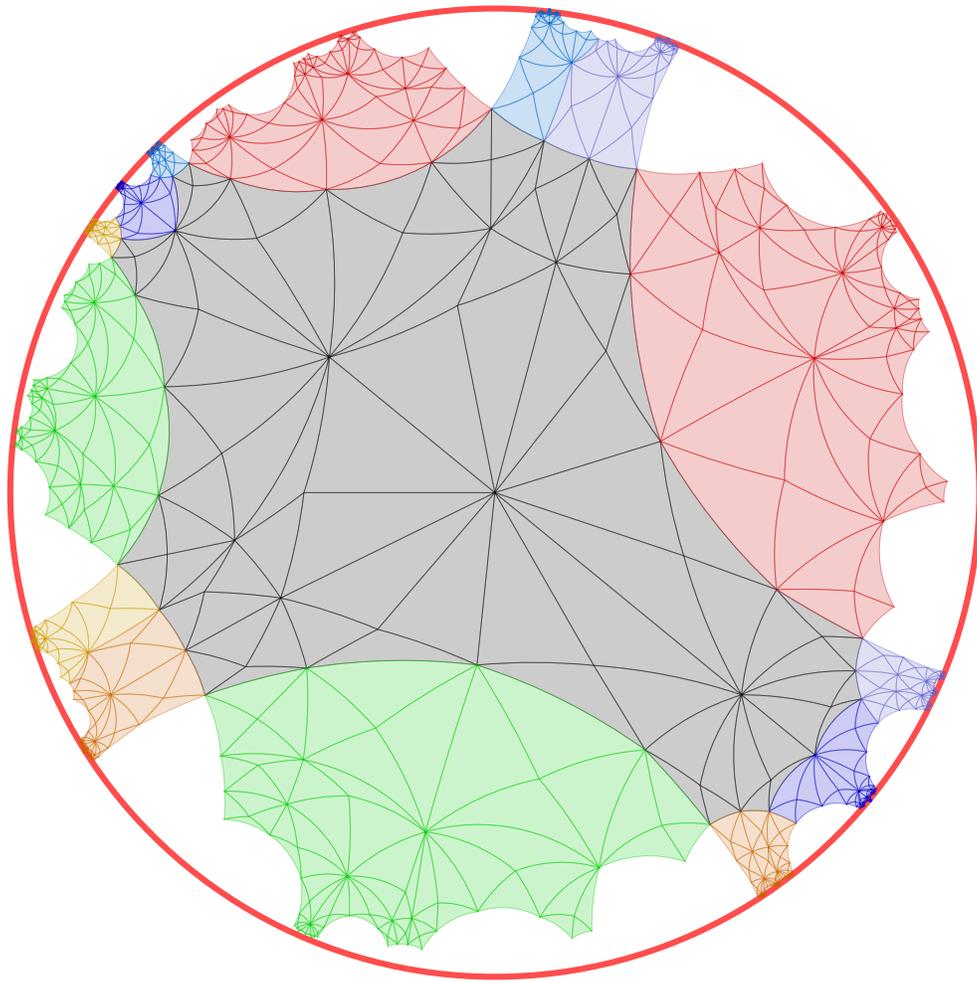


Figure 2.1.8: A partial tiling of the Poincaré disk.

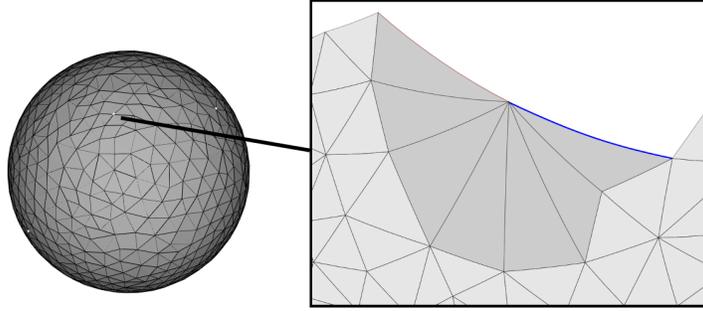


Figure 2.1.9: A sphere (left) may admit a hyperbolic metric given enough cone points (white vertices). Once it has been metrized, every region can be conformally mapped to the hyperbolic plane. Right: A cone point of order $k = 2$. The total angle around the cone point is $2\pi/k$. The blue and orange colored edges do not represent boundaries but rather are identified with each other.

which means that there is only a total angle of $2\pi/n = 2\pi/2 = \pi$.

A hyperbolic orbifold is a topological entity that can be associated with any arbitrary metric, and in particular, it can admit a hyperbolic metric given the right topology. In general, the number and order of the cone points may vary, which, together with the genus of the surface, alters the inherent topology of the surface as reflected by the Euler characteristic. A surface with a sufficient number of cone points admits a hyperbolic metric. In particular, a compact 2-dimensional orbifold has an Euler characteristic χ given by

$$\chi = \chi_0 - \sum_i \left(1 - \frac{1}{m_i}\right) \quad (2.1.17)$$

where χ_0 is the Euler characteristic of the underlying topological manifold and m_i are the orders of the cone points. For example, for a topological sphere (genus 0, $\chi_0 = 2$), we would require at least 5 cone points of order 2 in order to have a hyperbolic orbifold.

Orbifolds and Homotopy

In our method, a surface map between two manifolds is recast as a map between two orbifolds. The idea is to have a matching number of corresponding points specified so that a continuous map between the two orbifolds can be established such that the specified points are in

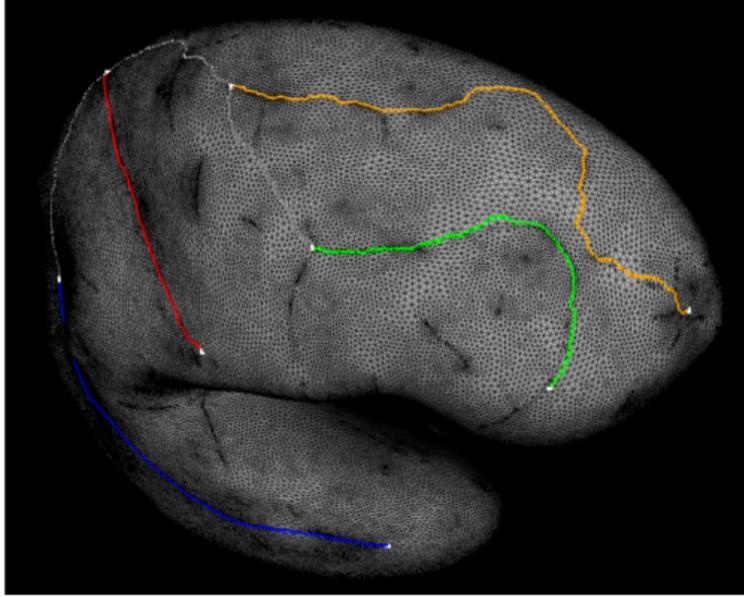


Figure 2.1.10: An orbifold with cone points shown as white vertices that are connected by a marking tree. The particular connectivity of the tree specifies the homotopy class of mappings to other orbifolds with equivalent marking trees. The marking tree may be varied to select a different homotopy class.

correspondence. In order to establish such a map, however, we will need to specify a *marking tree* on the orbifold, which is tree connecting the cone points – see Figure 2.1.10. With an equivalent marking specified on each orbifold, this fixes the *homotopy class* of maps for that particular pair of surfaces. This homotopy class contains the set of all maps that can be obtained from this initial map by fixing the cone points and continuously deforming the remainder of the map.

2.1.7 Harmonic Map – Existence and Uniqueness

Consider a pair of orbifolds that have been marked with an equivalent marking tree. As the marking tree fixes the homotopy class of maps between the orbifolds, we can consider an arbitrary representative map f_0 that continuously maps from one orbifold to the next. By continuous deformation of any part of the map, we can obtain any other map in the homotopy class. Hyperbolic geometry ensures theoretically that the resulting map achieves a unique global minimum. The Eells-Sampson theorem [15] states that there exists a harmonic map

f_h (i.e. a map that minimizes overall squared average length distortion) in the homotopy class of maps between any pair of hyperbolic surfaces, and Hartman [22] proved that such a map is unique. Thus, we can envision a process of relaxation that takes as input any map f_0 and produces by iterative deformation the unique harmonic map f_h .

2.2 Related work and applications

We describe here some of the work towards conformal planar parameterizations that serve as the basis for many state of the art surface mapping methods we utilize in order to establish correspondences between surfaces. The discussion of works specifically related to surface mapping methods described in this thesis will be left to their respective chapters.

Recall the discussion of circle packing associated with the triangulation of a Riemannian surface. Thurston’s proposal of a circle packing scheme with a particular set of tangencies, which allows for conformally parameterizing a surface mesh, sparked research in the area of conformal maps and serves as a theoretical lynchpin for the computational methods for conformal maps. Hurdal and Stephenson [25] introduced an alternative computational algorithm for computing a circle packing subject to a specified set of tangency constraints, which as a result is able to compute a circle packing as a discrete approximation to the Riemann map. While it is an improvement in the proposed scheme of Thurston, one drawback is that the method may not be able to solve for a circle packing for triangulations with very small triangles. In 1986, Hamilton [21] proposed Ricci flow on Riemann surface with metric g_{ij}

$$\frac{dg_{ij}}{dt} = -2Kg_{ij} \tag{2.2.1}$$

and proved that a solution converges to a constant curvature metric (i.e. the uniformization metric) as time goes to infinity. From this continuous formulation came analogous formulations in discrete settings: Luo [42] defines a convex energy to evolve the curvature of a piecewise flat metric to one of constant curvature. Bobenko et al. [9] derived a parallel no-

tion of discrete conformal equivalence based on hyperbolic polyhedra and introduced a convex energy that, if minimized, produces a uniformization metric. The two share in common that both utilize conformal factors to scale the metric to achieve the prescribed curvature at each vertex: [9] gives the energy explicitly, while [42] gives it as a path integral. Furthermore, both methods share the limitation that while the energy is convex, there are cases where an admissible conformal factor does not exist. The work in this thesis builds upon the discrete conformal method of [9], and computational details can be found in Section 4.3.2.

The discrete Ricci flow method [12] is a flow analogous to the continuous Ricci flow but defined on a triangulated surface. By encoding curvature in the circle radii at each vertex, the metric is computed in the space of circle packings. Note that the two theories of circle packing and discrete conformal equivalence are tied by a variational principle based on hyperbolic polyhedra, in which circle packing deals with hyperbolic polyhedra with prescribed dihedral angles, whereas discrete conformality deals with prescribed metric. The requirement of this method, as with the conformal factors-based method, is that the mesh is of sufficient quality, typically Delaunay meshes being assumed. This can be a challenge, especially when working with meshes acquired from upstream data sources that do not necessarily come with triangulations of uniformly good element density or shape quality. Literature suggests [57] edge flips may remedy triangles that are near singular, and recent work attempts to build edge flips into the discrete problem formulation [62] to address the existence problem, but we find the problem is not simply an issue of mesh element quality (Figure 4.5.1 will be discussed later). In practice, remeshing methods are relied on in combination with manual mesh manipulation to obtain admissible meshes. While remeshing is a research area of its own, there is publicly available software implementing techniques which we make use of and refer to in Appendix B.

We note that surface registration, in the cases in which we are interested, is just part of the larger problem of multiple registration, in which we need to align multiple datasets, bringing them into correspondence in a common reference system. As having this correspondence is

often a key component of further downstream task in an application, the distinction between surface alignment/registration and surface mapping to some degree speaks to the nature of the data and its intended application [64]. On one hand, there may not be a well-defined correspondence to be had. For example, incomplete or noisy range data can be captured from multiple vantage points, meaning that “corresponding” meshes may be marred by topological noise, and surface correspondence may not exist in the sense of being one-to-one. Perhaps it is desirable to use a partial mesh to query a shape database for intrinsically similar parts, in which case a topologically robust method [11] advantageous. Surface mapping assumes that the subjects of comparison are topologically matched and typically deal with intrinsic models, hence appealing to isometric and conformal approaches. This seems to be a reasonable assumption when considering the specimens in biomedical image analysis: putting aside the difficulty of ascertaining correspondences, biological models on the scale of organs and bone structures can be considered as non-rigidly deformable surfaces. A one-to-one correspondence can reasonably be established for the purpose of statistical analysis across time or population.

Chapter 3

Hyperbolic Orbifold Maps

3.1 Introduction

Cortical surface matching— establishing point-wise correspondences between cerebral cortex surfaces— is a crucial step in MRI-based studies of brain morphology. Algorithms typically aim to induce a surface-to-surface mapping that minimally distorts morphological features. It is also desirable to use information provided by experts to guide the mapping. This information can consist of landmarks, labeled as points or curves on the surfaces, that are required to correspond to each other. Our goal is a cortical surface matching method that exactly matches point landmarks while insuring that the mapping minimizes distortion.

We quantify surface distortion in terms of *conformality*, or angle preservation. Conformal maps have been studied intensely due to their ability to preserve key shape properties of biological specimens [25], and because conformal maps from any genus-zero surface to the sphere, and from any higher-genus surface to a surface of constant curvature, provably exist regardless of surface morphology [7] [65]. However, point landmarks are difficult to incorporate into conformal maps. A conformal map to the sphere, for example, is uniquely determined by the mapping of exactly three surface points; matching more than three requires sacrificing either conformality or exact landmark matching. For this reason, we seek maps that minimally

deviate from conformality, using the Dirichlet energy of the mapping to quantify this deviation. Specifically, we conformally map each cortical surface to a hyperbolic orbifold, and find Dirichlet energy minimizing maps between the orbifolds that exactly align an arbitrary number of point landmarks. We show that the Dirichlet energy has exactly one, unique, global minimum over the relevant set of orbifold-to-orbifold maps, making it computationally robust.

Our approach is summarized in Figure 3.3.1. Given two triangulated cortical surfaces, we first alter the distance metrics on the two surfaces using conformal transformations [57]. This results in hyperbolic metrics on both surfaces with singularities at a finite number of isolated cone points, near which the metric behaves as though the surface is shaped locally like the vertex of a cone. One such cone point is located at each of the point landmarks. We then calculate an initial mapping from one surface to the other that exactly aligns the corresponding point landmarks. There is exactly one deformation of the initial map that maintains these point landmark matches while obtaining a minimum of the Dirichlet energy among the set of maps reachable by continuous deformation, *i.e.* within the homotopy class of the initial map. Given the uni-minimal nature of the Dirichlet energy landscape, finding the energy minimizing mapping within the homotopy class is straightforward using gradient descent.

Using brains from a large epidemiological study [13] with manual point landmarks, we assessed whether exact landmark matching (versus approximate *e.g.*, [41]) results in greater surface distortion, and less utility in practical situations, that counterbalance the theoretical advantage of guaranteed globally-optimal mapping. We also assessed whether abandoning truly conformal mapping for Dirichlet energy minimizing mapping results in noteworthy practical disadvantages. To do this we compared our method (`OrbifoldExact`) to two competing methods that minimized landmark mismatch in a least-squares sense. An *orbifold least squares* method (`OrbifoldLS`), inspired by earlier work [41], balanced a tradeoff between Dirichlet energy minimization and landmark mismatch in a least squares sense. A *conformal least squares* method (`ConformalLS`) found the conformal map in the spherical

domain that minimized landmark mismatch [41]. We compared the methods in terms of point landmark mismatch, surface distortion, mismatch of novel (*i.e.*, not used to define the mapping) point landmarks, and ability to re-capitulate known population-level associations between cortical thickness and age [50]. Finally, we assessed whether the behavior of OrbifoldLS is stable with respect to critical but difficult-to-set operating parameters.

3.2 Prior work

Prior landmark-based cortical surface matching methods begin by finding harmonic energy minimizing mappings to spherical [41] or Euclidean [5] canonical domains as an initial step, or finding initial conformal maps to canonical Euclidean annuli [70] or hyperbolic “pairs of pants” [69]. One such canonical domain is then mapped onto the other in a way that encourages landmark matching. Using a Möbius transformation for this mapping [41] insures conformality but it is restricted to either exactly match only 3 points, or inexactly match a larger number. Harmonic maps are more flexible, but guarantee neither conformality nor exact landmark matching [5]. Quasi-conformal maps have bounded angle distortion, but the practical utility of implementations at the time of these experiments [71] was not clear. Methods that cut the brain surface and map the cut to the boundary of a canonical domain have the additional limitation that conformality is lost along the cut [70], and there is an arbitrary decision about how exactly to map out the cut to the boundary. Note that while we specify paths between point landmarks that are similar to cuts, these instead constitute a *marking*, *i.e.* a landmark ordering convention that insures that the eventual mapping comes from a natural homotopy class, *i.e.* that the mapping can be connected to the desired optimal mapping by some deformation.

Building on earlier work on mapping cortical surfaces to the hyperbolic disc for visualization [25], we leverage an earlier observation that there is a single globally optimal map between hyperbolic discs that minimizes the Dirichlet energy [35], and model the landmarks

as cone points in the hyperbolic metric to insure exact matching. Such cone point singularities have been considered previously to reduce area and length distortions during flattening to the plane, for applications such as texture mapping [66]. Other discrete conformal mapping methods [28] [57] incorporate cone point singularities into their conformal transformation of the surface distance metric without cutting.

3.3 Method

We begin with a pair of triangulated surfaces whose topology is spherical; in our experiments, each of these is the outer pial surface of a human cerebral cortex hemisphere output by commonly available software. Each surface has been annotated by an expert with a set of k point landmarks that are known to be in correspondence across surfaces. There are four key steps to our approach. The first, **orbifold mapping**, calculates conformally equivalent hyperbolic metrics on the surfaces, or equivalently, conformally maps the surface to a k pointed orbifold such that landmark points map to its cone points. Next, **marking** allows us to constrain the surface-to-surface mapping to a natural homotopy class of mappings: those that preclude reflections, complex surface folding, twisting, etc, in between exact point matches. For this step we define a tree that connects the landmark points analogously across surfaces. An **initial mapping** is constructed that maps the first tree to the second one, and extends that mapping continuously across the rest of the surfaces; this mapping belongs to the natural homotopy class. Finally, **energy minimization** is used to adjust this mapping so that it arrives at the unique Dirichlet energy minimizing map within this homotopy class that leaves the landmark matches fixed.

3.3.1 Orbifold Mapping

We use a conformal factors-based method [9] to calculate a conformal transformation of the surface distance metric. Briefly, a surface triangulation T is a set of vertices V , edges E , and

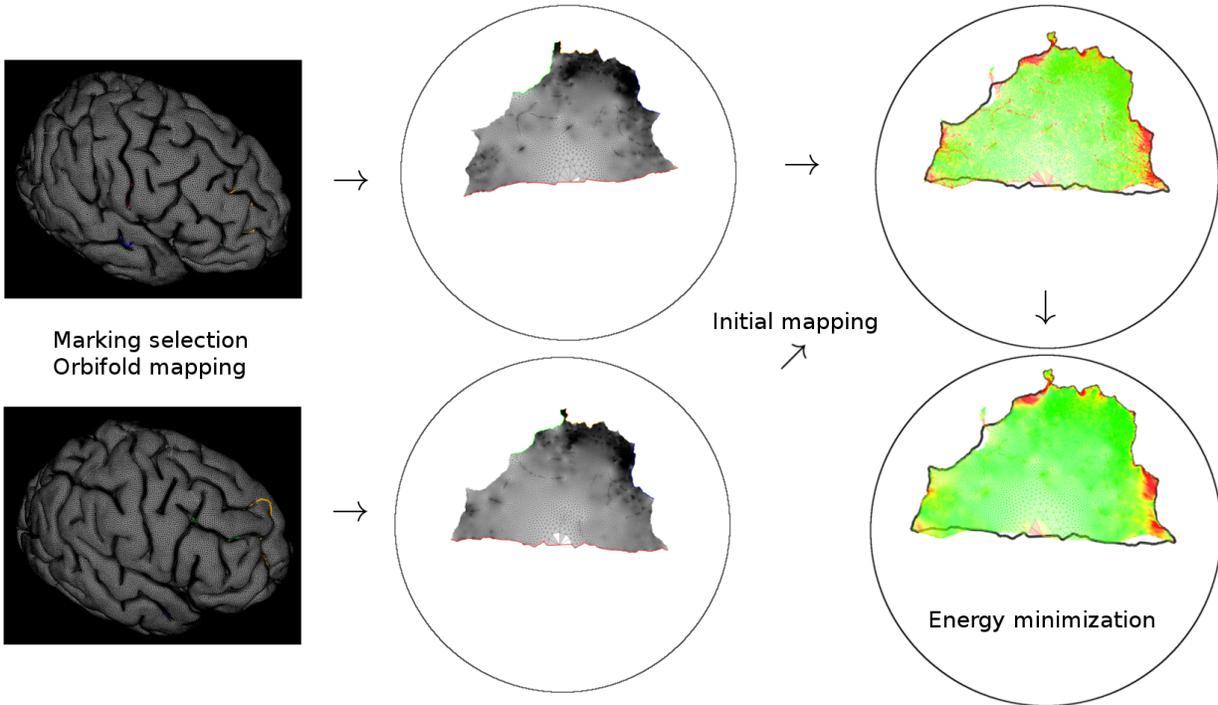


Figure 3.3.1: Overview of surface-to-surface mapping algorithm. Given two brain hemisphere surfaces (left), orbifold mapping constructs hyperbolic metrics on each surface. Then, marking selection identifies a natural homotopy class of mappings to optimize over. These steps allow us to map the surfaces to the hyperbolic plane (middle: sulci shown as colored boundary curves). Next, an initial surface-to-surface mapping is constructed in the hyperbolic plane (top right). This mapping introduces surface distortions in the form of dilatations (redder colors suggest greater distortion). Dirichlet energy minimization in hyperbolic space adjusts the mapping to obtain the global minimum of such distortions over the homotopy class (bottom right).

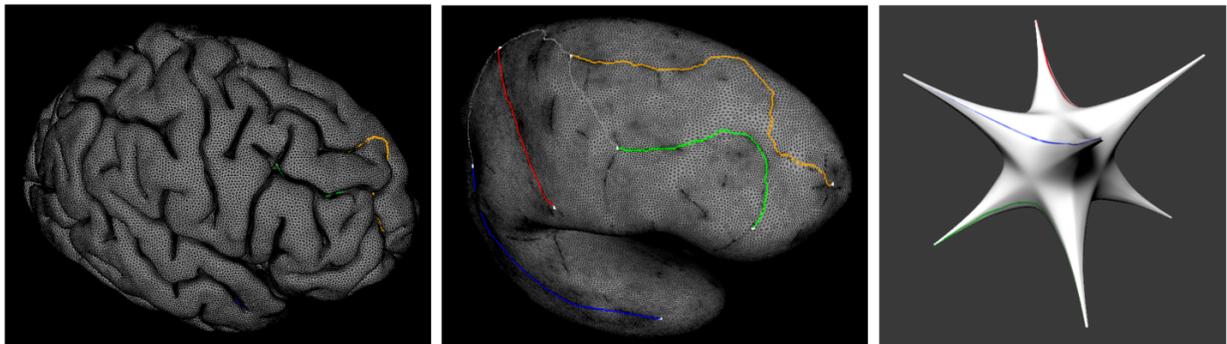


Figure 3.3.2: Identifying a cortical surface (left) with a hyperbolic orbifold (right). Landmark points (white dots) are selected and connected by a set of paths to form a marking tree (colored curves, shown on an inflated cortical mesh, center). The point-to-point distance metric on the mesh can be conformally transformed into a hyperbolic metric by imposing an angle constraint at each landmark point, effectively identifying the brain with the hyperbolic orbifold.

faces F , and a discrete metric $d : V \rightarrow \mathbb{R}$ assigns lengths to each edge such that the triangle inequality at each face is satisfied. Two combinatorially equivalent triangulations, one with a Euclidean metric d and the other with hyperbolic metric \tilde{d} , are discretely conformally equivalent if their respective metrics d and \tilde{d} are related by

$$\sinh \frac{\tilde{d}_{ij}}{2} = e^{\frac{1}{2}(u_i + u_j)} d_{ij} \quad (3.3.1)$$

where $u : V \rightarrow \mathbb{R}$ is an assignment of conformal factors to each vertex. This is an equivalence relation, and the set of discrete metrics in the same equivalence class is called a discrete conformal class. In this way, the conformal factors u define a conformal transformation of the Euclidean metric into a hyperbolic one.

Starting with a triangulation and the Euclidean metric (T, d) , we seek a discretely conformally equivalent triangulation with hyperbolic metric, such that the metric treats the k provided landmark points as cone points. One way to formulate this requirement is to consider the angle sum of vertex $v_i \in V$: this is the sum, over all mesh triangles that include v_i , of angles with v_i as the vertex. An angle sum of 2π means that the surface is locally flat at v_i , while angle sums less than 2π mean that the surface in the neighborhood of v_i more resembles the vertex of a cone. As described previously [9], we solve for conformal factors u such that the above equation is satisfied and each vertex is constrained to have a certain angle sum: for all point landmarks, the required sum is π to treat them as cone points; for all other points, the required sum is 2π to treat them as locally flat. For a genus zero surface to have a hyperbolic orbifold metric with cone point singularities, the sum of the angle defects ($2\pi - \text{angle sum}$) across all vertices must be greater than 4π . With an angle sum of π at the cone points and 2π everywhere else, this means we must have at least five cone points to insure that the resulting metric is hyperbolic. We use a trust region Newton's method [36] to minimize an energy function [9] to compute the conformal factors.

3.3.2 Marking selection

Once the hyperbolic metric is defined on both surfaces, we impose an ordering on the set of landmarks that constrains the surface-to-surface mapping to be simple and well-behaved, *i.e.* to map corresponding points onto each other exactly while inducing no gross foldings or twists to the rest of the mapping. To do so, we first draw vertex-constrained paths that connect one sulcal endpoint to the opposite endpoint. We then insert additional paths connecting endpoints across sulci until all endpoints are connected in a tree, called the *marking tree* (see Figure 3.3.2 middle). Note that paths must not intersect with themselves or other paths, but otherwise any tree of paths will suffice.

The purpose of the marking tree is to specify the homotopy class of the mapping *ie.* it allows us to specify an initial mapping such that one tree is mapped onto the other in a natural way, and the tree-to-tree mapping is extrapolated continuously over the rest of the surfaces. Optimization is then constrained such that the mapping remains within the same homotopy class. Thus, this step effectively allows us to rule out unnatural mappings. Together with the fact that under hyperbolic space, there exists a unique harmonic map that minimizes Dirichlet energy in each homotopy class, the mapping can be refined to obtain the globally unique map in the sense of minimum Dirichlet energy.

3.3.3 Initial Mapping

Because there is a single, unique global minimum for Dirichlet energy within each homotopy class, the initial mapping that is optimized to minimize Dirichlet energy is arbitrary– the only requirement is that it belong to the homotopy class of mappings that map one surface onto the other in a simple, reasonable way (*i.e.*, with no complex folding or twisting of the surface in between landmarks). We parameterize this initial mapping in the Euclidean disc to make use of existing computational methods [35] that insure the mapping is in the correct homotopy class, but we note that the optimization itself is governed by the hyperbolic metric induced upon the surface as described above. The initial mapping is constructed by assigning

each edge of the marking tree to a side of a regular polygon in the Euclidean disc and filling in the remainder of the mapping by minimizing harmonic energy. One such parameterization is performed for each surface; the overlay of the two Euclidean polygons provides the initial surface-to-surface mapping. The corners of the polygon are identified with the landmark points in the order that they appear in a traversal of the marking tree, and thus one marking tree maps to the other in a simple way while filling in the remainder of the mapping in a smooth, reasonable manner. We emphasize that while this approach effectively cuts the spherical-topology surface open along the marking tree, resulting in a topological disc in 3D that is then flattened into a disc contained in the plane as in prior work [35], this is solely for the purpose of establishing an approximate initial mapping that is then optimized based on the hyperbolic metrics described above; the optimized mapping does not contain discontinuities or other distorting artifacts along the marking tree edges. We also emphasize that this initial mapping is arbitrary: other methods may be applied without impact to the optimality of the final mapping.

3.3.4 Energy minimization

Suppose one surface contains edges e_{ij} that connect vertices v_i to v_j , w_{ij} are the cotangent weights $w_{ij} = 0.5(\cot \alpha + \cot \beta)$, where α and β are the two angles opposite the edge e_{ij} , and the other surface lies in the hyperbolic plane. The initial mapping f maps the points of this surface onto the other surface such that point landmarks are kept in correspondence. The Dirichlet energy of f under the hyperbolic metric can be approximated as follows:

$$E(f) = \frac{1}{2} \sum_{e_{ij}} w_{ij} \|f(v_j) - f(v_i)\|_{\mathbb{H}}^2 \quad (3.3.2)$$

Given hyperbolic orbifold structures defined on each surface and an initial map between them, the theorems of Eells and Sampson [15] and of Hartman [22] imply that there is a unique harmonic map that minimizes Dirichlet energy in the homotopy class of maps that

can be realized by deforming the original map.

To compute the Dirichlet energy minimizing map, we re-parameterize the the initial 2D polygon-to-polygon mapping in the Euclidean disc to the hyperbolic disc (specifically the Poincare disc). We then use steepest descent to minimize the Dirichlet energy: this amounts to adjusting surface vertex positions in the Poincare disc but constraining the vertices corresponding to landmark points to stay fixed and in correspondence. To overcome numerical issues, we follow a prior approach by [35] optimizing the mapping of each surface vertex one at a time: this point and its surrounding one-ring of mesh faces (its local "chart") is translated to the Poincare disc origin, where the hyperbolic metric is well approximated by a corresponding Euclidean metric. The position of that point is then adjusted to minimize the Dirichlet energy and translated back to its original position. See [63] for implementation of local charts.

3.4 Experiments

3.4.1 Data

We obtained brain MRI of 50 healthy elderly subjects from a prior study [13], identified gray matter voxels [17], used BrainVisa to convert each hemisphere’s cortical gray matter mask into matching inner and outer pial surface meshes [44], from which we removed small or slivery mesh triangles [19]. Cortical thickness was estimated at each outer pial surface vertex using a “normal-average” approach [30]. A set of 16 sulcal endpoints were annotated on each outer pial hemisphere by an expert rater using a validated protocol [56].

3.4.2 Competing methods

Experiments compared the method described above, termed `OrbifoldExact`, to two competing methods that strike a different balance between landmark matching and surface distortion. One competitor, `OrbifoldLS`, is identical to `OrbifoldExact` except that point land-

marks are not constrained to be fixed during energy minimization, and the energy function balances a tradeoff between landmark mismatch and Dirichlet energy:

$$E(f) = (1 - \lambda) \frac{1}{2} \sum_{e_{ij}} w_{ij} \|f(v_j) - f(v_i)\|_{\mathbb{H}}^2 + \lambda \sum_{v_i \in L_1} \|f(v_i) - L_2(v_i)\|_{\mathbb{H}}^2 \quad (3.4.1)$$

where L_1 is the set of landmark points on the source surface and $L_2(v_i)$ is the landmark point on the target surface corresponding to landmark point v_i on the source surface.

The other competitor, `ConformalLS`, first conformally maps each triangular mesh onto the unit sphere [57], then solves for a Möbius transformation (*i.e.*, a conformal mapping of the first sphere to the second one) that minimizes point landmark mismatch in the same least squares sense as in `OrbifoldLS`.

3.4.3 Performance measures

Surface distortion. `OrbifoldLS` and `OrbifoldExact` are able to induce surface distortions in the form of *dilatations*– stretches that transform local circles to local ellipses under the mapping– while `ConformalLS` precludes such dilatations by construction. For the former, we compute a discrete approximation of dilatation [51] at every mesh triangle and report summaries of dilatation over all vertices. We also show brain surfaces color-coded by dilatation under various mappings.

Landmark mismatch. `OrbifoldLS` and `ConformalLS` allow imperfect matching of point landmarks, while `OrbifoldExact` requires exact landmark matches by construction. We report the mean Euclidean distance between corresponding landmark points under the mapping for the former two approaches.

Strength of expected associations. We selected one of the 50 left outer pial surfaces as a canonical brain surface and used each of the three techniques to map the remaining 49 surfaces onto it. The mappings allowed us to transfer cortical thicknesses from the 49 surfaces onto the canonical one, and interpolate the thicknesses to the positions of canonical

mesh points. This resulted in 50 cortical thicknesses (one per subject) at each canonical mesh point. We calculated a linear regression model at each mesh point to assess the strength of association between that point’s local cortical thickness and the age of the corresponding subjects. The p values for these regressions were corrected for multiple comparisons [73], and the p values at the vertices were interpolated across intervening mesh faces. The surface area that had $p < .05$ was then calculated. Numerous studies (*e.g.*, [50]), agree that the thickness of the cortical mantle reduce with age, so we seek mapping methods that give rise to a statistically significant relationship with age across the largest possible cortical area. Results are summarized in Table 3.4 and Figure 3.5.2.

3.4.4 Experimental settings

Comparison of 3 methods. Given 16 point landmarks on a hemisphere surface, we consider two experimental settings. In the first, we use the full set of landmarks to define the mappings, and evaluate surface distortion, landmark mismatch, and strength of expected associations on appropriate methods. In the second, we cross-validate: we use 14 of the 16 landmarks to define the mapping and evaluate landmark mismatch for the remaining two.

Practical limitations of OrbifoldLS. OrbifoldLS includes two operating parameters that are difficult for a user to know how to set optimally: the marking tree and λ . We assessed whether settings of these parameters impact landmark mismatch and surface distortion by running OrbifoldLS over a range of settings for both and assessing variability in both performance characteristics.

3.5 Results

Landmark mismatch for OrbifoldLS and ConformalLS is in Table 3.1. Mismatch is substantial for each method, averaging greater than 5 *mm* for most sulci (note that mean mesh edge length is 0.55 *mm*). This motivates exact landmark matching as in our method: the

Sulcus	ConformalLS	OrbifoldLS, $\lambda = 0.1$	OrbifoldLS, $\lambda = 0.5$
Central	19.95	7.62	6.29
Precentral	20.01	7.80	6.33
Postcentral	21.42	9.56	7.40
Cingulate	26.61	7.16	6.11
Intraparietal	25.54	8.66	7.25
Superior Temporal	29.40	8.16	6.46
Superior Frontal	24.12	6.45	5.28
Inferior Frontal	24.31	5.90	4.90

Table 3.1: Average landmark mismatches, in *mm*, for sulcal endpoints across 50 subjects. Note that landmark mismatch for `OrbifoldExact` is zero by construction.

mapping problem is difficult enough that least squares methods do not readily find a solution that matches landmarks closely. Dilatation for `OrbifoldLS` and `OrbifoldExact` are in Table 3.2, and anecdotal dilatation maps are in Figure 3.5.1. As expected, `OrbifoldExact` gives rise to dramatic maximum dilatations, which occur at the isolated point landmarks where the orbifold construction has dramatically changed the surface distance metric. But large dilatations do not broadly affect large surface regions (see Figure 3.5.1 and mean dilatations in Table 3.2), suggesting that requiring exact landmark matches does not preclude broadly well-behaved mappings.

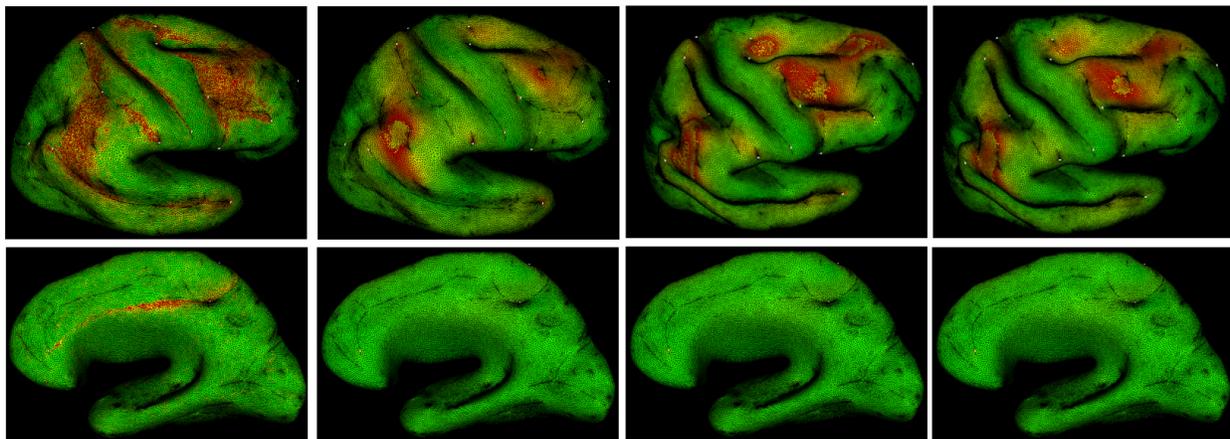


Figure 3.5.1: Mesh face dilatations for one example pair of surfaces (greener/redder indicates lesser/greater dilatation). Recall that dilatation is exactly one for `ConformalLS`. Left to right, columns show dilatation of the initial mapping, `OrbifoldExact`, `OrbifoldLS` with $\lambda = 0.5$, and `OrbifoldLS` with $\lambda = 0.1$. White dots indicate landmark points.

Table 3.3 shows landmark mismatch for the left-out landmark in cross-validation. `OrbifoldLS`

Method	Mean	Median	Max	Method	Age
OrbifoldExact	1.489	1.300	2089	ConformalLS	450.0
OrbifoldLS, $\lambda = 0.1$	1.452	1.283	320	OrbifoldExact	3885.8
OrbifoldLS, $\lambda = 0.5$	1.503	1.288	732	OrbifoldLS, $\lambda = 0.1$	4020.8

Table 3.2: **Left:** Mean, median, and maximum distortion from conformality (dilatation) across all mesh faces on all 50 surfaces, for three competing methods. Note that dilatation of 1 is realized in an isometry and by construction is 1 everywhere for **ConformalLS**. **Right:** Total area (mm^2) of cortex showing statistically significant evidence of a linear relationship with cortical thickness.

conferred no notable benefit over **OrbifoldExact**, again suggesting that requiring exact landmark matches draws no appreciable cost in terms of practical performance. **ConformalLS** provides superior matching of the left-out point for approximately half of the landmarks, but matching is similar to **OrbifoldExact** or worse for the other half. This suggests that **ConformalLS** offers no obvious performance advantage over **OrbifoldLS** in terms of matching previously-unseen landmarks, and the landmarks used to define the mapping again exhibit substantial mismatch.

Table 3.2 summarizes the cortical surface area significantly associated with age. **OrbifoldExact** and **OrbifoldLS** are highly similar in this regard, while **ConformalLS** lags far behind. This suggests that **OrbifoldLS** and **ConformalLS** holds no obvious practical advantage in terms of relevant applications that might compensate for their other theoretical or computational limitations.

Sulcus	ConformalLS	OrbifoldLS, $\lambda = 0.1$	OrbifoldLS, $\lambda = 0.5$	OrbifoldExact
Central	12.49 (20.90)	25.59 (7.22)	25.61 (6.53)	25.18
Cingulate	28.96 (21.23)	26.65 (6.93)	26.65 (6.27)	25.20
Inferior Frontal	24.31 (21.11)	24.57 (6.87)	24.56 (6.03)	22.82
Intraparietal	30.44 (21.21)	31.37 (7.08)	31.35 (6.18)	30.39
Postcentral	21.62 (20.28)	28.92 (7.01)	28.92 (6.20)	30.93
Precentral	17.04 (21.05)	25.84 (6.83)	25.85 (6.01)	26.12
Superior Frontal	23.86 (20.77)	36.75 (5.79)	35.89 (4.95)	35.90
Superior Temporal	39.57 (23.52)	33.84 (7.69)	33.85 (6.81)	31.16

Table 3.3: Mean sulcal endpoint deviations in mm of the left-out sulcus using a map created by applying the given method. Numbers in parentheses indicate the landmark deviation (mm) averaged over the remaining seven sulci.

Variability in landmark mismatch for **OrbifoldLS** is shown in Table 3.5. Landmark

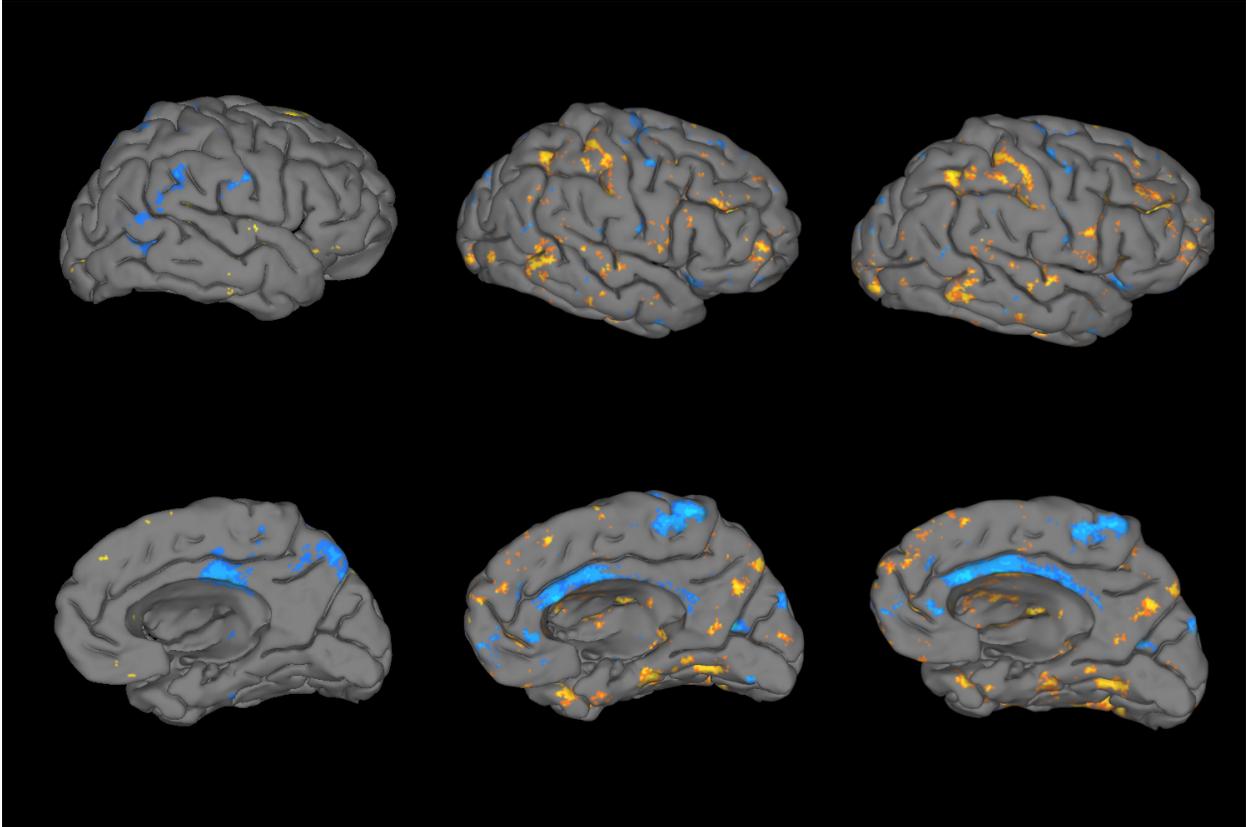


Figure 3.5.2: Strength of association results can be improved by `OrbifoldExact`. Map of areas on the cortex that show a statistically significant linear relationship between age and cortical thickness ($p \leq 0.05$). From left to right, the columns result from maps generated using 1. spherical map, 2. optimal Dirichlet energy map, and 3. combined energy map with $\lambda = 0.1$.

Method	Age	Gender
<code>ConformalLS</code>	450.0	800.0
<code>OrbifoldExact</code>	3885.8	1824.0
<code>OrbifoldLS</code> , $\lambda = 0.1$	4020.8	1991.0

Table 3.4: Strength of association results can be improved by `OrbifoldExact`. Total area (mm^2) of cortex showing statistically significant evidence of a linear relationship with cortical thickness. matching varies with respect to marking tree and λ (i.e. weighting parameter to prioritize landmark matching over harmonic energy), although users cannot know *a priori* what setting is favorable. So while matching of `OrbifoldLS` and `OrbifoldExact` are similar, variability due to arbitrary parameter settings reduces the robustness of `OrbifoldLS`. Indeed, this demonstrates the difficulty in least squares approaches as optimizing an energy function with local minima is subject to initial conditions.

Marking tree	Deviation (mm)	λ	Deviation (mm)
A	2.05	0.1	2.05
B	0.23	0.5	0.24
C	0.53	0.8	0.027

Table 3.5: Mean landmark deviations with respect to marking tree and λ . Experiment was conducted with eight point landmarks. We set $\lambda = 0.1$ when varying the marking tree, and we fix marking tree A when varying λ . Tree A is the same as shown in Figure 3.3.2. Trees B and C connect the landmarks instead in a single path.

In conclusion, our method provides a surface-to-surface mapping that exactly matches point landmarks and arrives at the global minimum of a particular surface distortion energy. Experiments suggest that neither requiring exact landmark matches, nor failing to require conformality, reduce the practical performance of the method, suggesting usefulness in practice.

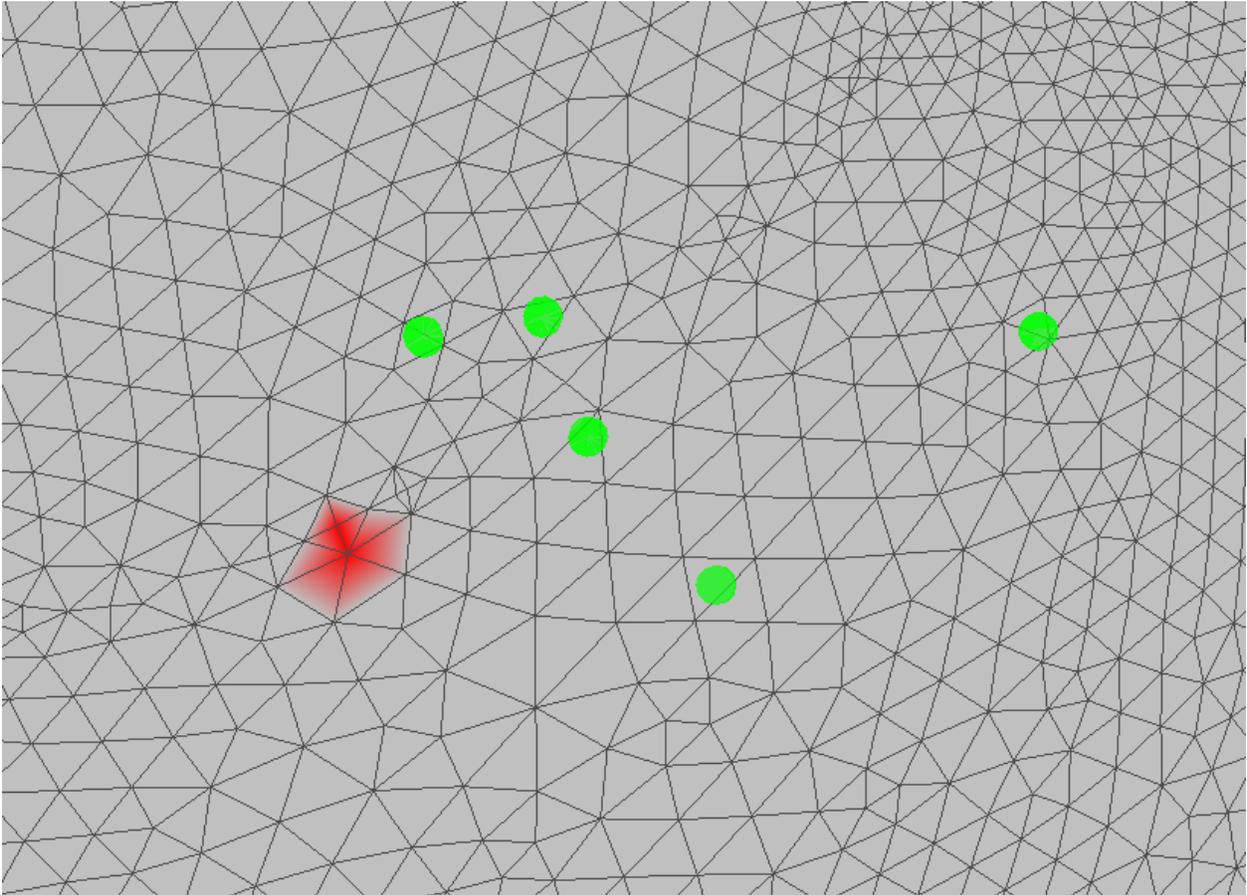


Figure 3.5.3: Whereas there exists a unique map from optimizing the Dirichlet energy, the map that results from optimizing the combined energy is dependent on initial conditions. Here is the result of fixing $\lambda = 0.1$ (λ is a weighting parameter determining how heavily landmark matching should be favored rather than harmonic energy) but varying the marking tree. Green dots are images under various mappings of one point from the source surface, which corresponds to the point on the target surface colored in red.

Chapter 4

Curve Extension For Hyperbolic Orbifold Maps

4.1 Introduction

Geometric analysis of biological surfaces, such as those extracted from medical imaging modalities or 3D surface scanning techniques, is a prevalent and important application. Such analysis can help differentiate between groups of surfaces [76], correlate shape properties with external variables of interest [69], and track shape change over time [74]. Solving these problems can shed light on the biology of organism development, changes due to environmental and genetic variables, and degeneration due to disease.

One predominant approach to surface analysis begins by establishing a *map* between the surfaces. Explicitly, a map is a correspondence between points on one surface and points on the other. It is imperative that such maps reflect biological realities and identify homologous biological features to each other. For example, in most brain MRI studies, each brain contains the same set of broad neuroanatomical regions and it is essential that correspondences properly identify those regions to each other across individuals [10] [47]. A mapping induces a deformation of one shape as its points are placed into correspondence with

those on the other shape. Quantifying the degree of distortion induced by a mapping is one way of quantifying shape-to-shape differences, which is useful for a variety of applications. For example, in protein folding [31], the shape of a probe protein can be compared to those of a library of known ligands to identify protein-ligand docking candidates; and quantified shape differences between monkey skull fossils [72] can elucidate and/or corroborate the interspecies relationships that have been empirically modeled by scientists.

Mapping-induced surface distortion can be quantified in terms of changes to geometric features such as angles and edge lengths on a triangular mesh, or equivalently in terms of metric distortion: changes to point-to-point distances on the surface under a given distance metric. A central goal of principled mapping techniques is the minimization of such metric distortion, which when considered over the entire surface quantifies a map’s distance from isometry, or metric-preserving map. While there are other ways to quantify and optimize a map’s distance from isometry [11] [23] [34] [61], we take the approach of harmonic energy, a classic approach [48] that can equivalently be seen as the average of edge length distortion over the surface. Metric distortion can be quantified via properties other than length; for instance, conformal maps are maps that do not distort angle. Each metric has its relative merits that depend on the context of the application. In this paper, we consider the harmonic energy of the map as a distortion metric for optimization and shape comparison.

We take as input two topologically equivalent triangulated surfaces, each provided with landmark curves consisting of non-closed paths of adjacent vertices (that is, the ends of the paths do not connect to form a loop). We propose a novel construction that allows both surfaces to be conformally mapped to canonical *hyperbolic orbifold* domains. Intuitively, orbifolds are surfaces that resemble a pillow in shape, that is, with the exception of some singular points, the neighborhood of a surface point resembles two-dimensional Euclidean space. Theoretically, there exists a unique map between such hyperbolic surfaces that exactly maps curve endpoints to corresponding curve endpoints, maps the interior points of each curve to the corresponding interior of the curve, and minimizes distortion of triangle edge

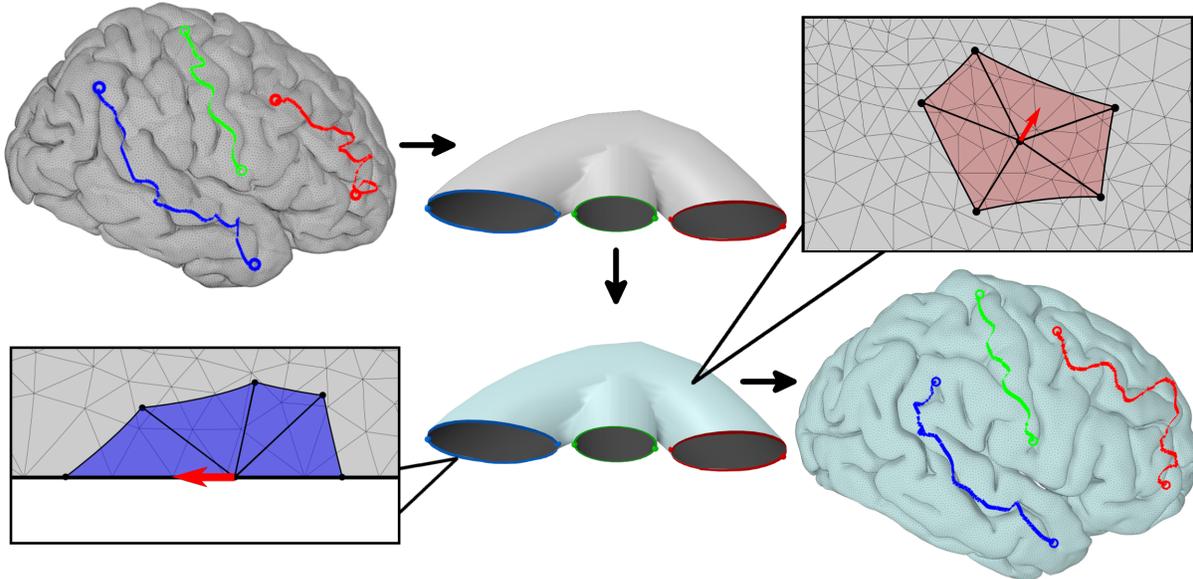


Figure 4.1.1: Brains have anatomical landmarks which must be mapped in correspondence (the frontal, central, and temporal sulci are major brain landmarks shown in red, green, and blue, respectively). We conformally map them to orbifolds (i.e. surfaces with cone singularities at landmark endpoints). The surface with holes above show half of the doubled surface that forms this orbifold – by symmetry, we need only retain this portion for the map. An initial map between orbifolds keeps curve landmarks in correspondence (round points indicate curve landmark endpoints that are mapped exactly), and minimizing Dirichlet energy obtains a minimum distorting map that maintains landmark correspondence.

lengths. We implement an algorithm that approximates this map. Figure 4.1.1 gives an overview of the approach, which improves upon prior work by optimizing over the entire surface simultaneously, including the boundaries.

The Riemann uniformization theorem states that every simply connected Riemann surface can be conformally mapped to either the sphere, the Euclidean plane or the hyperbolic plane. In particular, the Poincare disk model of the hyperbolic plane provides a convenient frame of reference for computing harmonic distortion. The approach taken in this paper utilizes a conformal map to a canonical hyperbolic domain, which facilitates our process for minimizing distortion.

In this setting of minimum-distortion mappings, it is desirable to be able to specify the mapping of certain regions. It is often possible to annotate anatomical surfaces with anatomical landmarks, either through automated procedures or manual tracing by an expert.

The goal is a map that identifies corresponding landmarks to each other while minimally distorting a surface as it is being mapped. These landmark constraints can serve as either a soft constraint in which landmarks are encouraged to align, or they may be a hard constraint which must be satisfied in the construction of the map. Some approaches employ the strategy of a regularization, heuristically balancing between goodness-of-fit with the landmarks and overall distortion [41]. In this work, we utilize landmarks as hard constraints on the mapping problem, as our framework can guarantee that a global optimum exists that also satisfies the constraints exactly.

4.2 Prior work

We focus on landmark-based methods, in which an explicit but sparse set of corresponding points, curves, or patches specified across the surfaces constrain the final correspondence. These landmarks are placed manually, based on domain knowledge, or automatically [27]; landmark locations are either taken at face value or location uncertainty can be estimated [8]. In contrast to mapping techniques that lift the surfaces into a high-dimensional embedding space [55] [38] [68] [20], we utilize a canonical domain as a means for establishing a correspondence. Several approaches utilize the plane [53] and round sphere [61] [34] as an intermediate domain for surface mapping, but we restrict our attention to the hyperbolic domain, which has seen relatively less exploration in terms of application.

Among the methods that utilize the hyperbolic domain, there are a few variants that have been applied to surface mapping. Shi et al. [54] map brains with curve landmarks by cutting the mesh along the landmarks and treating them as boundaries of a corresponding auxiliary mesh. Harmonic relaxation is then performed to optimize the map between these auxiliary meshes; however, the map is simply fixed along the boundary, and it loses out on an opportunity to optimize for distortion in those regions. Tsui et al. [67] parameterized brains with a set of point landmarks to a canonical hyperbolic orbifold domain. Again, harmonic

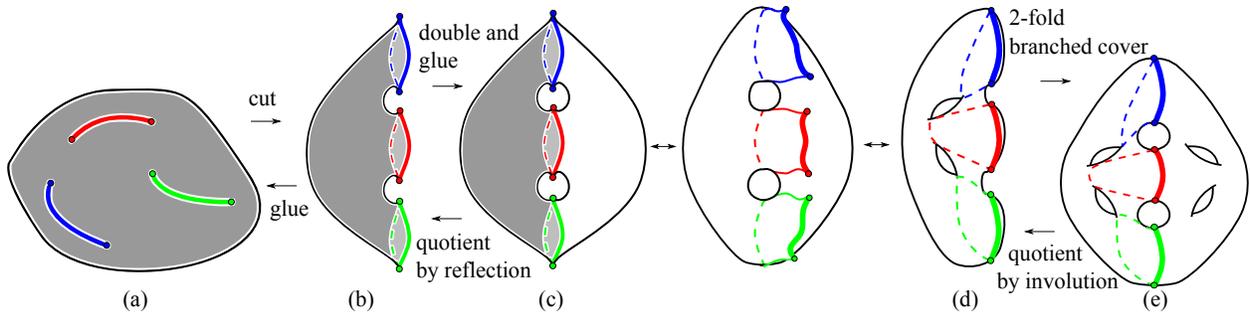


Figure 4.2.1: A closed surface with curve landmarks (top left) corresponds to an orbifold with index-4 cone points. The mesh is cut along each curve landmark, producing a surface with holes (b). A copy of the surface with holes with reverse orientation and reflectional symmetry across the curves is glued to the first surface with holes to produce a doubled surface. The curve endpoints form index-2 cone points on the doubled surface, giving a new orbifold (c). Topologically, the orbifolds from (c) to (d) are equivalent. The Eells-Sampson theorem applies to a 2-fold branched cover of the orbifold branched over its cone points (e), which has rotational symmetry about the landmarks and is itself a manifold surface, in which the cone points are replaced by points with a rotational symmetry.

relaxation is performed to optimize an initial orbifold map, for which theory guarantees there exists a unique solution. Recently, there has been work by Aigerman and Lipman [2] to compute orbifold embeddings in the Euclidean domain, which can be applied to construct surface maps that align cone points.

Using hyperbolic orbifold domains is important because any surface of complex topology (with genus greater than one) or more than four landmark points maps to a hyperbolic rather than spherical or planar orbifold. Current orbifold surface mapping approaches have been limited to point landmarks (as in Tsui et al) or have unnaturally restricted the correspondence along curve landmarks (as in Shi et al). To the best of our knowledge, this paper is the first to apply orbifolds to achieve exact matching of curve landmarks, while allowing the correspondence along the curves to be optimized.

4.3 Method

4.3.1 Theory and Curve Landmarks Extension

Our approach consists of mapping two surfaces to a canonical *orbifold* domain with singularities, and then optimizing a surface map between the surfaces to minimize length distortion.

We begin with the theoretical motivation for the general idea of using orbifolds as a canonical domain, and then we will proceed to discuss our extension of the idea to handle curve landmarks. Hyperbolic geometry ensures theoretically that the resulting map achieves the global minimum of the length distortion we have formulated. A hyperbolic orbifold is a surface in which every point has a neighborhood that is the quotient of hyperbolic space \mathbb{H}^2 by a finite metric preserving group action; it can intuitively be viewed as a pillow-like surface with cone points at the corners of the pillow. This orbifold can then be uniquely identified with a particular hyperbolic surface—its two-fold branched cover—along with an involution on this two-fold branched cover. The Eells-Sampson theorem [15] states that there exists a harmonic map (i.e. a map that minimizes overall squared average length distortion) between any pair of hyperbolic surfaces, and Hartman [22] proved that such a map is unique. Further, the harmonic map exactly maps point landmarks to point landmarks, since uniqueness implies that it commutes with the involution.

While orbifolds allow for point landmark matching (by identifying landmarks with cone points), we propose the following novel extension of the orbifold construction to enable surface mapping with curve landmarks (see Figure 4.2.1). A surface with genus g and marked with a set of k curve landmarks, whose endpoints do not coincide, can be identified with a genus g surface with k boundaries by cutting along the curves. If we take two copies of the genus g surface with k holes, reverse the orientation of one copy (that is, reflect it across the cuts), and glue them along the boundaries, we obtain a closed, genus $2g + k - 1$ surface with reflectional symmetry across the boundary curves – we call this surface the **doubled surface**.

Note that there exists a map from the doubled surface to the original surface which is two-to-one everywhere except at the point landmarks, where it is one-to-one. Because the neighborhood of any point in the interior of a curve landmark is now symmetrical, minimizing distortion by optimizing the harmonic energy can only cause a displacement of the point along the curve. This lets the optimization preserve the curve landmarks while still optimizing the mapping along the curves.

While cutting the curve landmarks creates a duplicate of every vertex along the curve interior, the original surface can be recovered by gluing the pairs back together. We preserve the correspondence between these pairs during the optimization by maintaining equivalence across the cut. We still need to identify the landmark curve endpoints on the two surfaces. By making them cone points, we make the doubled surface into an orbifold, and given a map between two such orbifolds, the Eells-Sampson theorem states there is a unique harmonic map in the homotopy class established by that map. Our method approximates this harmonic map on the doubled surface and projects it to the original surface, yielding a surface map where the curve endpoints are placed in exact correspondence, and curve interior points correspond to curve interior points. We describe in the next subsection how to construct this approximated harmonic map.

4.3.2 Algorithm

We begin with a pair of triangulated surfaces without boundaries and with the same topology (see, *e.g.*, Figure 4.1.1). Each surface is annotated with a set of k curve landmarks that anatomically correspond across surfaces. As an initial step, we generate spheres with holes and doubled surfaces as described in Section 4.3.1. Three key steps follow in our approach. The first, **orbifold mapping**, calculates conformally equivalent hyperbolic metrics on the doubled surfaces, effectively mapping the original surface to a $2k$ pointed orbifold such that curve endpoints map to orbifold cone points and landmark curves map to geodesics connecting said cone points. Next, **marking and initial mapping** is the process of computing

an arbitrary continuous, bijective initial map from one surface with holes to the other such that curve endpoints map to curve endpoints and curve interior points map to curve interior points. Finally, **energy minimization** is used to adjust this initial mapping so that it arrives at the unique harmonic energy minimizing map within this homotopy class that leaves the landmark matches fixed. The harmonic energy, defined below, quantifies edge length distortion induced by the mapping.

Orbifold mapping

We use the notion of *discrete conformal equivalence* as in Bobenko et al. [9]: two combinatorially equivalent triangulations, one with a Euclidean metric ℓ and the other with hyperbolic metric $\tilde{\ell}$, are discretely conformally equivalent if the respective edge lengths ℓ_{ij} and $\tilde{\ell}_{ij}$ between points i and j are related by

$$\sinh \frac{\tilde{\ell}_{ij}}{2} = e^{\frac{1}{2}(u_i+u_j)} \ell_{ij} \quad (4.3.1)$$

where $u : V \rightarrow \mathbb{R}$ assigns conformal factors to each vertex. Bobenko et al. [9] propose the problem of solving for a set of conformal factors u satisfying the above equation, subject to the constraint that each vertex satisfies a particular total angle condition θ . We compute a discrete conformal map of the doubled surface S_i with k landmark curves to the $2k$ -pointed orbifold $O_i, i = 1, 2$ by solving an instance of this discrete conformal mapping problem, with the constraints that the angle sum θ is 2π at every vertex except at the $2k$ curve endpoint vertices that are mapped to index-4 orbifold points, at which θ is set to $\pi/2$. The orbifold that we consider is the surface formed by taking two copies of the mesh, cutting along the landmark curves, and gluing along the holes, so effectively, this imposes the constraint that the cone points have angle π constraint. Note O_i is embeddable in the hyperbolic plane \mathbb{H}^2 since we have computed a hyperbolic metric. The conformal factors that satisfy the orbifold

angle constraints can be found by minimizing the nonlinear energy function

$$E_{T,\Theta,\lambda}^h(u) = \sum_{ijk \in T} 2\hat{V}_h(\lambda_{ij}, \lambda_{jk}, \lambda_{ki}, -u_i, -u_j, -u_k) + \sum_{i \in V} \theta_i u_i \quad (4.3.2)$$

where

$$\lambda_{ij} = 2 \log \sinh \frac{\ell_{ij}}{2} \quad (4.3.3)$$

is the logarithmic length of the edge ij , the auxiliary function $2\hat{V}_h$ is given by

$$2\hat{V}_h(\lambda_{12}, \lambda_{23}, \lambda_{31}, \lambda_1, \lambda_2, \lambda_3) = \quad (4.3.4)$$

$$\alpha_1 \lambda_1 + \alpha_2 \lambda_2 + \alpha_3 \lambda_3 + \alpha_{12} \lambda_{12} + \alpha_{23} \lambda_{23} + \alpha_{31} \lambda_{31} \quad (4.3.5)$$

$$+ \mathbb{J}(\alpha_1) + \mathbb{J}(\alpha_2) + \mathbb{J}(\alpha_3) + \mathbb{J}(\alpha_{12}) + \mathbb{J}(\alpha_{23}) + \mathbb{J}(\alpha_{31}) \quad (4.3.6)$$

$$+ \mathbb{J}\left(\frac{1}{2}(\pi - \alpha_1 - \alpha_2 - \alpha_3)\right), \quad (4.3.7)$$

where we consider a hyperbolic triangle ijk with logarithmic lengths λ_{ij} scaled by the conformal factors λ_i and λ_j , and α_k is the hyperbolic angle opposite of side ij . Supposing the lengths are given by a, b, c in counterclockwise order, where a corresponds to edge ij , then the angle α_k is given by

$$\alpha_k = \tan^{-1} \frac{\sinh((a-b+c)/2) \sinh((a+b-c)/2)}{\sinh((-a+b+c)/2) \sinh((a+b+c)/2)} \quad (4.3.8)$$

Here, \mathbb{J} denotes Milnor's Lobachevsky function:

$$\mathbb{J}(x) = - \int_0^x \log |2 \sin t| dt. \quad (4.3.9)$$

Finally, the following quantities are defined

$$\alpha_{23} = \frac{1}{2}(\pi + \alpha_1 - \alpha_2 - \alpha_3), \quad (4.3.10)$$

$$\alpha_{31} = \frac{1}{2}(\pi - \alpha_1 + \alpha_2 - \alpha_3), \quad (4.3.11)$$

$$\alpha_{12} = \frac{1}{2}(\pi - \alpha_1 - \alpha_2 + \alpha_3). \quad (4.3.12)$$

The derivative of $E_{T,\theta,\lambda}^h$ is determined by the partial derivative of u_i :

$$\frac{\partial}{\partial u_i} E_{T,\theta,\lambda}^h = \Theta_i - \sum_{jk:i,j,k \in T} \tilde{\alpha}_{jk}^i \quad (4.3.13)$$

where $\tilde{\alpha}_{jk}^i$ is the angle opposite of side jk in hyperbolic triangle ijk (use Equation (4.3.8) with lengths given by $\tilde{\ell}$), and Θ_i is the sum of angles around vertex i . Also, the second derivative at u is given by

$$\sum_{i,j \in V} \frac{\partial^2 E_{T,\theta,\lambda}^h}{\partial u_i \partial u_j} du_i du_j = \quad (4.3.14)$$

$$\sum_{ij \in E} w_{ij}(u) \left((du_i - du_j)^2 + \tanh^2\left(\frac{\tilde{\ell}_{ij}}{2}\right) (du_i + du_j)^2 \right) \quad (4.3.15)$$

We use the first and second derivatives to perform nonlinear optimization of the energy function using the TRON Newton trust region solver [36]. According to [9], if the method converges, then the conformal factors gives the unique orbifold metric for the mesh that satisfies the angle constraints that we specified for each vertex as mentioned above. The sufficient conditions for a solution to exist are not discussed by Bobenko et al. [9], but we found that sufficiently dense and well-conditioned triangle meshes can in practice be metrized, occasionally with the aid of preprocessing (discussed in Section 4.5).

Marking and initial mapping

Similar to prior work [67] [3], we define a *marking*, which is a tree that connects the landmark curves analogously across surfaces in order to constrain the surface map to a natural homotopy class of maps: those that preclude reflections, complex surface folding, twisting, etc., in between exact landmark matches. An initial map is constructed that maps the first tree to the second one, and extends that map continuously across the rest of the surface; this map belongs to the natural homotopy class and preserves landmark correspondence. The curve landmarks should consist of edges of the triangulation. The curve landmarks must be disjoint, that is, they should all have distinct endpoints and they should not cross. In addition, the curve landmarks should use only one edge of each triangle. This is because we will reflect and glue the triangles along the cuts, and including two consecutive edges from a triangle would produce a degenerate situation with the triangle and its reflection adjacent along two edges. We ensure this by replacing any two consecutive edges from one triangle with the unchosen third edge of that triangle.

For the datasets that we experimented with, we automatically join the curve landmarks up into a tree by computing geodesic vertex paths that connect a particular set of endpoints such that the selected geodesic paths neither intersect existing curve landmarks nor each other. Cutting the mesh along the tree yields a polygon whose number of sides depends on the number of landmarks. We parameterize the topological polygon to a regular polygon with the Euclidean harmonic map. We generate such a parameterization for both the source and target surface in corresponding fashion, making sure to consistently orient the polygon, and the initial map is determined by overlaying the polygons. Note that while there will likely be high distortion in this initial setting, the curve landmarks will be mapped onto corresponding curve landmarks, and their endpoints will be in exact correspondence.

Energy minimization

Next, we adjust this initial map so that it minimizes its harmonic energy under the hyperbolic metric: distortion in the lengths assigned to the edges under the hyperbolic metric estimated at the time of orbifold mapping. Our approach is based on the observation that the hyperbolic harmonic map can be viewed as a collection of locally Euclidean harmonic maps. If the map is represented as a function $f : O_1 \rightarrow O_2$, the harmonic energy of the map can be well approximated in a local neighborhood by:

$$E(f) = \frac{1}{2} \sum_{e_{ij}} w_{ij} (f(v_j) - f(v_i))^2, \quad (4.3.16)$$

where v_i and v_j are surface vertices, α, β are angles opposite the edge e_{ij} connecting v_i to v_j , and $w_{ij} = 1/2(\cot \alpha + \cot \beta)$ are called Laplacian cotangent weights. Thus, we adjust the vertex positions one at a time in a gradient descent approach; in particular, for each vertex v_i , we parameterize the image of its one-ring under the mapping f in the Poincare disk, making sure v_i is centered at the origin using the hyperbolic Möbius transformation

$$w = \frac{z - z_0}{1 - \bar{z}_0}. \quad (4.3.17)$$

Where e_{ij} is a boundary edge, the cotangent term belonging to the missing triangle is set to zero. For each vertex, we calculate the derivative of $E(f)$ in the neighborhood of that vertex and move the vertex in the direction of steepest descent. However, when the update is applied to curve landmark internal vertices, we project the update vector onto the curve tangent vector at that point, to ensure that the curve point remains on the curve. Additionally, we note that landmark curve interior points occur in pairs and require that they map to the same position on the symmetric boundary to preserve continuity. We update the pair of points simultaneously by averaging each individual tangential contribution to the gradient of $E(f)$ and apply a matching update step to both. When we achieve the energy minimum,

we project the resulting map onto the original surfaces by re-gluing the boundaries, and the resulting map maintains the landmark correspondence established upon initialization in addition to optimizing for edge length distortion.

4.4 Results

4.4.1 Datasets

We applied our curve constrained orbifold mapping approach to three biological surface data sets. In each case we started with triangle meshes derived from MRI or laser scanning devices, together with expert tracings of curve landmarks. We remeshed [19] the meshes to improve triangle quality. We then mapped pairs of surfaces to each other and evaluated the amount of surface distortion induced by the maps, the generalization of maps to yet-unseen curve landmarks, and the ability of map distortion measures to capture biologically relevant differences. Finally, we demonstrate the utility of this approach for cases of anatomical labeling uncertainty.

Framingham Heart Study Dataset We obtained brain MRI of 100 healthy elderly subjects (58 females; average age 62) from a prior study [13], identified gray matter voxels [17], and used BrainVisa to convert each hemisphere’s cortical gray matter mask into matching inner and outer pial surface meshes [44]. An expert rater annotated 8 sulci on each mesh, which is a subset of a validated protocol [56] including the following sulci: central, precentral, postcentral, cingulate, intraparietal, superior temporal, superior frontal, and inferior frontal.

Old World Monkey Skull Dataset We use 21 meshes constructed from laser scans of Old World monkey skulls [72]. The skulls are genus-4 closed surfaces and have been marked with 8 curve landmarks, including each of the large canines, each eye socket rim, each zygomatic arch, the neck rim, and the center line down the vault in the back. The set contains 7 skulls from the macaque species, 4 from the baboon species, and the remainder are close relatives in the evolutionary tree.

Face Recognition Grand Challenge (FRGC) Dataset We reconstructed 910 spherical face meshes using the ball pivot algorithm [6] from a database of range images [18] and mark each with four curve landmarks (the lower eyelids, the center line down the nose, and the top of the lower lip). 419 are female examples, and expressions include happy, shocked, angry, sad, and neutral.

4.4.2 Comparison methods

We compare our method, `OrbifoldCurves`, with related approaches based on conformal maps:

1. **HyperbolicPants:** To evaluate the method’s ability to reduce surface distortion through relaxation of curve interior point mappings, we compare to the `HyperbolicPants` method. This method also maps each surface to a canonical domain through hyperbolic metric estimation, with exact matching of curve landmarks [54]. However, this method applies an initial, fixed mapping of curve landmarks to boundary curves in the canonical domain, while our approach maintains fixed curve endpoints but allows curve interiors to slide to reduce surface distortion.
2. **ConformalLS:** To evaluate the enhancement in landmark matching accuracy provided by moving from Euclidean to hyperbolic conformal mapping, we compare to `ConformalLS`, a spherical conformal method described in that aligns landmarks in a least-squares sense [67]. This method computes an optimal spherical conformal map by composing a conformal map to the round sphere with a Möbius transformation that best minimizes point landmark mismatch in a least-squares sense. So as to emulate curve matching for a more apt comparison to `OrbifoldCurves`, we sampled each curve landmark uniformly by arc length to arrive at 10 point landmarks per curve that drove the Möbius transformation.
3. **OrbifoldPoints:** To evaluate the value added by exact matching of entire curves over

solely curve endpoints, we compare with `OrbifoldPoints` [67]. As described above, this method uses the same hyperbolic metric estimation and relaxation approach as `OrbifoldCurves` to map surfaces to an orbifold domain in which the curve endpoints are mapped to cone points. However, in this approach curve interior are not used in any way to compute the map.

4.4.3 Energy minimum robustness to differences in initial conditions

Our first experiment examined the ability of the algorithm to reach the same global minimum regardless of initial conditions. In the brain data set, we randomly selected a pair of brains and mapped them to each other using `OrbifoldCurves` multiple times, each with a different marking tree used to generate initial maps. In each run, we observe a monotonic decrease in the Dirichlet energy with respect to iteration, which is expected from the theory. The resulting maps appear similar to each other on visual inspection, and by computing the variance of the coordinates of each vertex resulting from each map, we see that the maps are overall stable with respect to different initializations. This suggests that our implementation is able to closely approximate the unique harmonic hyperbolic map and behaves as expected from theory. Figure 4.4.2 gives an example of stable solution being reached from varied initialization.

4.4.4 Comparing distortion and landmark matching across methods

For each method and each data set, we generate maps from a prototype surface to all other surfaces in the dataset using the full set of curves. We then assess how much surface distortion each method produces in the maps (see Table 4.1). We use *dilatation* as a discrete estimate of local angle distortion. As in [51], we consider the large and small singular values of the Jacobian of the restriction of the map from original to hyperbolic surface on a particular

triangle t . Supposing the side lengths are given as a, b, c , then the maximum and minimum singular values are given, respectively, as

$$\Gamma = \sqrt{\frac{1}{2} \left((a+c) + \sqrt{(a-c)^2 + 4b^2} \right)}, \quad (4.4.1)$$

$$\gamma = \sqrt{\frac{1}{2} \left((a+c) - \sqrt{(a-c)^2 + 4b^2} \right)}. \quad (4.4.2)$$

Then dilatation is given by the ratio Γ/γ . A value of 1 represents an absence of angle distortion, as is the case in a conformal mapping. We approximate this in hyperbolic space by placing one of the vertices of the image triangle on the origin in the Poincare disk and computing using Euclidean distance. *Elastic energy*, defined as

$$L(f) = \sum_{e_{ij} \in E} \left(\frac{\|f(v_j) - f(v_i)\|}{\|v_j - v_i\|} - 1 \right)^2, \quad (4.4.3)$$

estimates length distortion, with 0 representing the minimum, as in an isometric mapping [33]. For `OrbifoldPoints` and `ConformalLS`, which do not require exact matching of the entirety of corresponding curves, we quantify curve-to-curve mismatch by average closest-point distance. Closest-point distance from a point and a polyline is given by the minimum distance between that point and any point defining the polyline, and so the closest-point distance between two polylines is the pointwise closest-point distance averaged over all points.

First, as expected the proposed `OrbifoldCurves` achieves lower elastic energy than `HyperbolicPants` due to its added ability to reduce elastic energy by sliding interior curve points. Second, `ConformalLS` and `OrbifoldPoints` do not happen to accidentally achieve near-zero landmark matching error; thus the requirement of exact landmark matching is relevant to real-world performance. Finally, the requirement of exact curve matching (`OrbifoldCurves` and `HyperbolicPants`) induces greater length distortion as expected, but the increase in length distortion is not exorbitant.

As a followup experiment, we consider the relative distortion of the mapping when re-

stricted to a region of interest surrounding the curve landmarks themselves between the `OrbifoldCurves` method and the `HyperbolicPants` method. For each brain map in the Framingham dataset, we compute the harmonic energy of a local region consisting of all connected vertices that are at most 10 edges away from a landmark vertex (approximately 5mm in geodesic distance). On average, the distortion in `OrbifoldCurves` tends to be less than that of `HyperbolicPants`. The result is summarized in Figures 4.4.4. In a majority of maps, the harmonic energy of `OrbifoldCurves` is lower than `HyperbolicPants` in the neighborhood of the landmarks. Furthermore, quality measures not specifically optimized also show improvement: all `OrbifoldCurves` maps show lower elastic energy and generally lower dilatation in the neighborhood of the landmarks (see Figure 4.4.5). This suggests that while the maps are overall very similar, relaxing the vertices along the landmark curves allows the mapping method to further optimize for distortion near the landmarks.

Method	Dilatation	Elastic Energy	Landmark mismatch
Brains			
<code>OrbifoldCurves</code>	3.286	77620	0
<code>HyperbolicPants</code>	3.485	80774	0
<code>OrbifoldPoints</code>	2.896	53062	1.960
<code>ConformalLS</code>	1	45764	6.823
Skulls			
<code>OrbifoldCurves</code>	1.349	16508	0
<code>HyperbolicPants</code>	1.352	16539	0
<code>OrbifoldPoints</code>	1.446	16048	0.012
<code>ConformalLS</code>	N/A	N/A	N/A
Faces			
<code>OrbifoldCurves</code>	1.114	4527	0
<code>HyperbolicPants</code>	1.115	4665	0
<code>OrbifoldPoints</code>	1.105	4509	0.429
<code>ConformalLS</code>	1	3534	2.400

Table 4.1: Overall average map distortion metrics for maps generated from mapping a single probe surface to all other surfaces in the dataset.

4.4.5 Generalization of maps to unseen landmark mismatch across methods

We evaluated whether the introduction of landmark sliding sacrificed the generalizability of the maps, i.e. its ability to match yet-unseen anatomical landmarks by performing the following leave-one-out cross-validation approach.

1. Leave out one landmark l_i .
2. For each method k and each pair of surfaces S_a, S_b , compute map $f_k : S_a \rightarrow S_b$ using all landmarks l_j where $i \neq j$.
3. Compute the closest-point distance $d(f_k, l_i)$ by computing the distance from the image $f_k(v)$ of each point v in l_i from the source surface S_a and the closest point on the corresponding landmark on the target surface S_b .
4. Repeat previous steps for each other landmark.

On the resulting map, we compute the average squared closest-point distance between the left-out landmark from the first mesh and the left-out landmark mapped from the second mesh onto the first mesh. We compare the averages across each method and for each dataset, except for the face dataset, which due to its size, we randomly selected a subset of 100 face meshes on which we perform the experiment. This leave-one-out error is highly similar across all hyperbolic methods and higher in **ConformalLS** across all data sets. See Figure 4.4.3 for brain results, which showed the most variation between the four comparison methods. This suggests that the generalizability of the map is preserved even after introducing interior curve point sliding.

4.4.6 Biological relevance of map distortion measures

To understand how useful our method might be in biomedical applications, we calculated surface maps and used distortion as an index of shape difference. We then investigated

whether this shape difference measure captured relevant biological differences between groups of surfaces. Each data set can be divided into multiple groups: young (less than 62 years old) vs. old brains; male vs. female brains; macaque vs baboon monkey skulls; faces with neutral vs. non-neutral expressions (happy, sad, angry and shocked); female vs. male faces.

Initial probe analysis

We selected a probe surface from one group and mapped it to all other surfaces in the data set to determine the degree to which within-group shape differences were lesser than between-group shape differences. Table 4.2 summarizes the results. In all but one case the within-group mappings have lesser length distortion than the between-group mappings on average (the male vs. female face maps are the exception). The results are intuitive: neutral face to shocked face mappings have the highest distortion because shocked faces involve the greatest deformation of both the eyes and mouth (Figure 4.4.6, left), and mappings from macaques to baboons capture the more elongated brain case and snout of the latter (Figure 4.4.6, right). These results suggest that the method is capable of identifying biologically-driven shape differences between surfaces.

Followup pairwise analysis

As a followup, we examined the harmonic map distortions of all pairwise mappings within each dataset to examine whether the observed trends extend to the whole dataset, and whether there are overarching effects within and between groups of surfaces. The harmonic energies of the matrix of pairwise mappings are visually summarized in Figure 4.4.6.

We fit a linear mixed model to explain the harmonic energy of the mapping with the class of the mapping. The class of the mapping is defined by the mutually shared property of the two mapped surfaces, or is defined to be an in-between mapping for unmatched surface pairs. For example, a mapping from a face with a happy expression to another face with a happy expression is put in the ‘happy’ group, whereas if the target surface has a non-happy

expression, it is put in the ‘between’ group. We control for individual subject variability in the specification of the model. We test whether the average energy of the between-group mappings differ from the within-group mappings, and we also investigate the relative difference of the individual groups of mappings with respect to each other.

When comparing groups of expressions of faces, we find that the harmonic energy of mappings between different expressions averages to 1.767, and the within-group mappings are significantly lower than the harmonic energy of between-group mappings (all p-values at most $p = 0.05$). From lowest to highest mean harmonic energy relative to in-between mappings, the groups are ordered as follows: none (-0.0334), sad (-0.030), angry (-0.023), happy (-0.012), shocked (-0.008). The happy and shocked groups show high dissimilarity between all other groups of expressions, but they also show significant within-group dissimilarity. This is attributed to the variance in these groups of expressions. The angry, neutral, and sad groups show relatively more within-group similarity, as well as similarity amongst themselves. Indeed, the difference from a neutral facial expression and either a sad or angry one can be subtle, as is the case in this dataset. Lastly, they exhibit high dissimilarity with the happy or shocked groups, which is intuitive and can be seen in Figure 4.4.6 where two clusters seem to form in the face matrix.

When comparing monkey skull mappings grouped by species, both macaques and baboons exhibits within-group similarity. The average between-group harmonic energy is 5.581. We find that the harmonic energy of within-group mappings is significantly lower than the harmonic energy of between-group mappings (-0.132 for macaques, $p = 0.0173$; -0.266 for baboons, $p = 0.000387$). Furthermore, we find that the mean harmonic energy of within-baboon mappings is 0.134 lower than those of macaques ($p = 0.0573$).

When comparing groups of brains partitioned into old and young as described above, we find that the average harmonic energy of maps between young and old brains is 13.602. When the maps are grouped by age, the average harmonic energy of within-old mappings is significantly less than between-young-and-old mappings (-0.5901, $p=0.0136$) but only numerically

less for within-young mappings (-0.3836, $p=0.1022$). There was no statistically significant difference between the harmonic energy of within-young and within-old maps ($p = 0.447$). Grouping by gender and comparing between- and within-group mappings, we find that the between-group average energy is 13.6293, and the within-group energy is significantly lower than the harmonic energy of between-gender mappings – 0.5916 less for within-female ($p = 0.0099$) and 0.4406 less for within-male ($p = 0.0777$). The difference between the harmonic energy of within-male and within-female maps were not significant ($p = 0.5823$).

Group	Elastic Energy	Dilatation
Brains		
Female*	75630	3.009
Male	80472	4.095
Young*	74843	2.651
Old	80658	4.323
Faces		
Neutral*	3196	1.101
Sad	3404	1.103
Angry	3929	1.112
Happy	5631	1.150
Shocked	9857	1.168
Male*	4268	1.116
Female	3826	1.113
Skulls		
Macaque*	9193	1.236
Baboon	17430	1.331

Table 4.2: Average groupwise distortions for maps generated from mapping a probe surface to all other surfaces in the dataset. Figure 4.4.1 shows a snapshot of the probe surfaces for each dataset.

4.4.7 Mapping in the presence of anatomical uncertainty

To assess the relative performance of least squares landmark matching methods with exact landmark matching methods, we compare surface maps computed by `ConformalLS` and `OrbifoldCurves` methods. For a randomized subset of the Framingham brain dataset where each example mesh was duplicated twice, a group of three expert tracers trained on a common sulcal tracing protocol [56] traced eight major sulcal landmarks. We choose a particular set

of landmarks on the target surface, then we compute the surface map for all sets of landmarks on the source surface. Figure 4.4.8 shows the landmarks mapped onto the target surface for the least-squares and exact method, respectively. We visually compare maps with the exact- and least-squares matching methods to assess differences in correspondence quality in the face of user variability across the approaches. We find that the least-squares method tends to oversmooth, with landmarks often coming out of the sulcal pits. In contrast, the exact method fits contours and complicated features more closely, and as such, the variability is more representative of the tracer variability rather than an artifact of the matching method (for example, compare the superior temporal sulcus, blue curves, in each method). To quantify the controlled variability that we observe in the figure, we define hyperbolic closest-point distance (HCP distance) from a particular point p to a reference set of points Q is the shortest distance $d(p, q)$ to any $q \in Q$. Then the HCP distance between point sets P and Q is the average HCP distance of all $p \in P$ to Q . we compute the average HCP distance between each forward-mapped sulcus and the tracing on the target surface used to fix the mapping. Overall, the HCP distance for `ConformalLS` is higher than our `OrbifoldCurves` method. The HCP distances are broken down by landmark in Figure 4.4.7. This is consistent with our observation that there is less extraneous variability and uncertainty with an exact method.

4.5 Discussion

Overall, the maps that we construct with the proposed method are an improvement over existing conformal map-based approaches in the sense that curve landmarks are matched exactly, while length and angle distortion is comparable. By the dual surface construction, our method computes a unique optimal map without fixing a particular boundary condition (i.e. fixing the mapping of the curve landmarks). As evidenced by the favorable distortion metrics of `OrbifoldCurves` over `HyperbolicPants`, the impact of boundary relaxation is

less distortion in the vicinity of the curve landmarks. We note that the harmonic energy is an approximation, computed in the hyperbolic domain by applying equations for the Euclidean harmonic energy. As a consequence, while our `OrbifoldCurves` method performs favorably to `HyperbolicPants` in terms of harmonic energy in a majority of maps, there are some maps where the energy is apparently higher due to the harmonic energy of the more distorted `HyperbolicPants` maps being underestimated. Indeed, the `OrbifoldCurves` maps with slightly higher harmonic energy do perform better when considering either elastic energy or dilatation as mapping quality.

Theoretically, the harmonic energy of a unique mapping can be used to capture the difference between surfaces, and to a certain extent, this is demonstrated in our experiment, most notably in the skull dataset. The elongated braincases of baboon skulls give an increase in harmonic energy when mapping to macaques. This is in contrast with the within-group maps, which showed some more similarity indicated by lower harmonic energy. Differences were less pronounced in the face dataset, as the differences between certain expressions such as neutral, sad, and angry are more subtle. Only larger contortions brought on by smiles and shocked expressions seemed to show as harmonic energy differences between groups. Finally, there did not appear to be substantial clustering with the brains. Typically, incorporation of other biomarkers, such as those derived from 3D MRI or genetics can be more indicative of conditions. The use of harmonic energy of the mappings for shape analysis may have its limitation in the scale of differences it is able to detect, but the mappings themselves may be of use for other surface-based statistical analyses.

In this study, we used harmonic energy, dilatation, and elastic energy as measures of surface similarity, but strictly speaking, the maps do not form metrics in the formal mathematical sense. That is, given a particular source surface and target surface, the energy of the forward map computed with the method is not necessarily the same as the inverse map, and this is one limitation of our method. Certainly, from the viewpoint of shape space analysis and treating surfaces as points in shape space that are some distance apart, this

would be a useful feature to have, and indeed, there is recent work that incorporates this forward and backward symmetry into the metric optimization. For instance, Shi et al. [53] jointly estimate the forward and inverse maps in their optimal surface map formulation. Koehl and Hass define a *symmetric distortion metric* [34], which is a metric on the space of genus-zero Riemannian surfaces that is also invariant to mesh triangulation, and optimizes this constructing a spherical conformal map.

A common complaint against exact landmark matching is that the positions of underlying anatomical landmarks are uncertain. We advocate explicitly representing the space of possible landmark locations via sets of candidate landmarks, constructed either from automated methods [8] or by multiple manual tracings. Multiple mappings, each of which exactly map candidate landmarks to each other, can be used to propagate uncertainty in the landmarks to uncertainty in the mapping in a principled way. In contrast, approximate landmark matching as in `ConformalLS` does not incorporate landmark uncertainty information and gives no quantification of how landmark uncertainty impacts mapping.

As our method uses the formulation of discrete conformality in Bobenko et al. [9] to solve for a set of conformal factors that constitutes a hyperbolic uniformization metric, there are a number of necessary conditions for a solution to be possible. For surfaces with genus less than two, a number of cone point singularities with a total angle defect of at least 4π must be specified in order for a hyperbolic metric to be admissible. However, that may not be sufficient for there to be a solution to be guaranteed by the nonlinear optimization. In the literature, it has been noted that the key is to use reasonably well-conditioned triangulations in the sense that the triangles are neither too close to acute nor obtuse. Such a triangle is close to violating the triangle inequality and thus is near a singular point that the solver is unable to handle, and the method for rectifying this is to either flip an edge, collapse an edge, carrying out further subdivision, or some combination of these to eliminate the ill-conditioned triangles. Indeed, the method we use to compute orbifold metrics relies on well conditioned surface meshes that lack slivery or obtuse triangles. Nevertheless, there are

some pathological cases where the triangles are all relatively well-conditioned but a discrete conformal metric is unable to be solved. The example in Figure 4.5.1 was created from an initial mesh of genus two that admits a discrete conformal hyperbolic metric. A patch of the surface was subdivided and a vertex with valence 3 was extruded to form a local pyramid. This subdivided mesh continues to admit a hyperbolic metric until the sides of the pyramid approach equilateral triangles, in which case no valid solution is possible. In practice we found that such mesh configurations occurring in our anatomical datasets did not represent significant geometric features and manually removed them, but we note that this is an issue when working with the conformal factor approach. Finally, we note that recent work by Sun et al. [62] attempts to define an algorithm that incorporates diagonal edge flips together with vertex scaling for computing discrete conformal metrics, thereby addressing the problem of ill-conditioned triangles.

In conclusion, we use orbifold metrics to approximate the global minimum of a well-defined surface distortion energy and exactly match curve landmarks. This is a novel application of the orbifold canonical domain, which to the best of our knowledge has only been applied to align point landmarks. The method allows curve interior points to slide along each other during mapping to reduce surface distortion. Experiments suggest that this enhancement enables mappings that reduce surface distortion and identify relevant biological shape differences.

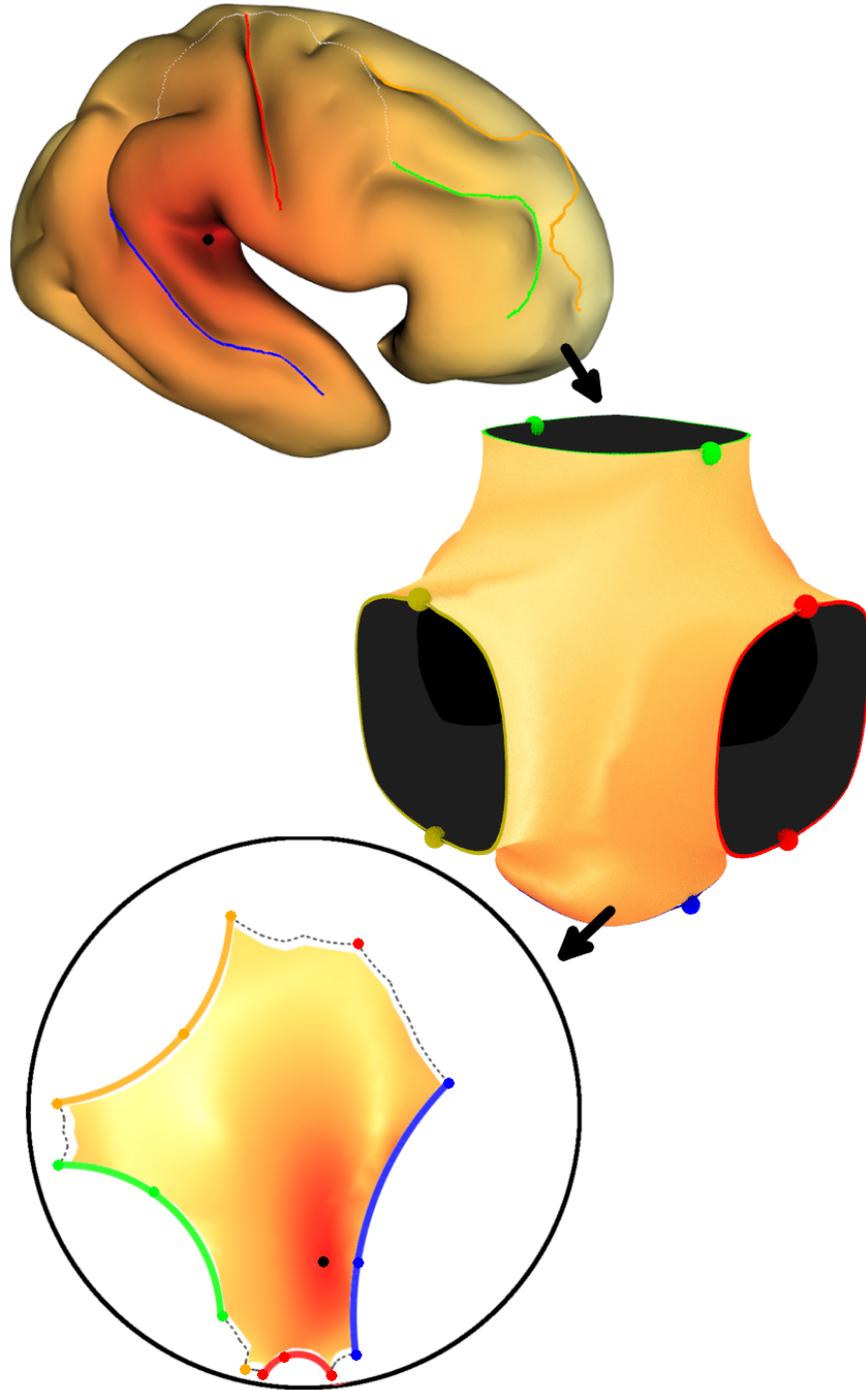


Figure 4.3.1: A schematic illustration of marking and initial mapping of a brain. In the above box, the brain is traced with four landmark curves, shown in the colored bold lines. Note that these curves should be viewed as holes, and the mesh with holes should be viewed as one half of the orbifold mesh formed by gluing two copies of this mesh along these holes. These are in turn connected up in a tree by additional vertex paths, shown in dotted lines. If we were to cut along this tree, the mesh unfolds into a k -gon. If we specify the same tree on the target surface, we have another k -gon which we can overlay on our initial k -gon to obtain an initial map. The meshes are colored based on the geodesic distance from the black point on the original surface to illustrate the change in metric.

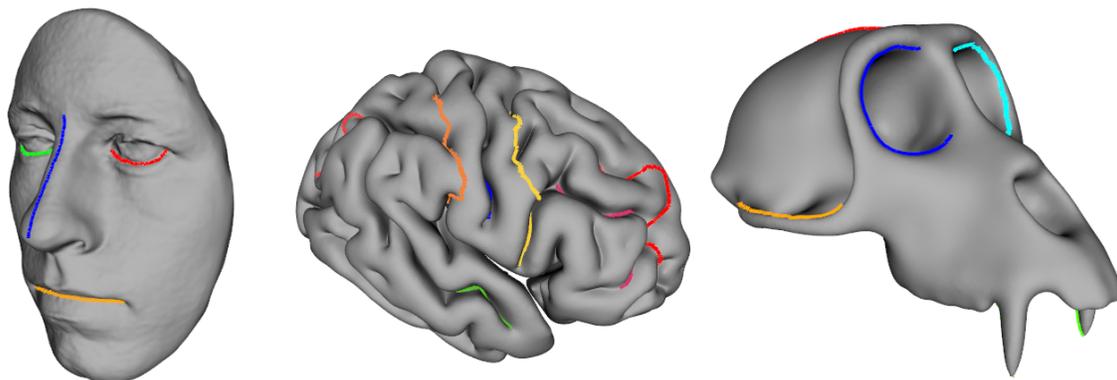
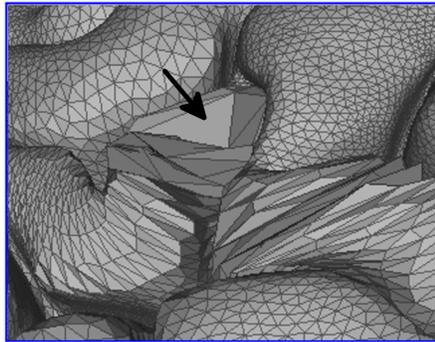
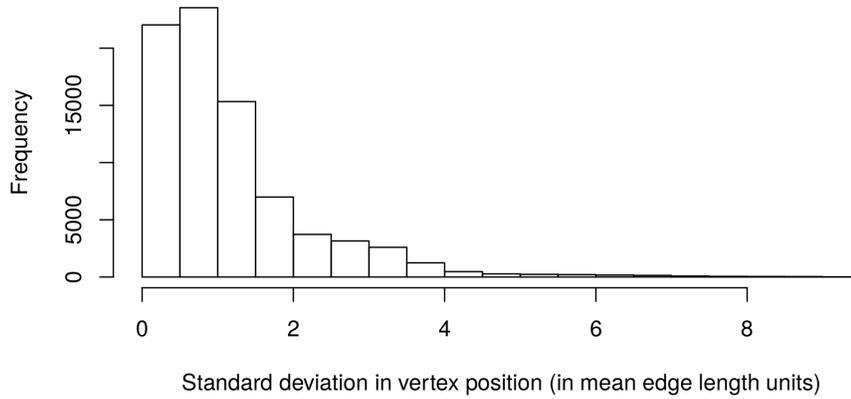
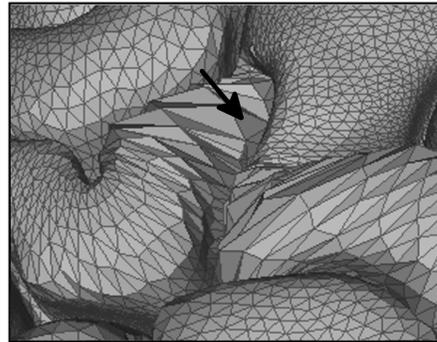


Figure 4.4.1: Example meshes from each dataset used to test mapping methods. Colored curves indicate expert traced landmarks. The brain and skull are slightly inflated for easier visualization of landmarks.

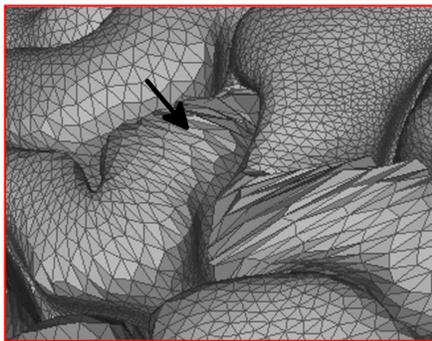
Standard deviation in final vertex position across different initial maps



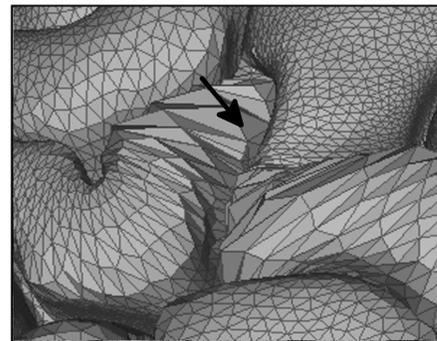
initial map 1



final map 1



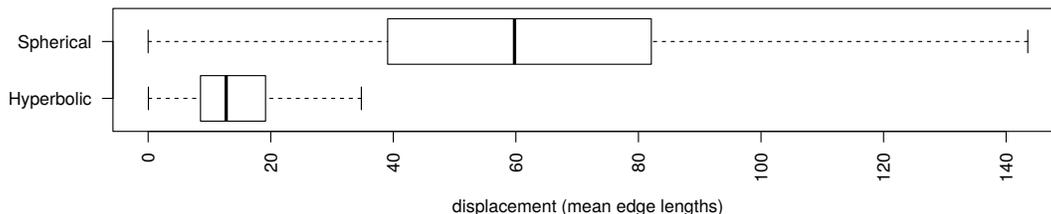
initial map 2



final map 2

Figure 4.4.2: We computed eight different initial maps on a pair of brains with respect to our method. Theoretically, irrespective of initialization the optimization should approximate the same minimal harmonic energy map. **Top:** The histogram of the variance with respect to each mesh vertex shows that, subject to varied initializations, vertices are generally placed consistently. **Bottom:** The boxes on the left show maps initialized by different marking trees, and on the right the similar approximate minimal harmonic energy maps after the optimization in Section 4.3.2

Overall landmark mismatch - all hyperbolic methods vs spherical



Mismatch at left-out landmarks relative to OrbifoldCurves

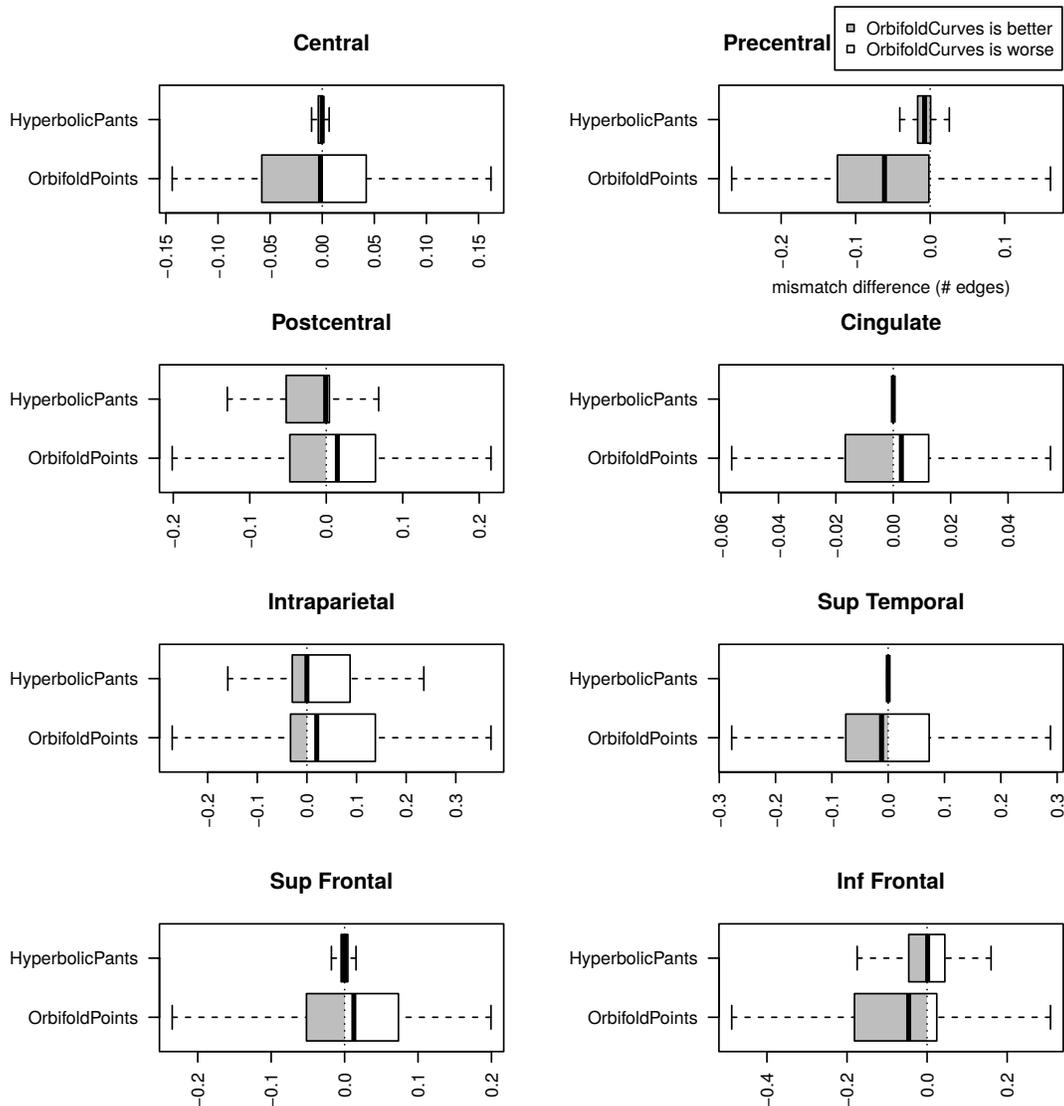


Figure 4.4.3: Boxplots show leave-one-out landmark mismatches for the brain dataset. Errors are given as average number of edges from nearest corresponding landmark vertex. The top boxplot shows that the hyperbolic methods performed similarly and overall better than the spherical-based method, `ConformalLS`. The following eight boxplots show errors relative to `OrbifoldCurves`, where the shaded portions indicate where `OrbifoldCurves` produces less error at that landmark. We note that face and skull dataset demonstrate the same relative performance between methods and are omitted for space.

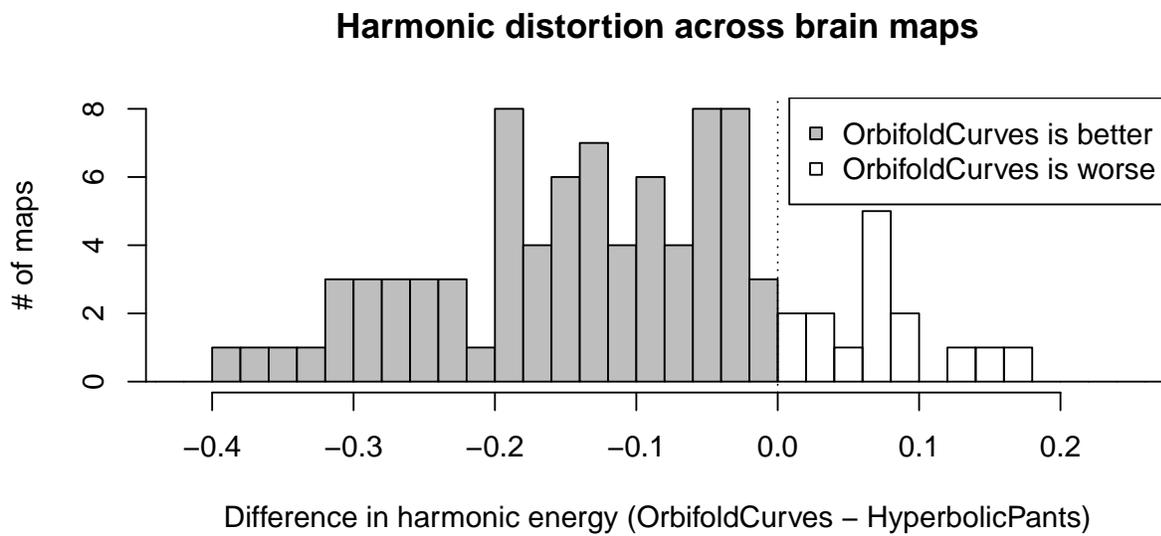


Figure 4.4.4: This shows the difference between `HyperbolicPants` and `OrbifoldCurves` in terms of harmonic energy across all maps in the Framingham brain dataset. Energy is totalled across all edges in a 10-ring neighborhood of the curve landmarks, which corresponds roughly to a 5mm radius. The shaded bars highlight the portion of maps where `OrbifoldCurves` achieves lower distortion. Note that the all `OrbifoldCurves` maps also outperforms `HyperbolicPants` in elastic energy.

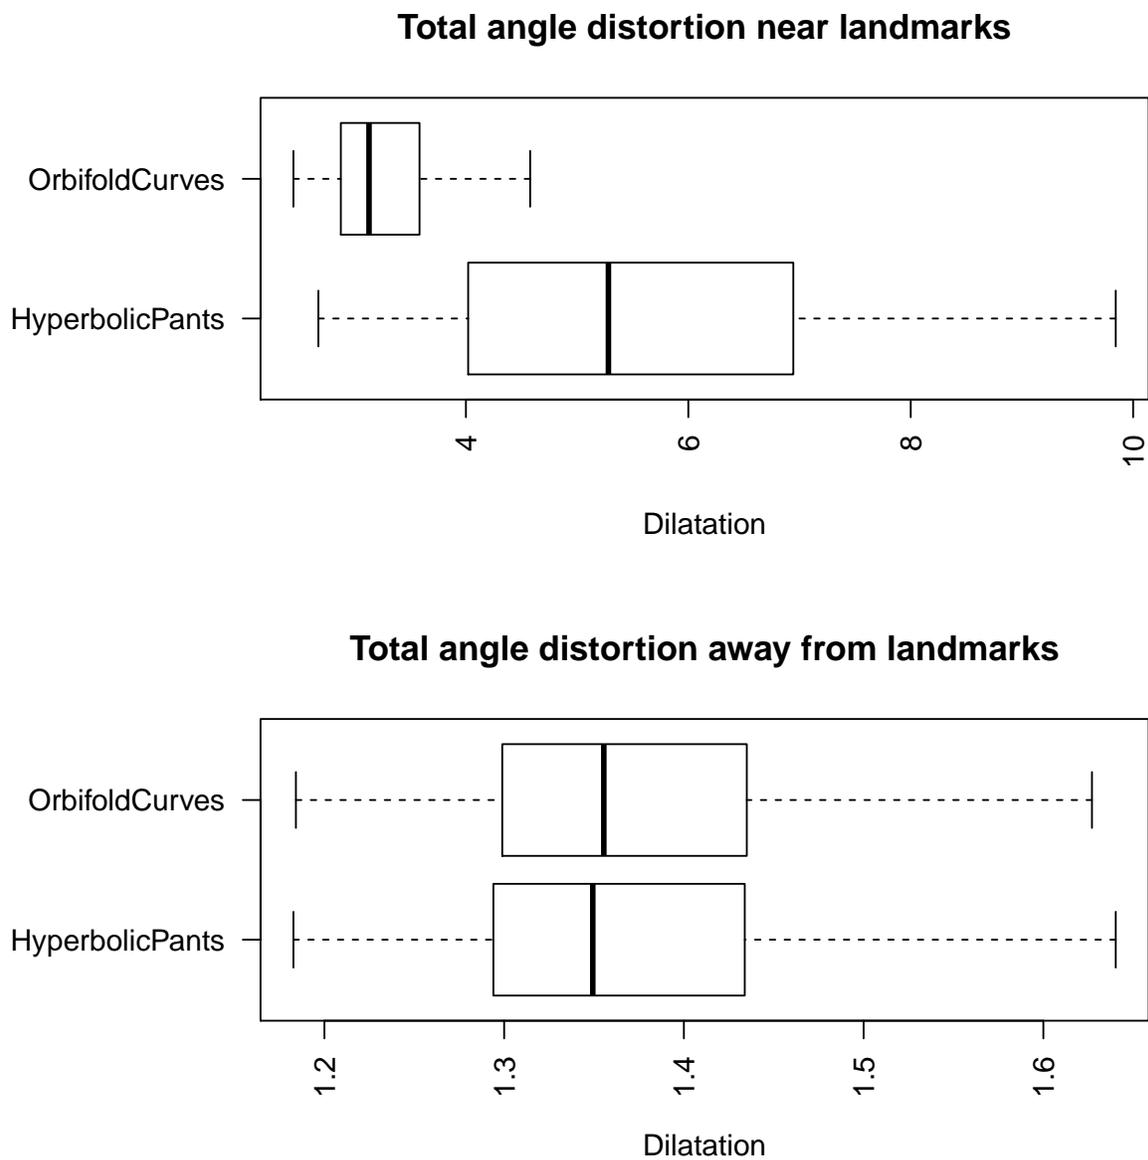


Figure 4.4.5: This shows the difference between `HyperbolicPants` and `OrbifoldCurves` in terms of local dilatation across maps in the Framingham brain dataset. Dilatation values range from 1 to ∞ ; 1 means the map locally exhibits no angle distortion and an infinitesimal circle maps to another circle, while 2 means the circle is mapped to an ellipse whose major axis is twice as long as its minor axis, and so on. Regions within a 10-edge neighborhood of the curve landmarks are considered near landmarks, and otherwise are considered away from landmarks.

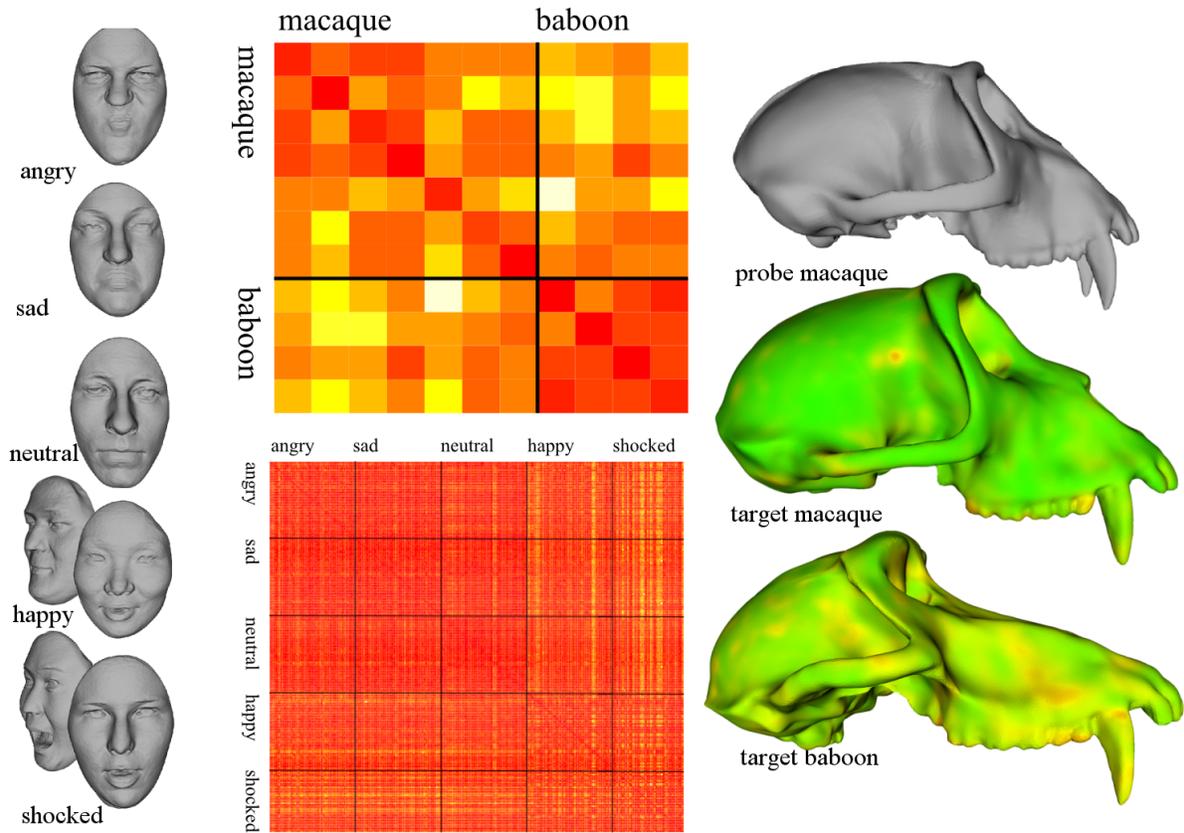


Figure 4.4.6: Visualization of the distortion of all pairwise mappings. Each cell represents the total harmonic distortion of a pair of surfaces, with values ranging from low (red) to high (white). The top matrix represents the pairwise matrix of 235 faces, grouped into five roughly equal sized blocks of expressions. The bottom matrix is the pairwise mappings of monkey skulls, ordered from left to right by 7 macaques and 4 baboons. **Left:** From top to bottom, they are representative models of the respective expressions: angry, sad, neutral, happy, and shocked. Two shocked examples show some of the variability that is evident by the brighter colors in the lower-right block of the face matrix. **Right:** Probe macaque mapped to a different example macaque (middle) and baboon (bottom). Colors visualize dilatation on the map to macaque (average 1.25) and baboon (average 1.45), respectively. This example suggests that closer species relationship may be reflected in the fossil geometry and captured by our method.

Relative difference of landmark mismatch across tracer maps

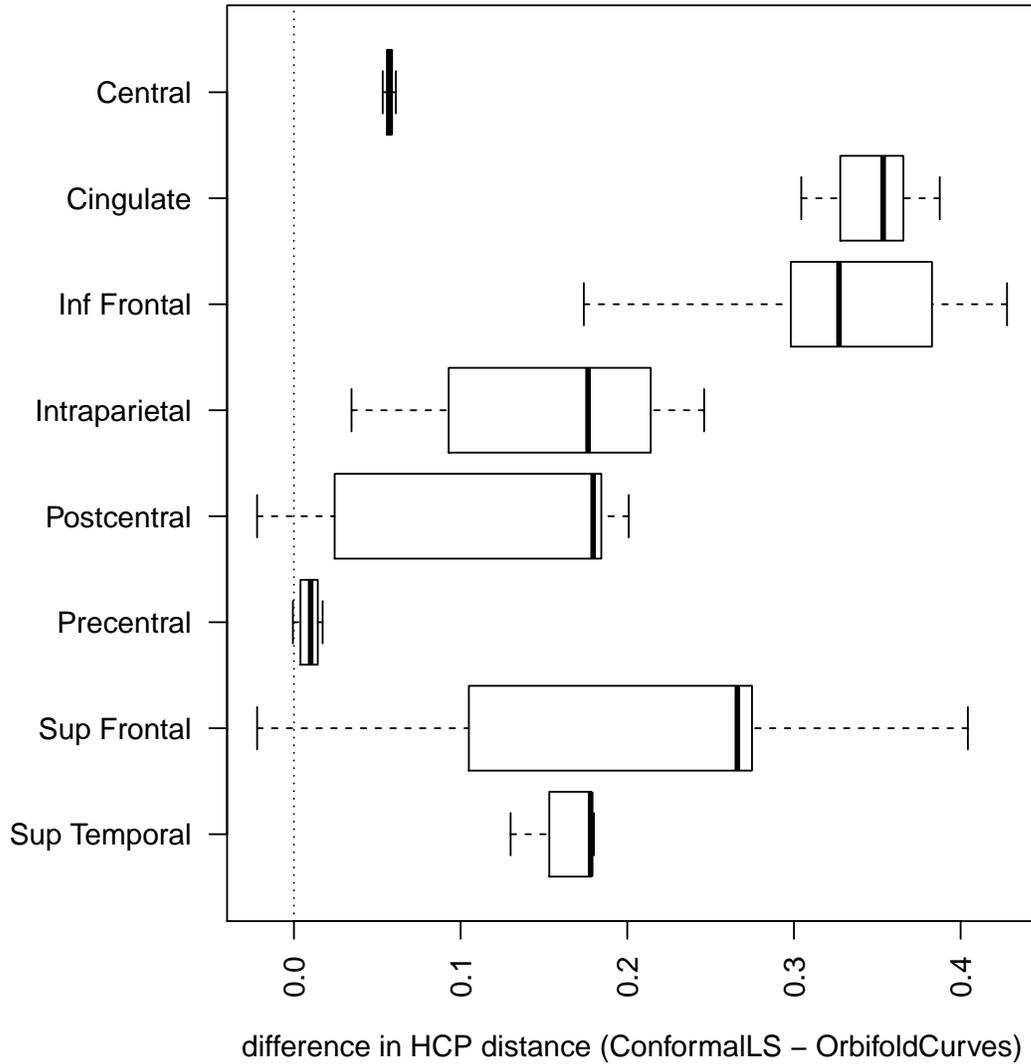


Figure 4.4.7: Differences of hyperbolic closest-point distance from `ConformalLS` to `OrbifoldCurves` methods, broken down by landmark. Boxes plotted to the right of the dotted line indicate that across all landmarks, `ConformalLS` shows greater mismatch than `OrbifoldCurves`. Overall, an exact matching method leads to results that have less extraneous variability and uncertainty in the landmarks than the results produced by a least-squares matching method.

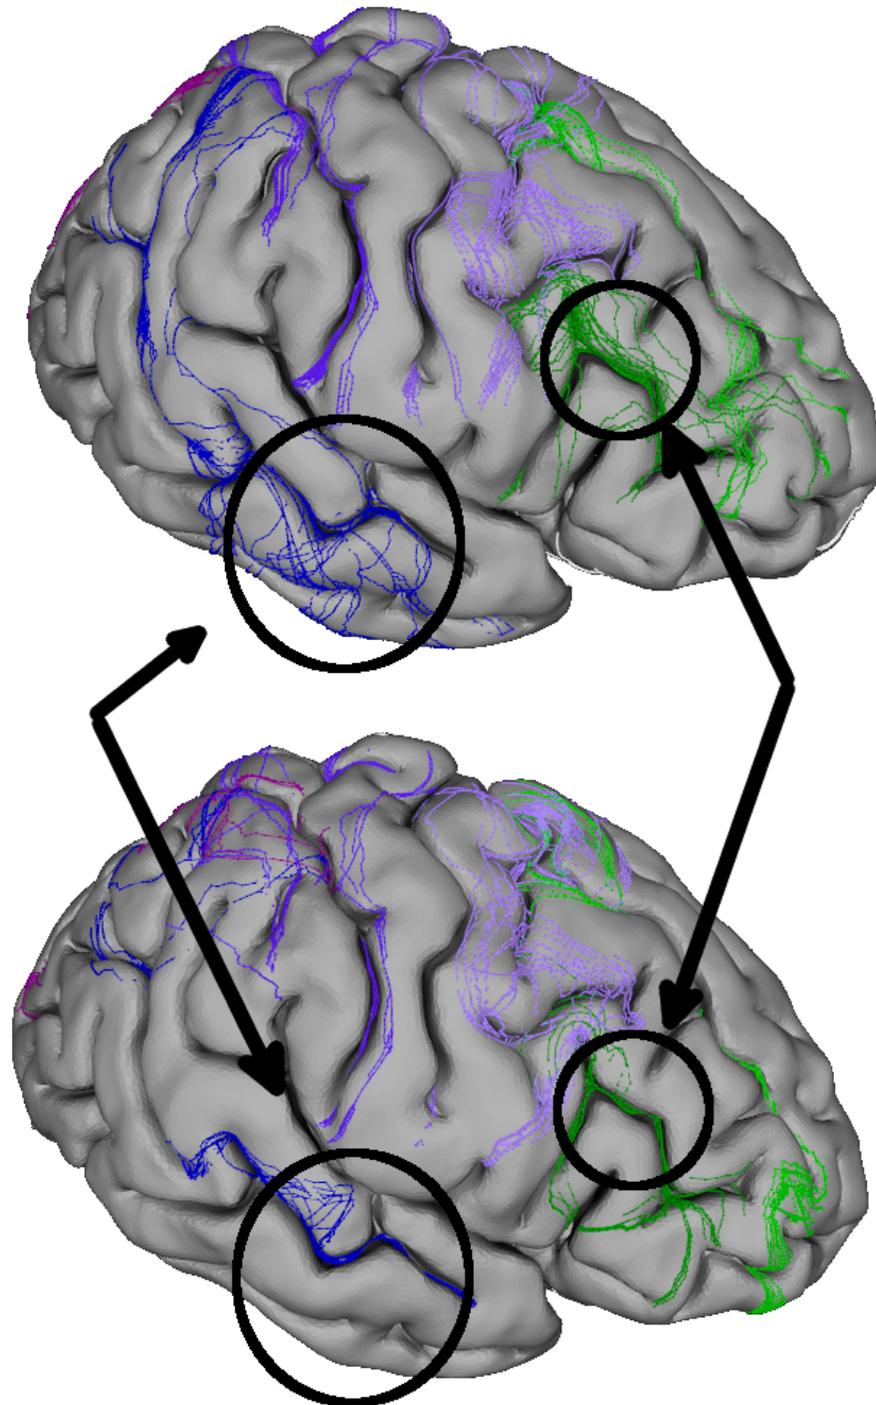


Figure 4.4.8: Multiple tracings mapped onto a single brain with `ConformalLS` (top) and `OrbifoldCurves` (bottom). Notice that the landmark variability is accounted for as tracer variability in the exact matching method, while the mapping of least squares is less predictable. Variability is less readily accounted for in a least-squares approach, which does not incorporate landmark uncertainty information and gives no quantification of how landmark uncertainty impacts mapping. Arrows point to corresponding regions where exact landmark matching imparts visible improvements in controlling sulcal mapping.

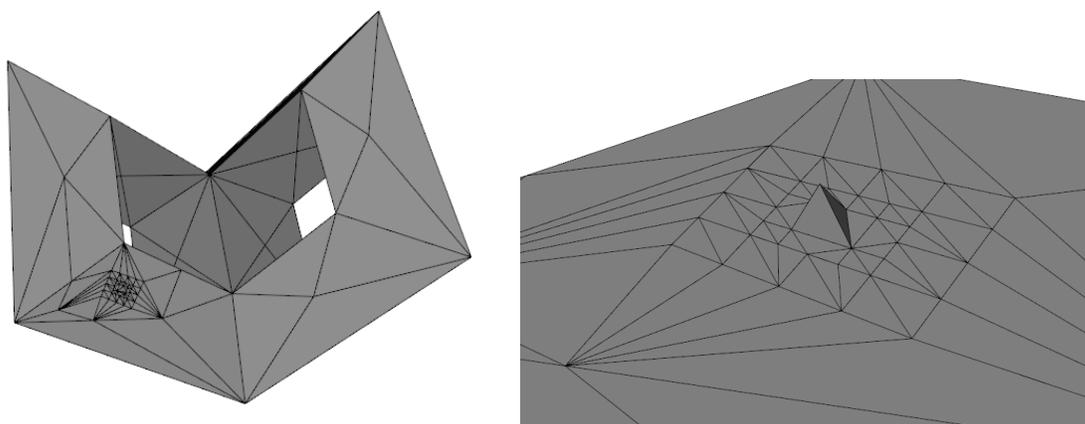


Figure 4.5.1: Visualization of a pathological mesh configuration in which the discrete conformal map cannot be computed without intervention. The triangles forming the pyramid in the bottom frame are close to equilateral, but their presence prevents the solution of an admissible set of conformal scale factors.

Chapter 5

Towards Angle Distortion-Minimizing Maps

Map smoothing typically employs metric distortion relaxation, notably harmonic energy relaxation, which aims to minimize the overall length distortion in a map. While it is understood that angle and length serve as complementary pieces of information that characterizes intrinsic geometry of a shape, there has been little empirical work done in the literature to compare metric distortion-optimizing maps such as harmonic maps with maps that optimize for angle distortion. In this chapter, we consider the idea of minimizing an energy that optimizes the map for angle distortion, subject to the landmark constraints imposed on the map. The approach follows the framework that we have setup for ourselves in the previous chapters (see Chapters 3 and 4) where we start from an initial correspondence and iteratively improve the solution until we reach an energy optimum, for some particular quantification of angle distortion as energy functionals. We propose three angle optimizing methods characterized by three different energy functionals, and we quantify and compare performance of each method in terms of each method's ability to minimize *dilatation* across the produced map. In Section 5.1, we review the theory of angle distortion optimizing maps, as well as current related work in the field. Then, Section 5.2 details our approach to the problem of angle-

distortion minimizing maps. Finally, we talk about the results of experiments that compare our angle-optimizing approach with a number of related angle-optimizing approaches in the literature, as well as the standard harmonic energy approach, in Section 5.3. We find that our proposed method is comparable to existing methods for angle optimization, with the `OrbifoldAngleShare` method giving better results, albeit with issues of reduced mesh resolution characteristic of conformal approaches. We show experimentally that mesh resolution artifact can be remedied with the introduction of additional landmark points but note that this is a limitation of the method. We compare harmonic (length) optimization to angle optimization for qualitative differences in output maps to assess the potential of gleaned independent shape information. Finally, we conclude from the results that, for reasons of simplicity and robustness, harmonic energy relaxation is preferable in practice.

5.1 Theory and Prior Work

5.1.1 Dilatation

We will quantify angle distortion in terms of *dilatation*, which can be defined concisely as the ratio of the singular values of the Jacobian of a linear mapping. A mapping is discretely represented by piecewise-linear maps of triangles, so consider that the action of a map f as restricted to a triangle A in \mathbb{R}^2 is a linear map to another triangle in \mathbb{R}^2 as in Figure 5.1.1. Note that even if the mapping domain is in \mathbb{R}^3 , we can always consider this picture in \mathbb{R}^2 without loss of generality because there exists a conformal parameterization that will send a triangle in \mathbb{R}^3 to \mathbb{R}^2 . We can then compute the dilatation of each individual triangle under the mapping by applying the following equations [51]. First, refer to the notation for a triangle mapping f in Figure 5.1.1. Treating the q_i 's as vectors, and denoting the area of the source triangle A , we can write the partial derivatives in terms of the parameterization

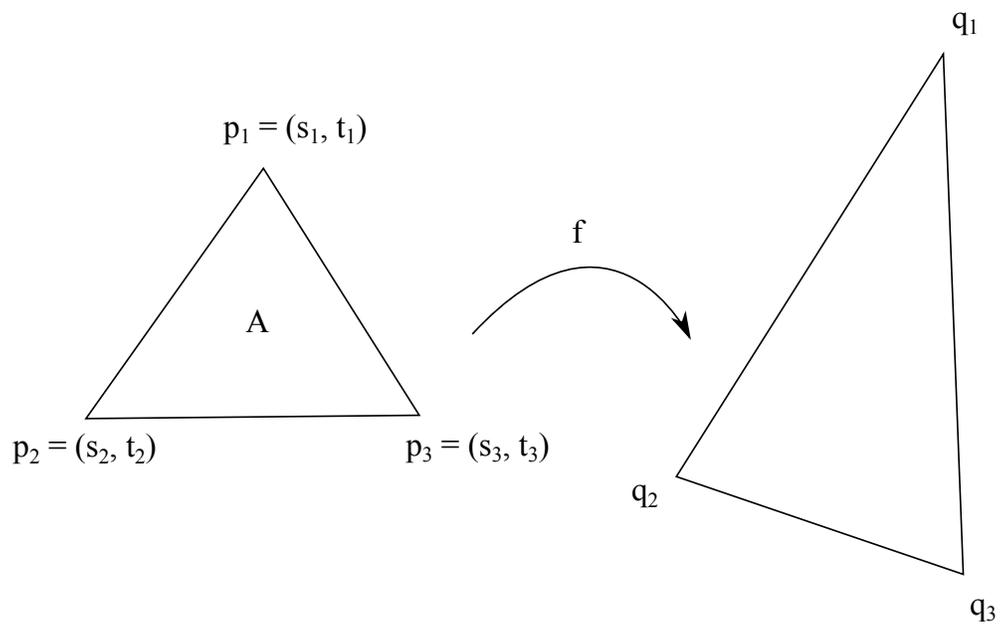


Figure 5.1.1: Set up for computing dilatation of a triangle mapping. Dilatation is the ratio of the singular values of the Jacobian of the linear mapping f .

by

$$f_x = \frac{(q_1(t_2 - t_3) + q_2(t_3 - t_1) + q_3(t_1 - t_2))}{2A}, \quad (5.1.1)$$

$$f_y = \frac{(q_1(s_3 - s_2) + q_2(s_1 - s_3) + q_3(s_2 - s_1))}{2A} \quad (5.1.2)$$

From here, we consider the large and small singular values of the Jacobian formed from these partial derivatives as follows. Letting $a = f_x \cdot f_x$, $b = f_x \cdot f_y$, and $c = f_y \cdot f_y$, then the maximum and minimum singular values are given, respectively, as

$$\Gamma = \sqrt{\frac{1}{2} \left((a + c) + \sqrt{(a - c)^2 + 4b^2} \right)}, \quad (5.1.3)$$

$$\gamma = \sqrt{\frac{1}{2} \left((a + c) - \sqrt{(a - c)^2 + 4b^2} \right)}. \quad (5.1.4)$$

Then dilatation is given by the ratio Γ/γ . A value of 1 represents an absence of angle distortion, as is the case in a conformal mapping. We note that translation and uniform scaling of the triangles of this 1-ring in a Euclidean parameter space does not change dilatation, so without loss of generality, we can consider the coordinates of the center of the 1-ring to coincide with the origin. This does not hold in hyperbolic space, but we approximate this configuration in hyperbolic space by parameterizing the 1-ring vertices on the Poincare disk with the center vertex on the origin, connecting the sides as Euclidean line segments. For small 1-ring neighborhoods, the hyperbolic lengths are equivalent to Euclidean lengths up to a constant factor.

5.1.2 Conformal maps

In any discussion of low angle distortion maps, we must begin by mentioning *conformal maps*, which are angle-preserving maps. The Riemann uniformization theorem states that, depending on the topology of the surface, it is always possible to map that surface onto either the round sphere, the Euclidean disk, or the hyperbolic disk in such a way that

no angle distortion is introduced. To simplify the general mapping problem between two surfaces, we consider the surfaces to be conformally mapped onto a canonical space (i.e. sphere, disk, or hyperbolic disk). Euclidean planar domains (i.e. disk) can subsequently be mapped conformally, for example, by Schwarz-Christoffel mappings, and particularly with the sphere, the class of Möbius transformations can map one sphere to the other while preserving angles, and specifying the mapping of any three distinct points determines the Möbius transformation. The inherent limitation to note here is that three points exactly determine a conformal transformation, and any attempt to align additional points must be done at the expense of conformality or landmark alignment. We will not consider pure conformal maps in this chapter’s experiments.

5.1.3 Prior approaches

While a benefit of the conformal map lies in its simplicity, it also means that in the context of constrained surface mapping, we are unable to exactly match more than three points while preserving all angles at the same time. Methods in the literature have taken the approach of either computing the optimal conformal map that minimizes a particular matching energy or, starting from a conformal map, iteratively deforming the map so as to trade off between constraint satisfaction and conformality preservation. For the former case, Lui et al. [41] computed an initial conformal map of topological spherical surfaces to the canonical round sphere. From there, an energy functional that balances between harmonic energy and landmark match is used:

$$E(f) = \frac{1}{2} \sum_{i,j} \omega_{ij} \|f(v_i) - f(v_j)\|^2 + \frac{\lambda}{2} \sum_{k=1}^n \|f(q_k) - p_k\|^2 \quad (5.1.5)$$

where the set of landmarks q_1, \dots, q_n are identified with corresponding points on the target surface $L(q_i) = p_i$ and λ is a Lagrange multiplier that specifies the relative importance of landmark matching. This makes use of the fact that for the spherical case, a harmonic

map, that is, a map that minimizes harmonic energy as defined by the first term of Equation (5.1.5), remains a conformal map. Therefore, minimizing the equation above results in a map that is no longer conformal, but attempts to fit the landmark constraints in a least squares sense.

In the work by Koehl and Hass [32], the approach was to not rely on explicit landmarks, but rather to define a quantity of elastic energy as measured over the entire surface. In a later work [34], they extended their method so that this energy is independent of the triangulation of the surface and accounts for distortion in the backward direction in addition to the forward direction. In each instance, the method produces a fully conformal map in the following manner. Both surfaces are mapped conformally to the round sphere, and then the optimal Möbius transformation determines the sphere-to-sphere mapping that minimizes the *symmetric distortion energy*.

Beltrami coefficient

Recall from the preliminary discussion that quasiconformal maps are a flexible superset of conformal maps, defined by the unit of angular distortion known as the Beltrami coefficient μ . Specifically, a function $f : \mathbb{C} \rightarrow \mathbb{C}$ is quasiconformal if

$$\frac{\partial f}{\partial \bar{z}} = \mu(z) \frac{\partial f}{\partial z} \tag{5.1.6}$$

where μ satisfies $\|\mu\|_\infty < 1$ is called the *Beltrami coefficient*. Work in quasiconformal surface mapping has almost exclusively focused on computing a special type of quasiconformal map with a particular structure. Suppose the mapping $f : \mathbb{C} \rightarrow \mathbb{C}$ defined on the complex plane can be expressed as

$$f = h \circ A_k \circ g \tag{5.1.7}$$

where g and h are conformal maps and A_k is a linear stretch of factor k . Given this characterization, the Beltrami coefficient can be written in the form

$$\frac{f_{\bar{z}}}{f_z} = k \frac{\bar{\theta}}{\theta} \quad (5.1.8)$$

where k is a scalar function that captures the conformal distortion in the map, and θ is a complex function. A map f that satisfies these equivalent conditions is called a Teichmüller map. Weber et al. [71] attempted to discretize the *least squares Beltrami energy*

$$E_{LSB} = \int_D |f_{\bar{z}} - k \frac{\bar{\theta}}{\theta} f_z|^2 dA \quad (5.1.9)$$

which is derived directly from taking the difference between both sides of the Beltrami equation (Equation (5.1.6)). Here, the unknowns are f , θ , and k , and there is the constraint that $\theta_{\bar{z}} = 0$ ($\theta_{\bar{z}}$ is the Wirtinger derivative with respect to the conjugate of z) and f satisfies the boundary condition if given.

Lui et al. [40] represents a different school of thought, seeking to use the μ as a surface representation and optimize it indirectly by performing normalized smoothing of $|\mu|$ over the surface. They observe that it is not necessary to compute explicitly all the components of the map as in Weber et al. but by ensuring that the map exhibits a uniform distribution of Beltrami coefficient norm across the surface, it qualifies as a Teichmüller map. They demonstrate that for surfaces that are convertible to the topological disk, they are able to achieve maps that have a distribution of Beltrami coefficient norms that approximates the characteristic constant norm of the Teichmüller map.

All of these methods attempt to compute the Teichmüller map, which minimizes the L_∞ norm of $|\mu|$, but to the best of our knowledge, there is no method that attempts to compute a map that minimizes the L_2 norm of $|\mu|$, which would be analogous to what the harmonic map does for isometric maps. We want to be able to minimize angle distortion directly by iteratively updating the vertex positions.

Angle market share

Alternatively, angle distortion can be measured by considering the distribution of angle around a point. Intuitively, this is called *angle market share* [25], and how well the distribution is preserved gives a measure of how close the mapping is to conformal at a given point. According to Hurdal et al., this interpretation is more in keeping with the Riemann geometric point of view, as a discrete representation of a curved manifold does not necessarily have exactly 2π angle surrounding every point as in the continuous case, and so a proportional measure of distortion is more suitable. As such, a discrete map is conformal if we can preserve angle market share. In general, the distribution will not be preserved, and so we can formulate a local operator that updates the vertex position so as to restore the distribution to the extent that is possible, given the configuration of the one-ring neighbors. Unfortunately, although the metric has been defined and used to evaluate the conformality of certain mesh parameterization methods [26], there has been no line of research that optimizes the angle market share directly. One possible approach is due to Zhou and Shimada [78] formulates a torsional equivalent of uniform-weighted Laplacian smoothing to do *uniform angle smoothing*. The method is a local heuristic that does not explicitly optimize a global energy nor guarantee optimality.

Another work that also tackles the problem of angle distortion by treating the angles directly is presented by Sheffer et al. [52], who proposes to minimize the following functional defined in the space of angles of a triangulation:

$$E(\alpha) = \sum_{t \in T} \sum_{k=1}^3 \frac{1}{w_k^t} (\alpha_k^t - \beta_k^t)^2, \quad (5.1.10)$$

where α_k^t are the unknown planar angles and β_k^t are the optimal angles. Here, the weights are taken to be $\frac{1}{\beta_k^{t/2}}$ to reflect relative angular distortion. Some constraints, such as the triangular inequality, the requirement that the sum of all angles around each vertex add up to 2π , and that pairs of matched triangle edges are of the same length, are included as separate terms

and optimized with weighting factors to assure that the solution set of angles yields a suitable planar parameterization. Linear systems are solved once to compute a set of angles and then once again to reconstruct a set of planar coordinates for the vertices. The issue with this method is that the size of the system, $4f + 2v$ for $f = |T|$ number of mesh triangles and v the number of interior vertices, is large. This can be computationally mitigated with some matrix manipulation and rewriting of the optimization loop, but another issue is that the method may fail to find a valid configuration if the one-ring neighborhood of a vertex becomes too small.

MIPS

We would like a direct way to optimize dilatation, but dilatation itself is a quasiconvex function that is not easily minimized. However, we can derive an alternative energy functional in terms of dilatation that is better behaved.

In Hormann's dissertation [24], dilatation is used to define a class of shape deformation functionals. These functionals are

1. invariant to translation
2. invariant to orthogonal transformation (rotation and reflection)
3. invariant to uniform scaling

Letting dilatation be $\rho = \Gamma/\gamma$, then *MIPS energy* is defined as $f(\rho) = \rho + 1/\rho$. The function $f(\rho)$ is the well-known MIPS (most isometric parametrizations) energy and can be understood as representing Dirichlet energy normalized by area. Hormann showed that the function $f(\rho)$ is convex in the kernel of the one-ring of a vertex.[24] Therefore, as suggested by Figure 5.2.2, the energy is locally convex, so we can define a local operator to minimize our shape deformation metric for each one-ring neighborhood.

5.2 Methods

The following section details the experiments we performed to compare the relaxation method proposed in the previous section, which we will call `OrbifoldMIPS` for the remainder of the discussion, to similar relaxation methods that use a metric-based energy. We compute maps from surface A to surface B in the following way:

1. Compute hyperbolic metrics for A and B .
2. Compute an initial map from A to B such that the landmarks are in correspondence.
3. Perform `RelaxationMethod(i)` on each vertex until convergence of the map-wide dilatation.

Where `RelaxationMethod(i)` is one of the following methods, and whose details are described in the following subsections:

1. `OrbifoldPoints`
2. `OrbifoldBeltrami`
3. `OrbifoldAngleShare`
4. `OrbifoldMIPS`

For each of the relaxation methods, the same initial map is established using the hyperbolic orbifold metrization and initial mapping framework described in Chapter 3. This results in one optimal map for each method, which we compare quantitatively and qualitatively.

5.2.1 OrbifoldPoints

We optimize for harmonic energy by the iteratively relaxing each vertex according to the Euclidean Dirichlet operator. This is used as an approximation of hyperbolic harmonic energy, which as discussed previously in Chapter 3, has a unique global optimum within the

homotopy class fixed by the initial mapping. Harmonic energy quantifies metric distortion in a least-squares sense.

5.2.2 OrbifoldBeltrami

A Teichmüller map has constant Beltrami coefficient norm everywhere. Therefore, this method seeks to iterate on the current mapping, driving its Beltrami coefficient norm distribution, as measured across all mesh faces, towards a distribution that is constant. To this end, we minimize the norm of the Beltrami coefficient as a measure of angle distortion by way of the QC iterations as described by Lui et al [40] [39]. The formulation of the linear Beltrami solver in [40] is particular to 2D Euclidean meshes, so we adapt the method to the manifold as in Zhang et al [77]. Given the parameterization from the hyperbolic orbifold framework, we apply a vertex-based relaxation as follows. Consider triangle $T = [v_i, v_j, v_k]$ mapped under f , so $f(v_I) = w_I$ for $I = i, j, k$. Denote v_I, w_I as complex numbers i.e.

$$\begin{aligned} v_I &= (g_I, h_I) \\ w_I &= (s_I, t_I) \end{aligned}$$

The Beltrami coefficient can be written in terms of the partial derivatives of the linear mapping that acts on T . Denoting the partial derivatives as follows

$$J_f(x, y) = \begin{pmatrix} a_T & b_T \\ c_T & d_T \end{pmatrix} \tag{5.2.1}$$

The linear mapping can be expressed as

$$f|_T(x, y) = \begin{pmatrix} a_T x + b_T y + n_T \\ c_T x + d_T y + m_T \end{pmatrix}. \tag{5.2.2}$$

Consider that the Wirtinger derivatives are linear operators given by

$$\frac{\partial}{\partial z} = \frac{1}{2} \left(\frac{\partial}{\partial x} - i \frac{\partial}{\partial y} \right), \quad (5.2.3)$$

$$\frac{\partial}{\partial \bar{z}} = \frac{1}{2} \left(\frac{\partial}{\partial x} + i \frac{\partial}{\partial y} \right). \quad (5.2.4)$$

If we combine this fact with the Beltrami equation ((2.1.7)), the Beltrami coefficient can be put in terms of the partial derivatives in Equation (5.2.1):

$$\mu(T) = \frac{(a_T - d_T) + \sqrt{-1}(c_T + b_T)}{(a_T + d_T) + \sqrt{-1}(c_T - b_T)}. \quad (5.2.5)$$

Knowing the coordinates of the triangle, we can solve for the Beltrami coefficient explicitly in terms of the coordinates:

$$\begin{aligned} a_T &= A_i^T s_i + A_j^T s_j + A_k^T s_k \\ b_T &= B_i^T s_i + B_j^T s_j + B_k^T s_k \\ c_T &= A_i^T t_i + A_j^T t_j + A_k^T t_k \\ d_T &= B_i^T t_i + B_j^T t_j + B_k^T t_k \end{aligned} \quad (5.2.6)$$

where

$$\begin{aligned} A_i^T &= (h_j - h_k) / \text{Area}(T) \\ A_j^T &= (h_k - h_i) / \text{Area}(T) \\ A_k^T &= (h_i - h_j) / \text{Area}(T) \\ B_i^T &= (g_k - g_j) / \text{Area}(T) \\ B_j^T &= (g_i - g_k) / \text{Area}(T) \\ B_k^T &= (g_j - g_i) / \text{Area}(T) \end{aligned} \quad (5.2.7)$$

We can represent the Beltrami coefficient as the symmetric transformation matrix $A(T)$ if we let $\mu(T) = (\rho(T), \tau(T))$:

$$A = \begin{pmatrix} \alpha_1 & \alpha_2 \\ \alpha_2 & \alpha_3 \end{pmatrix}$$

where

$$\begin{aligned} \alpha_1 &= \frac{(\rho - 1)^2 + \tau^2}{1 - \rho^2 - \tau^2} \\ \alpha_2 &= \frac{2\tau}{1 - \rho^2 - \tau^2} \\ \alpha_3 &= \frac{1 + 2\rho + \rho^2 + \tau^2}{1 - \rho^2 - \tau^2} \end{aligned}$$

Definition 5.2.1. Consider a vector field $\vec{V} = (V_1, V_2)$ restricted to each triangle T of a mesh. The *discrete divergence* of that vector field at a particular vertex v_i on the mesh is given by

$$\nabla \cdot (\vec{V})(v_i) = \sum_{T \in N_i} A_i^T V_1(T) + B_i^T V_2(T). \quad (5.2.8)$$

Algebraically, we can verify that

$$\sum_{T \in N_i} A_i^T b_T = \sum_{T \in N_i} B_i^T a_T. \quad (5.2.9)$$

and, stated in terms of discrete divergence, this is

$$\nabla \cdot \begin{pmatrix} -u_y \\ u_x \end{pmatrix} = 0 \quad (5.2.10)$$

We can use the Beltrami equation to calculate A from μ and work out that

$$\nabla \cdot (A \nabla u) = \nabla \cdot \left(A \begin{pmatrix} u_x \\ u_y \end{pmatrix} \right) = \nabla \cdot \begin{pmatrix} -u_y \\ u_x \end{pmatrix} = 0 \quad (5.2.11)$$

So from this, we arrive at the following constraint (Equation 5.12 from [40]) introduced at each vertex of the mesh

$$\sum_{T \in N_i} A_i^T [\alpha_1(T) a_T + \alpha_2(T) b_T] + B_i^T [\alpha_2(T) a_T + \alpha_3(T) b_T] = 0 \quad (5.2.12)$$

The α 's come from the matrix representation A of the Beltrami coefficient $\mu(T)$ at triangle T as defined in the previous section.

Given a set of Beltrami coefficients assigned to the facets of a mesh, we would like to compute the vertex coordinates that satisfy the Beltrami equation, thus getting a quasiconformal map. For every vertex v_i , we get the equation

$$\begin{aligned} & \sum_{T \in N_i} A_i^T [\alpha_1(T) a_T + \alpha_2(T) b_T] + B_i^T [\alpha_2(T) a_T + \alpha_3(T) b_T] \\ &= \sum_{T \in N_i} (A_i^T \alpha_1(T) + B_i^T \alpha_2(T)) a_T + (A_i^T \alpha_2(T) + B_i^T \alpha_3(T)) b_T \\ &= \sum_{T \in N_i} (A_i^T \alpha_1(T) + B_i^T \alpha_2(T)) (A_i^T s_i + A_j^T s_j + A_k^T s_k) \\ &+ (A_i^T \alpha_2(T) + B_i^T \alpha_3(T)) (B_i^T s_i + B_j^T s_j + B_k^T s_k) \\ &= \sum_{T \in N_i} [(A_i^T \alpha_1(T) + B_i^T \alpha_2(T)) A_i^T + (A_i^T \alpha_2(T) + B_i^T \alpha_3(T)) B_i^T] s_i \\ &+ [(A_i^T \alpha_1(T) + B_i^T \alpha_2(T)) A_j^T + (A_i^T \alpha_2(T) + B_i^T \alpha_3(T)) B_j^T] s_j \\ &+ [(A_i^T \alpha_1(T) + B_i^T \alpha_2(T)) A_k^T + (A_i^T \alpha_2(T) + B_i^T \alpha_3(T)) B_k^T] s_k \end{aligned} \quad (5.2.13)$$

$$\begin{aligned} &= \sum_{T \in N_i} c_i s_i + c_j s_j + c_k s_k \\ &= 0 \end{aligned} \quad (5.2.14)$$

where N_i is the pizza of triangles centered at v_i , and the vertex coefficients c_I are fully determined. This map optimizes for angle distortion, and in particular aims to approximate the Teichmüller map, which minimizes the maximal angle distortion over the map.

5.2.3 OrbifoldAngleShare

We modify the algorithm described in Zhou and Shimada [78] in order to restore the angle market share. The following algorithm operates on each vertex v_i .

1. Consider edge e_{ji} to be a lever that pivots on v_j . We want to adjust the lever so that the angle to the left and right of e_{ji} is equal. That is, there is an rotation that makes e_{ji} an angle bisector. Compute the rotation and record where the vertex v_i goes, v_i^j .
2. Update the position of v_i to be the average of all the v_i^j 's.

$$v'_i = \sum_{j \in N(i)} v_i^j / |N(i)|$$

While the method tries to approach the problem of angle distortion directly by working with the angles on the triangulation, the method is only concerned with making the angles as uniform as possible. Therefore, we introduce the following modification of the algorithm that considers the angle market share of the source 1-ring and attempts to achieve the same distribution through the rotations. We will call this approach *angle-ratio smoothing*. Suppose we are looking at the image of the one-ring of vertex v_i under some mapping f onto a target surface – refer to Figure 5.2.1. We restore the ratio of the angle on either side of each edge under the mapping. So we adjust the position of v_i so that the ratio of the angles on either side of each edge on the target surface matches that ratio on the source surface.

1. Determine the ratio between the two adjacent angles $\gamma = \alpha_2/\alpha_1$ in the source surface.
2. Calculate the angle adjustment required to rotate the edge ij to achieve the ratio of

angles γ in the target surface

$$\beta_j = (\alpha_2 - \gamma\alpha_1)/(1 + \gamma)$$

3. Rotate the edge by β_j about v_j

$$x'_j = x_j + (x_i - x_j) \cos \beta_j - (y_i - y_j) \sin \beta_j,$$

$$y'_j = y_j + (x_i - x_j) \sin \beta_j + (y_i - y_j) \cos \beta_j$$

4. Update the position of v_i with the average of these rotated points

$$x_i^{t+1} = \frac{1}{N(v_i)} \sum_{j \in N(i)} x'_j$$

$$y_i^{t+1} = \frac{1}{N(v_i)} \sum_{j \in N(i)} y'_j$$

The method is a local heuristic that attempts to drive the the angle ratios on the mapped domain to match that on the source domain. In particular, we note that a drawback of the heuristic does not explicitly account for the angle ratios of the center vertex.

5.2.4 OrbifoldMIPS

From the first subsection, we have a definition from which to compute dilatation as a quantity of angle distortion, and we would like to use this as an optimization metric. Unfortunately, dilatation itself is a quasiconvex function, but we can derive an alternative energy functional in terms of dilatation that is better behaved. Letting dilatation be $\rho = \Gamma/\gamma$, then *MIPS energy* is defined as $f(\rho) = \rho + 1/\rho$. Hormann showed that the function $f(\rho)$ is convex in the kernel of the one-ring of a vertex.[24] Therefore, as suggested by Figure 5.2.2, the energy is locally convex, so we can define a local operator to minimize our shape deformation metric for each one-ring neighborhood by gradient descent with backtracking line search.

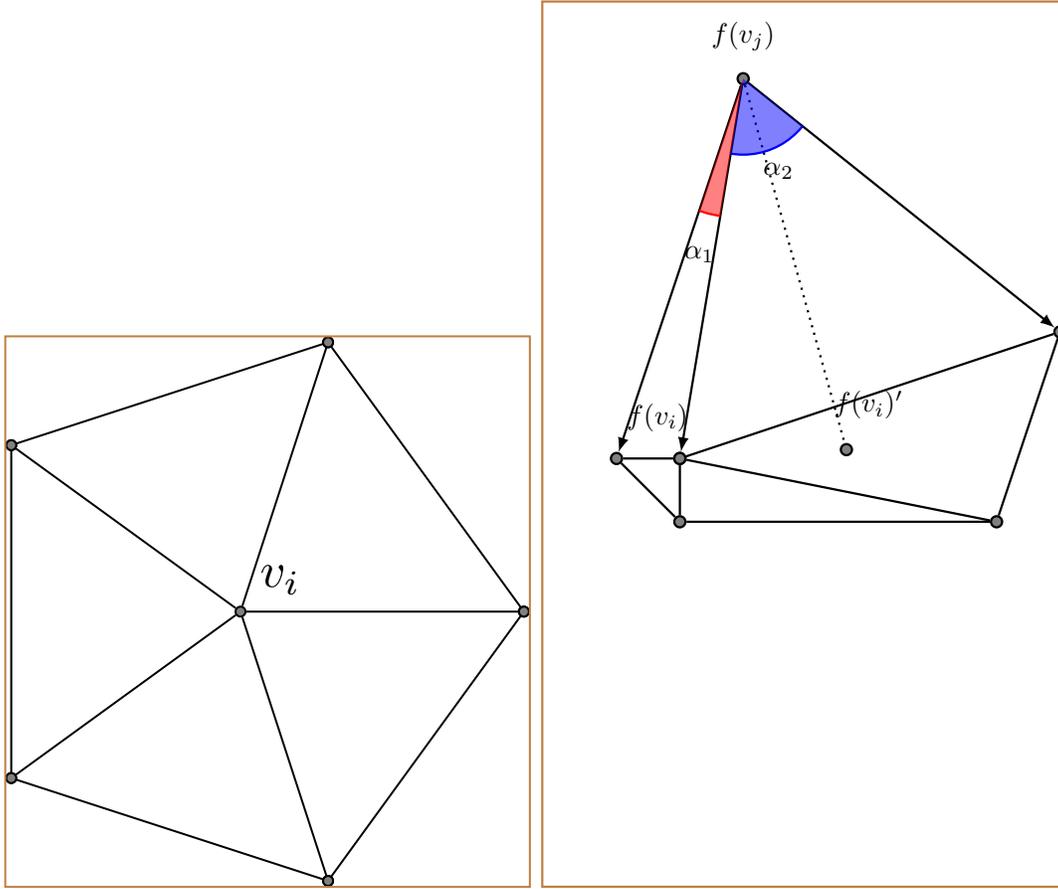


Figure 5.2.1: Illustrates the mechanic of Zhou and Shimada's torsion-based relaxation.

An explicit equation for dilatation can be computed in terms of the coordinates of the 1-ring neighbor vertices, however, due to the complexity of explicitly computing the gradient of this equation, we opt to determine the gradient by numerical differentiation using finite differences. We note that the gradient can be well-approximated with a fixed step size by utilizing dilatation's invariance to uniform scaling, and standardize the size of the 1-ring by scaling the length of the shortest 1-ring edge to unit length. In this way, we can use a fixed step size (say 0.01) with which to compute the first derivative by finite differences.

If we restrict the space of optimization to be in the convex region within the 1-ring such that placement of the center vertex anywhere in that region will not cause the orientation of any of the 1-ring triangles to be reversed. Thus, we define the iterative scheme `OrbifoldMIPS` to minimize MIPS energy distortion in a given discrete map. While solving for local opti-

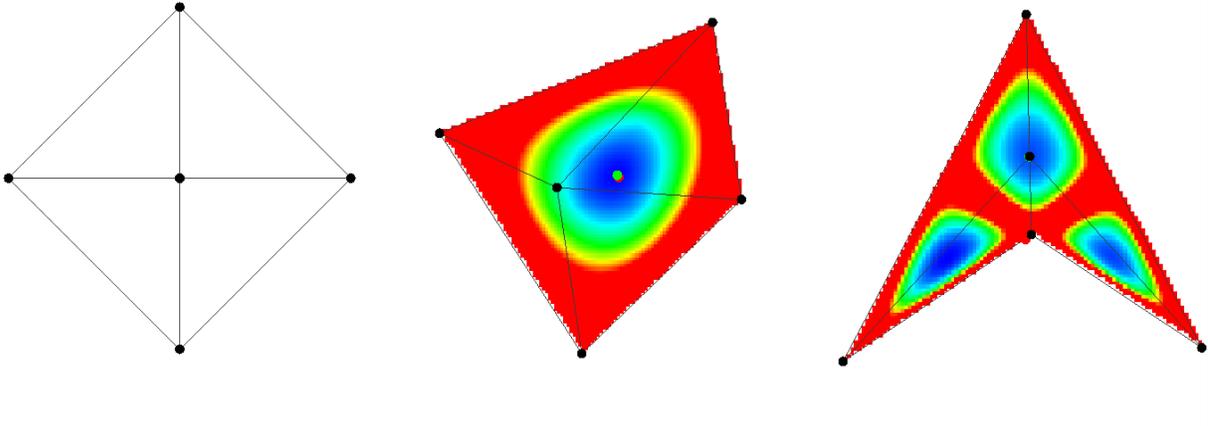


Figure 5.2.2: A one-ring neighborhood (left) is given two particular parameterizations (center, and right) and colored based on the sum of dilatations of the one-ring triangles. Note that the restriction to the subset of the one-ring interior where all triangle orientations are preserved (right) is still uniminimal.

mization problem at each vertex does not guarantee that we will arrive at a globally optimal solution, we can iterate until we converge to an overall energy minimum. The algorithm is as follows:

1. For each vertex v , do gradient descent to position the vertex to minimize local MIPS energy.
2. Repeat until the change in overall MIPS energy summed over all mesh facets is less than ϵ .

5.3 Results

5.3.1 Characteristic behavior of angle optimization on 2D grid

We start off with a toy example in 2D to explore the characteristic behavior of the angle relaxation methods. The initial condition is a 2D grid shown in 5.3.1. From this starting configuration, the map is relaxed by the respective methods. As we can see, the metric-based harmonic map optimizer distributes the edge-length errors, which is initially focused

on the right extreme of the grid, throughout the remainder of the mesh. Interestingly, the `OrbifoldBeltrami` method appears to be a more extreme version of the previous map, allowing for more drastic warping in order to reduce and focus the distribution of Beltrami coefficient norms.

5.3.2 Characteristic behavior of angle optimization on 3D surface

In this example, we extend the experiment from the previous section to the 3D case, seeing if the qualities observed in the 2D case carry over. We examine the behavior of the various angle relaxation methods, looking for qualitative differences in the output maps with two toy examples of surfaces that capture topologically what we expect to work with in real datasets (summarized in Figure 5.3.2).

1. **Pill** We start with a uniform randomly sampled round sphere that is partitioned into two hemispheres, and we select a total of six points, three in each hemisphere, which we treat as orbifold cone points so that the surface admits a hyperbolic metric. The sphere is mapped to a surface shaped like a capsule, which is formed by fixing one hemisphere of the round sphere while translating the free hemisphere away from the fixed hemisphere.
2. **Sprocket** In addition, we use a genus-2 surface and map it to a version of the same surface that has an elongated shaft portion. The higher genus allows us to compute a hyperbolic metric without the addition of any cone points or topological changes. The initial map is formed by taking the ends of the surface and projecting them to the corresponding ends on the target surface.

Starting from the initial configuration, we generate an optimized map by using the `OrbifoldPoints` and each angle optimization method. The relative distortion behavior of the maps is summarized in Table 5.1. Across both examples, the harmonic energy-based `OrbifoldPoints` method generates maps with lowest harmonic energy, as expected. For the pill example,

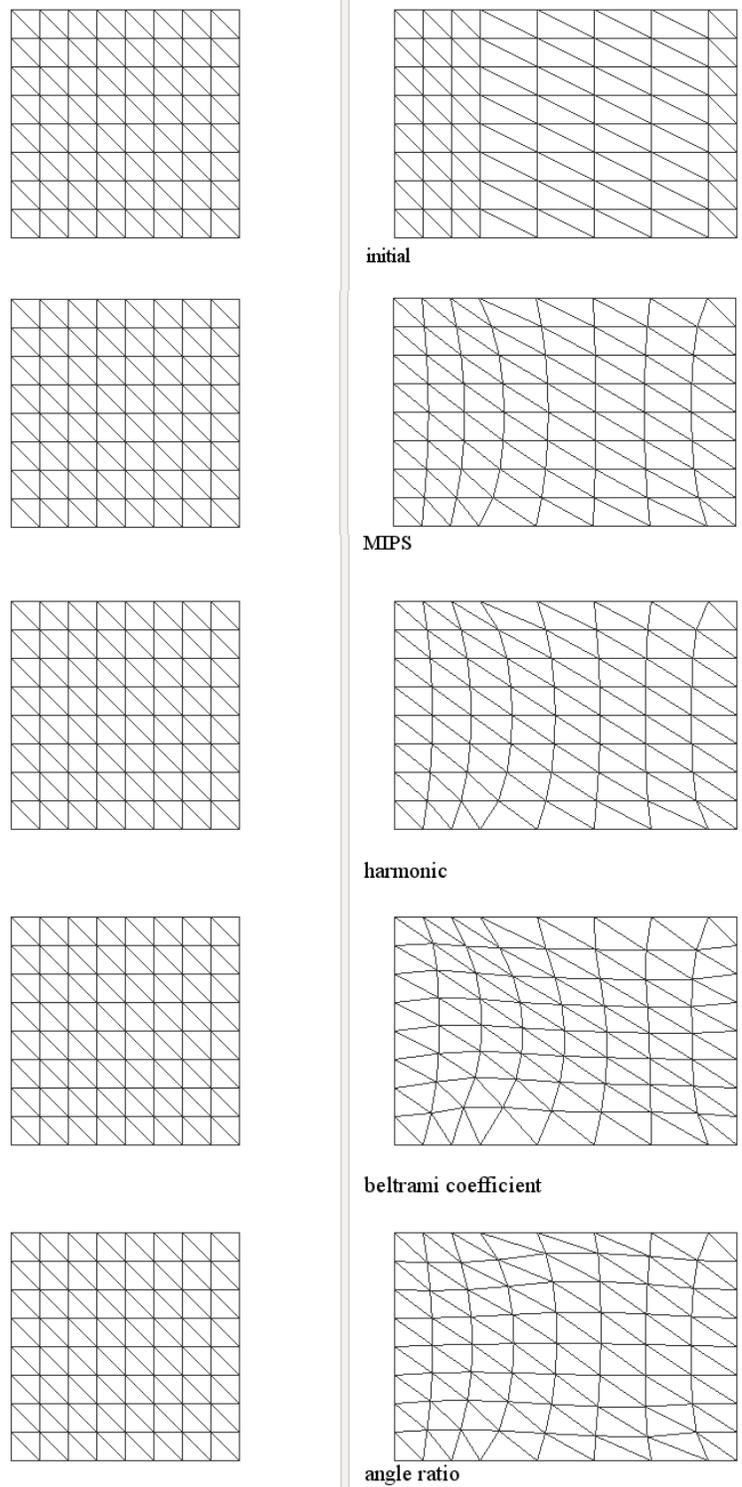
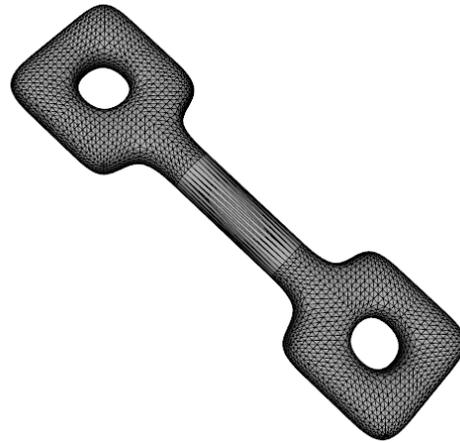


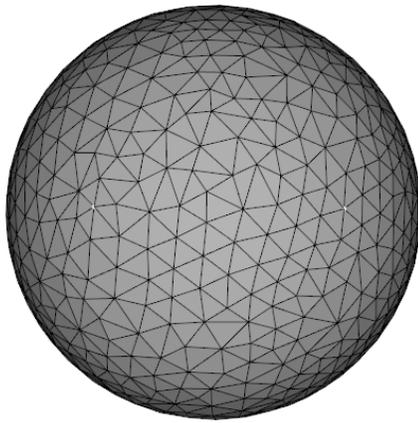
Figure 5.3.1: A comparison of the result of relaxing a 2D grid mesh where the initial condition is is the nonuniform stretch in the first row. From top to bottom are: (1) the initial configuration, (2) optimizing for MIPS, (3) harmonic energy, (4) Beltrami coefficient norm, and (5) angle market share.



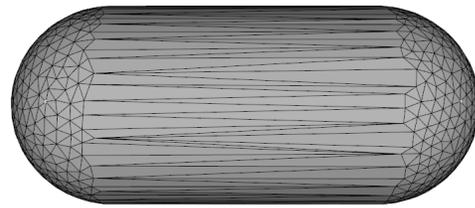
(a) A short sprocket.



(b) Mapped onto the target surface.



(c) A sphere with cone points selected as white dots.



(d) Mapped onto the target surface.

Figure 5.3.2: The source surface is mapped onto the target surface as shown, and then each relaxation algorithm is run on this initial configuration.

`OrbifoldAngleShare` produces lowest angle distortion across all metrics considered, while optimizing only for angle share. However, `OrbifoldAngleShare` does not perform as well in the sprocket example, possibly encountering a local minimum. For the sprocket example, dilatation and average Beltrami coefficient is lowest in the maps produced by `OrbifoldMIPS`. Overall, the results show that angle distortion can be and does get improved by the proposed method.

For this experiment, we observe whether the angle relaxation methods are able to arrive at a consistent solution under a different initialization. We use the sprocket mesh where the

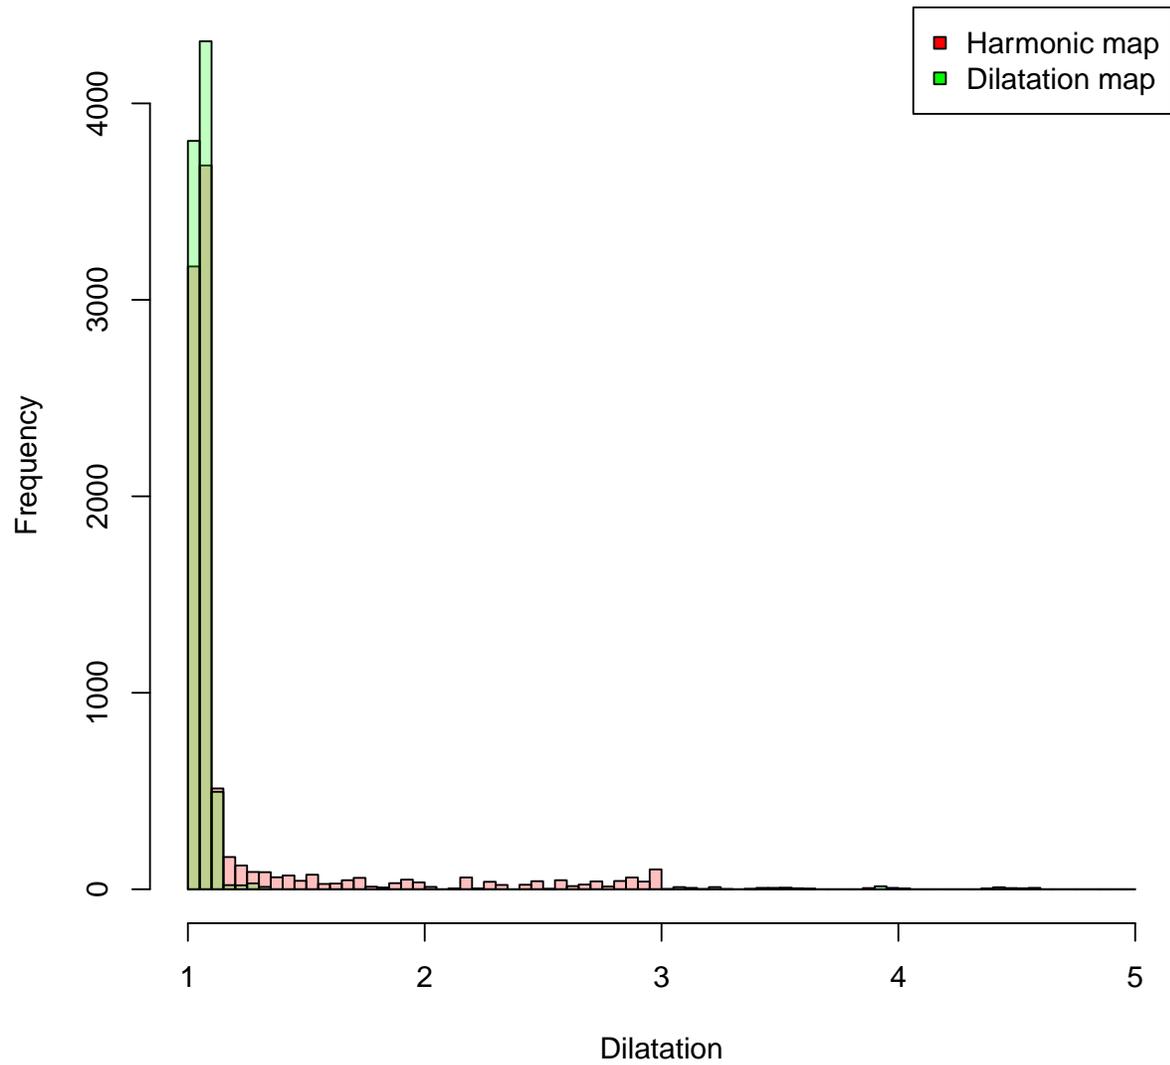


Figure 5.3.3: Histogram of dilatations for the sprocket example. The red bars belong to the harmonic map produced by `OrbifoldPoints`, while the green bars belong to the dilatation map produced by `OrbifoldMIPS`.

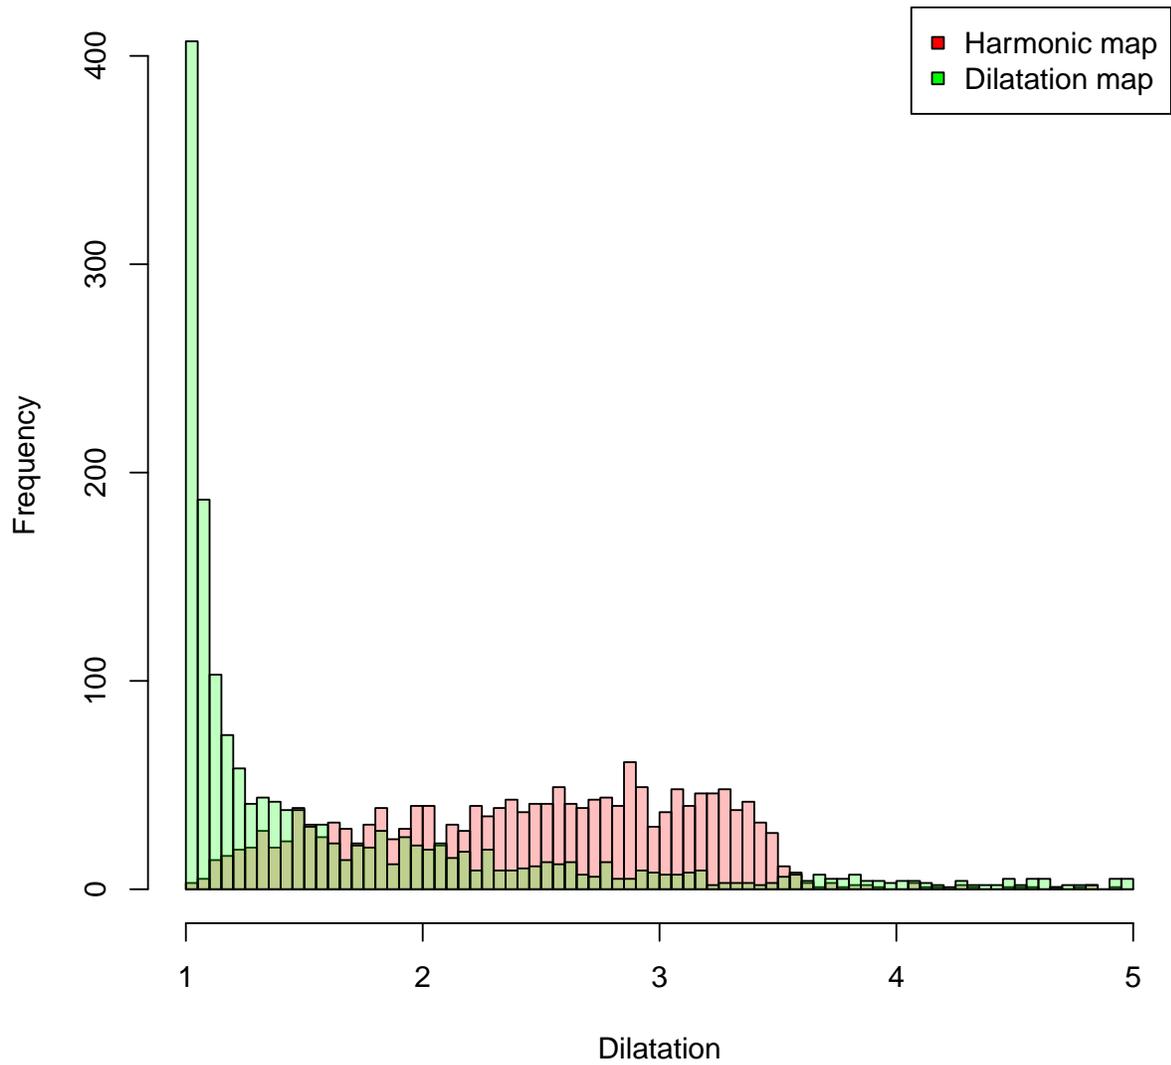


Figure 5.3.4: Histogram of dilations for the pill example. The red bars belong to the harmonic map produced by `OrbifoldPoints`, while the green bars belong to the dilatation map produced by `OrbifoldMIPS`.

	Dilatation	Beltrami coefficient (avg)	Angle share (avg)	Harmonic energy
Pill-h	5012	0.401	1.404	2.352
Pill-d	4524	0.378	1.387	3.238
Pill-b	5266	0.397	1.496	2.453
Pill-a	4479	0.298	1.381	2.926
sprocket-h	11550	0.114	1.036	3.386
sprocket-d	10916	0.061	1.041	3.946
sprocket-b	11247	0.096	1.035	3.406
sprocket-a	13085	0.148	1.069	3.545

Table 5.1: Table of distortions for the pill and sprocket toy example maps. The first column indicates the surface mapped and method used to produce the map (h - `OrbifoldPoints`, d - `OrbifoldMIPS`, b - `OrbifoldBeltrami`, a - `OrbifoldAngleShare`). Cells highlighted green indicate lowest distortion value for that group of maps.

initial map is twisted about the handles and the central shaft. The initial configuration and the respective results of each relaxation method are summarized in Figure 5.3.5.

5.3.3 Relative behavior of angle optimization across real-world datasets

For this experiment, we selected a representative subsample of the face and skull dataset, and additionally, a hand model. We then applied a structured deformation to form a test dataset. The examples differ in that protrusions in the mesh, such as noses and teeth, are made to be more extreme. We compute an initial map from the unaltered, original mesh to the deformed version of itself, then we perform relaxation in the hyperbolic parameter space separately for each method, optimizing for either harmonic energy or dilatation. Again, we explore the qualitative behavior of the angle optimization methods relative to each other and with respect to `OrbifoldPoints`. The following figures compare the resulting maps, and Table 5.2 quantifies the distortion behavior of these maps.

Overall, the angle optimizing methods have less dilatation than the harmonic map. `OrbifoldAngleShare`, `OrbifoldMIPS`, and `OrbifoldBeltrami` reduce the overall angle distortion across the surface with a modest tradeoff of harmonic energy. In particular, for the hand and face example, `OrbifoldBeltrami` exhibit the highest harmonic energy. Generally, `OrbifoldAngleShare` exhibits low values of angle distortion despite only explicitly

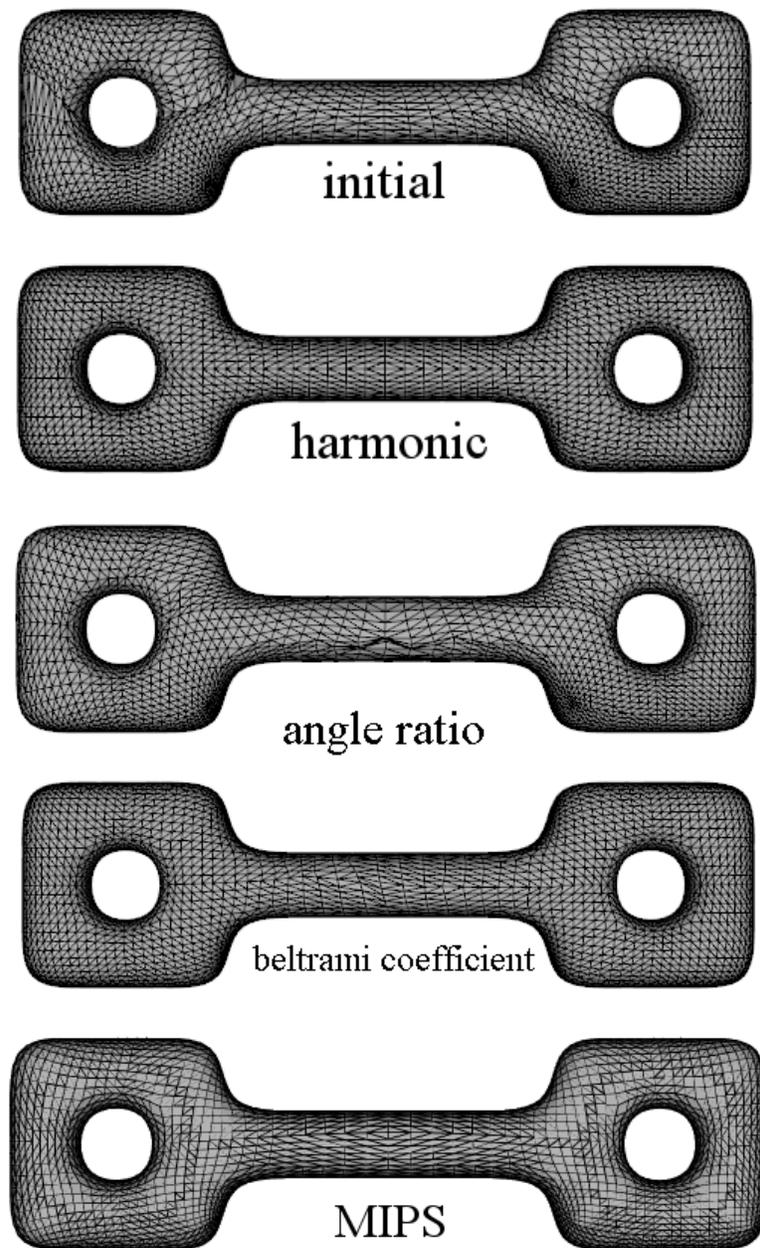


Figure 5.3.5: The initial map and results produced by the harmonic relaxation, angle ratio relaxation, quasiconformal relaxation, and dilatation relaxation, respectively.

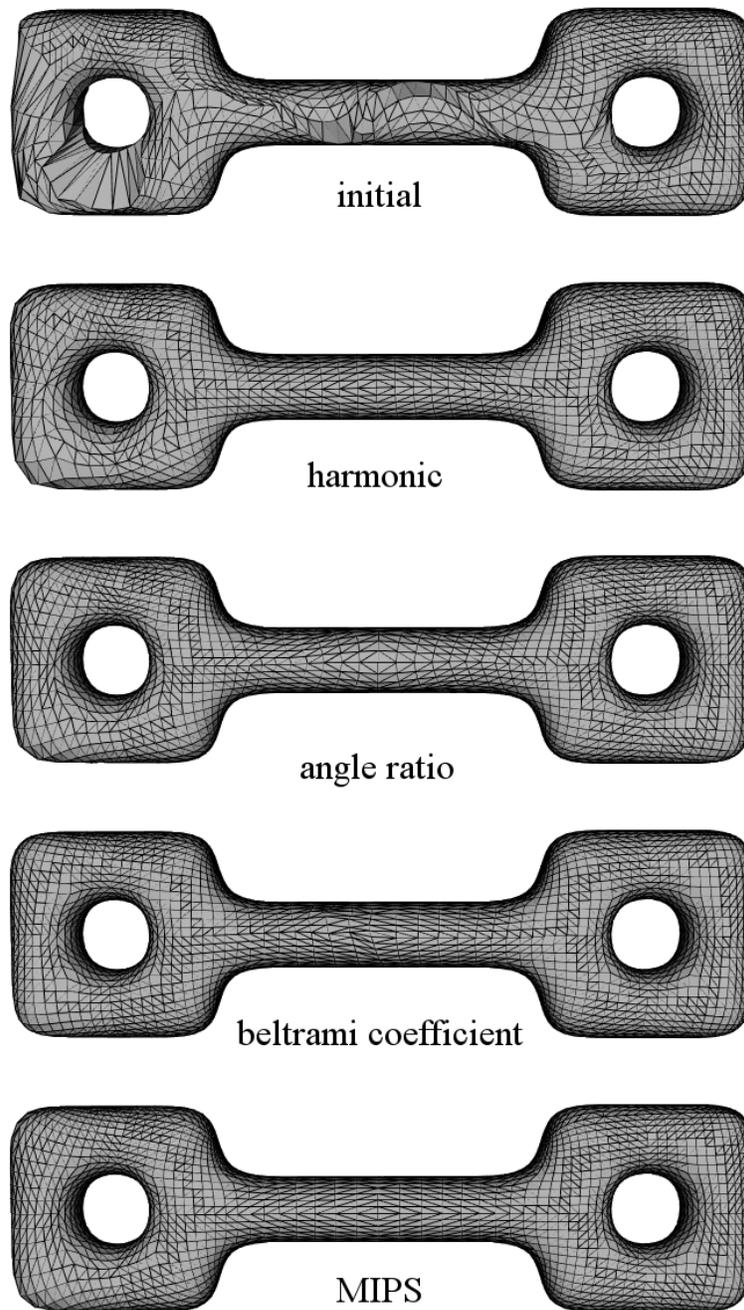


Figure 5.3.6: The initial map and results produced by the harmonic relaxation, angle ratio relaxation, quasiconformal relaxation, and dilatation relaxation, respectively.

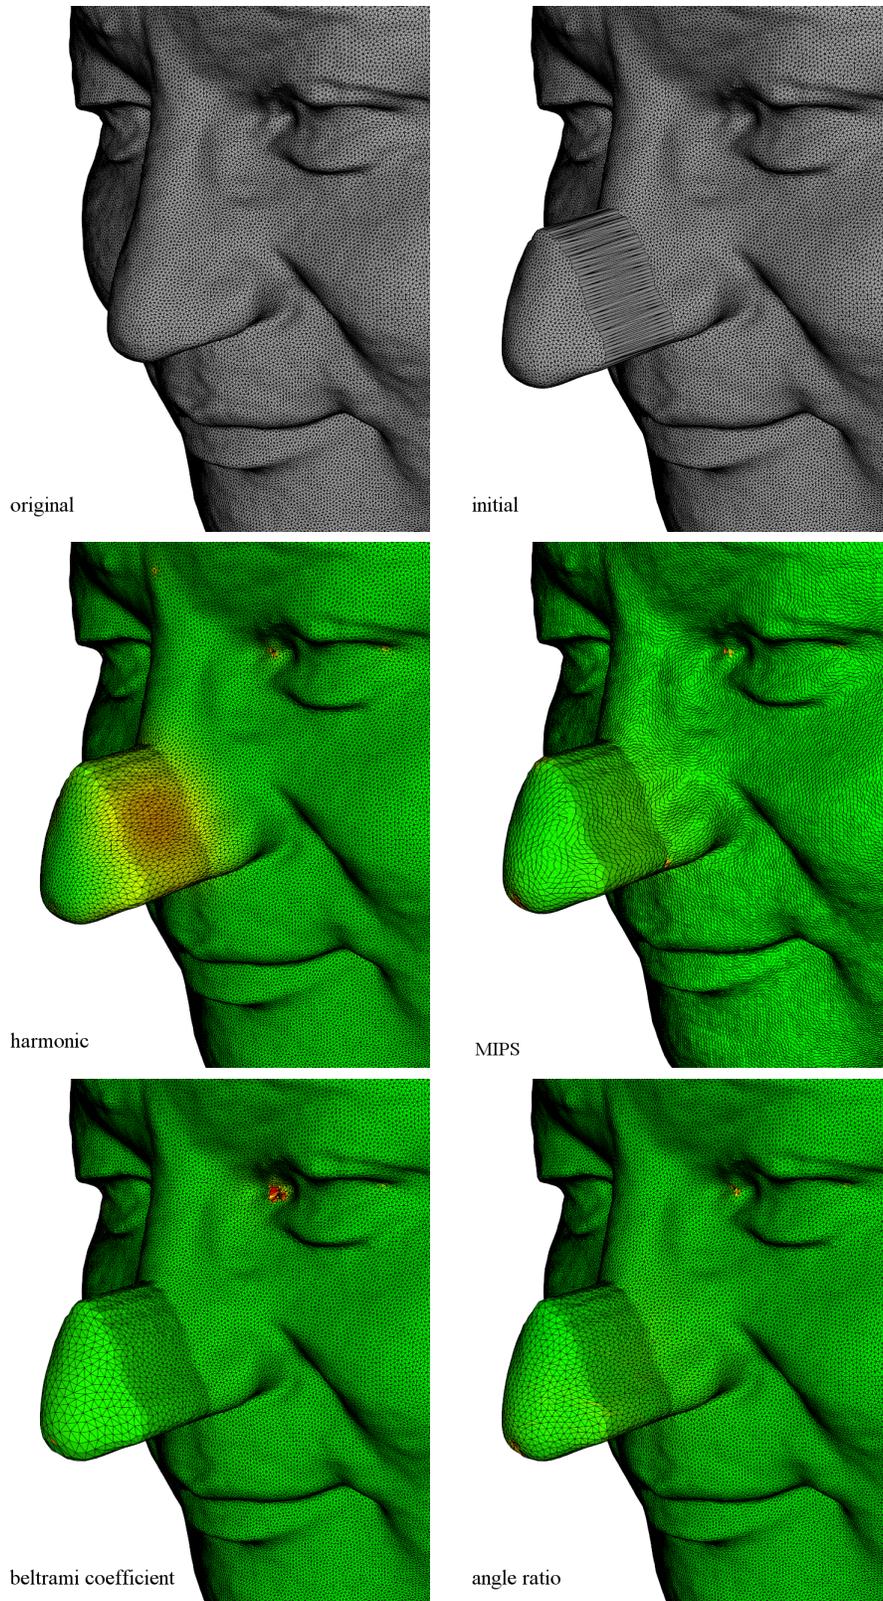


Figure 5.3.7: Face dilatation comparison. The original face mesh (upper left) is mapped onto a version of itself where the nose is extended (see upper right for initial correspondence). After relaxation, we compare the final correspondences produced by `OrbifoldPoints` (middle left), `OrbifoldMIPS` (middle right), `OrbifoldBeltrami` (lower left), and `OrbifoldAngleShare` (lower right), each colored by dilatation.

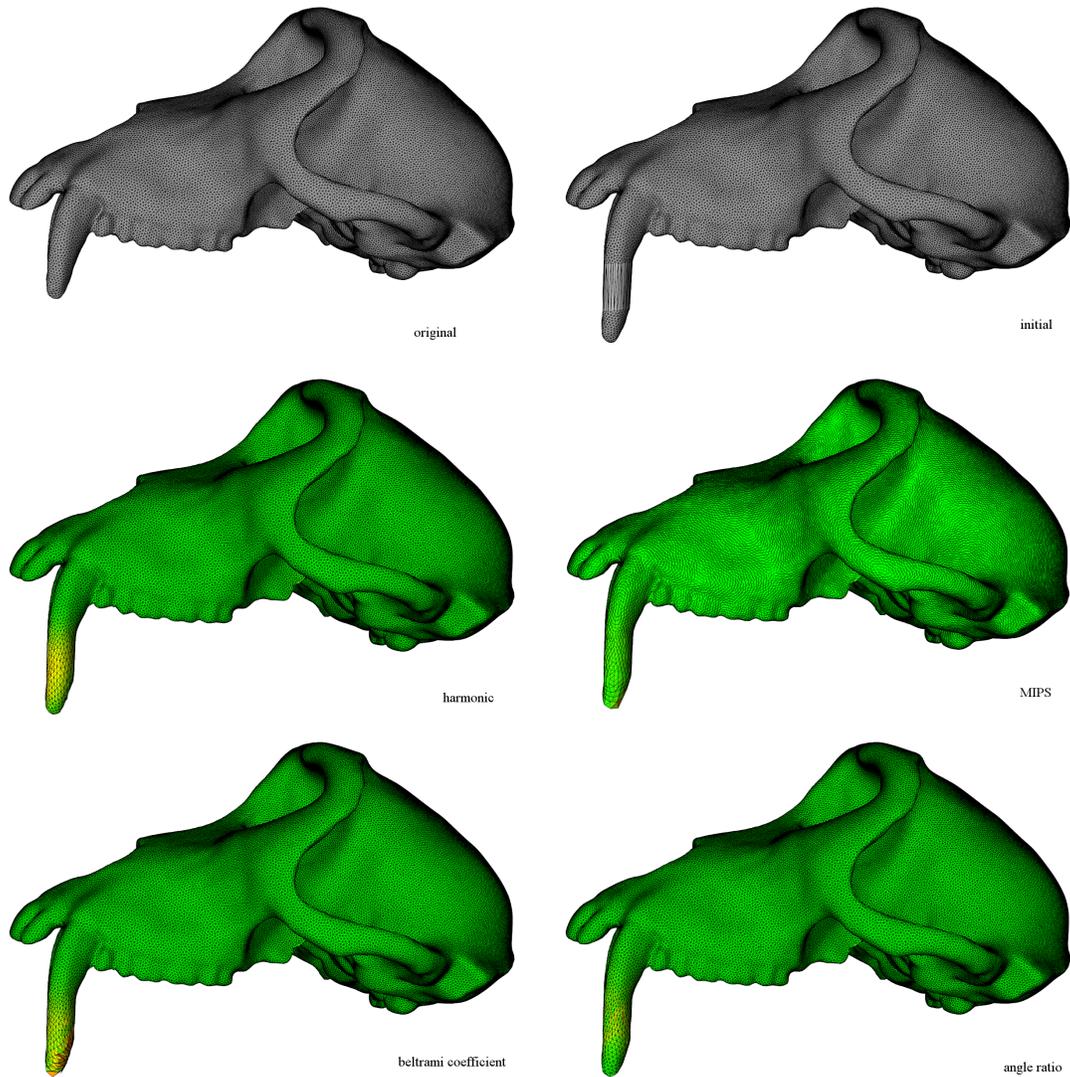


Figure 5.3.8: Macaque (*Macaca Pagensis*) dilatation comparison. The original macaque mesh (upper left) is mapped onto a version of itself where the canines are extended (see upper right for initial correspondence). After relaxation, we compare the final correspondences produced by `OrbifoldPoints` (middle left), `OrbifoldMIPS` (middle right), `OrbifoldBeltrami` (lower left), and `OrbifoldAngleShare` (lower right), each colored by dilatation.

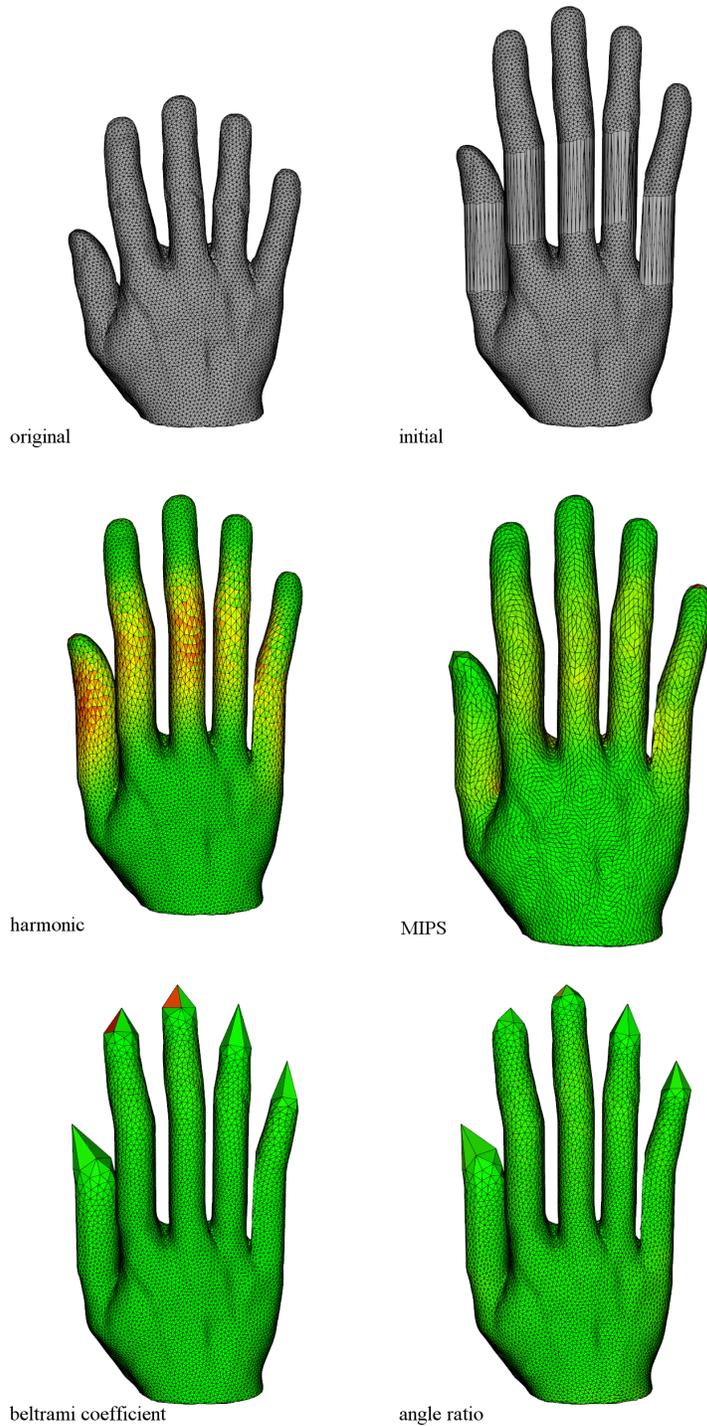


Figure 5.3.9: The original hand mesh (upper left) is mapped onto a version of itself where the fingers are extended (see upper right for initial correspondence). After relaxation, we compare the final correspondences produced by `OrbifoldPoints` (middle left), `OrbifoldMIPS` (middle right), `OrbifoldBeltrami` (lower left), and `OrbifoldAngleShare` (lower right), each colored by dilatation. Sparse sampling at the fingertips is an artifact of angle optimization: triangles grow arbitrarily near cone points, producing poor discretization. This artifact might be mitigated by conformal factor adjustment (see Section 5.3.5).

	Dilatation	Beltrami coefficient (avg)	Angle share (avg)	Harmonic energy
Skull-h	164637	0.0046	1.0025	10.4666
Skull-d	161489	0.0050	1.0006	10.4668
Skull-b	161251	0.0025	1.0025	10.4667
Skull-a	159981	0.0015	1.0001	10.4698
Face-h	140318	0.0265	1.0078	3.4621
Face-d	137724	0.0125	1.0030	3.4814
Face-b	144006	0.0218	1.0152	3.4985
Face-a	137292	0.0120	1.0020	3.4656
Hand-h	26704	0.1163	1.0497	11.0027
Hand-d	23893	0.0799	1.0173	11.4819
Hand-b	21809	0.0383	1.0746	11.7734
Hand-a	24275	0.0918	1.0148	11.1200

Table 5.2: Table of distortions for the real world example maps. The first column indicates the surface mapped and method used to produce the map (h - `OrbifoldPoints`, d - `OrbifoldMIPS`, b - `OrbifoldBeltrami`, a - `OrbifoldAngleShare`). The entries indicate the value of the distortion metric in the column for the map given in the row. Cells highlighted green indicate lowest distortion value for that group of maps.

optimizing angle share.

A notable defect is apparent in `OrbifoldBeltrami` near the landmark points, where the dilatation is concentrated. Fixing the landmark points hinders every method’s ability to minimize distortion, as landmark points are readily identified by points of high dilatation; however, `OrbifoldBeltrami` is impacted much more so than the other methods. A reason why this may be occurring is that the method is trying to push the mapping towards one where the Beltrami coefficient norm is constant across the entire map. This may be compounded with the fact that there may be inconsistencies in the transition between the orbifold parameterizations at the landmark points and normal manifold parameterizations at the points directly adjacent to the landmark points. Indeed, the Beltrami optimization scheme has only ever been used in 2D Euclidean and Euclidean manifold settings, so perhaps `OrbifoldBeltrami` should not be used on orbifolds with singular points.

The comparative metrics from the toy data in Table 5.1 favor no method in particular. However, when we consider typical data in Table 5.2, `OrbifoldAngleShare` demonstrates more consistency in performance across distortion metrics. Although the experiments thus

far show evidence that all three angle optimizing methods are susceptible to local minima dependent on the data, `OrbifoldAngleShare` appears to give overall improvement across all angle distortion metrics.

5.3.4 `OrbifoldAngleShare` is robust to length-only deformations

Consider the result is shown in Figure 5.3.10. To follow up on the experiment in the previous section, we investigate whether we can utilize an angle-based measure of distortion to gain some novel insight from how harmonic distortion informs us about the shape differences. We start with a hand mesh and create a copy where the fingers are deformed in such a way as to lengthen the fingers in one direction. With these two base meshes, we create intermediate hand meshes of varying finger length by linear interpolation. Then, we compute optimal maps with `OrbifoldAngleShare` and `OrbifoldPoints` from the original hand to each interpolated hand $i, i = 1..8$ and observe how the harmonic and angle distortions cluster, with the intuition being that lower distortions suggest that the original and target surfaces are closer with respect to the distortion metric. There is a clear correlation when looking at harmonic distortion, which is expected from an example that introduces a linear distortion in one direction, whereas there is less of a correlation with angle distortion. This is indicative of the method finding a local minima, which suggests that regardless of the elongating deformations of increasing magnitude being applied to the mesh, the energy landscape remains roughly the same. As such, we posit that angle relaxation may be robust to elongating deformations of the mesh.

5.3.5 Discussion

From the gallery of results, we see that the maps produced by the harmonic relaxation in `OrbifoldPoints` versus angle relaxation methods exhibit qualitative differences that are subtle in some cases and systematic in others. It can be suggested that angle distortion and harmonic distortion minimizing methods prioritize different areas of the mesh relative to

the introduced deformation. For example, in Figure 5.3.9, `OrbifoldPoints` spreads out the length distortion across the deformed canine, with the tip of the tooth and the base remaining relatively isometric, whereas `OrbifoldMIPS` concentrates distortion at the extremity of the canine. Without incorporating a term for length distortion into the energy function, this effect can be moderated to some extent by adjusting the conformal factors (see next section) as in Figure 5.3.13, but it is still present.

While the type of synthetic deformation introduced into each data sample is a directional stretch, the result for angle relaxation can vary from general shape smoothing, as in the skulls (Figures 5.3.8), to stratified layers of distortion as in the hands (Figure ??). Furthermore, as Figure 5.3.10 shows, angle optimization does not produce as steady of a progression of distortion as compared to harmonic optimization, given a steadily deforming sequence of shape. This may seem to indicate that the angle distortion-minimizing methods encounters local minima in both 2D and 3D settings. However, across all maps, the dilatations are uniformly lower in the angle-optimizing maps compared to the harmonic maps, suggesting that the methods may still be of use for optimizing maps that are close to an ideal configuration.

Conformal factor adjustment

It is well-known that the performance of conformal methods are heavily influenced by parameterization. As protruding regions on the mesh often exhibit large conformal factors, translating to extreme length distortion in order to achieve a conformal map into the parameter space. As a consequence, the iterative relaxation that attempts to dissipate distortion will be more heavily biased in regions with either high compression or expansion. The results of relaxation can thus be exaggerated, whether minimizing length or angle distortion as in Figures 5.3.11 and 5.3.12. While one of the main points of a method based on uniformization is in its ability to provide a unique parameterization, it is often valuable to consider measures to engineer a more uniform parameterization. In the literature, holomorphic one-forms [53] have been used to generate as uniform as possible parameter space coordinates from a

class of possible parameterizations. It is also suggested that additional point singularities be inserted in order to diffuse the conformal factors [57]. Indeed, this opens up a problem how to automatically optimize the conformal factors and whether there exists an optimal set of factors. On the other hand, there are approaches e.g. the conformal Wasserstein metric [37] [61] that utilize a surface’s “intrinsic” conformal factors as a shape signature, and artificial “uniformization” of the uniformization metric may be counterproductive in such cases. Therefore, the utility of conformal factor modification depends on the particular application.

In conclusion, we have implemented three alternate conceptions of angle distortion and found that their behavior produces variable results, particularly in relaxing 3D surface maps. Indeed, work in the literature for these methods have been limited to 2D surface meshes or Euclidean 2D parameterizations of surfaces, and this experiment shows empirically that there is a robustness issue, at least with the hyperbolic parameterization. While it is not guaranteed that any of the method variants will produce a unique solution as is theoretically the case in `OrbifoldPoints`, we see from the experiments on synthetic data that they successfully lower angle distortion and can provide insight that is complementary to harmonic distortion alone, potentially contributing to biologically relevant results. However, if we consider relative simplicity, optimality, and robustness under a variety of initial conditions, we conclude that a method that is based on harmonic relaxation such as `OrbifoldPoints` is best in practice.

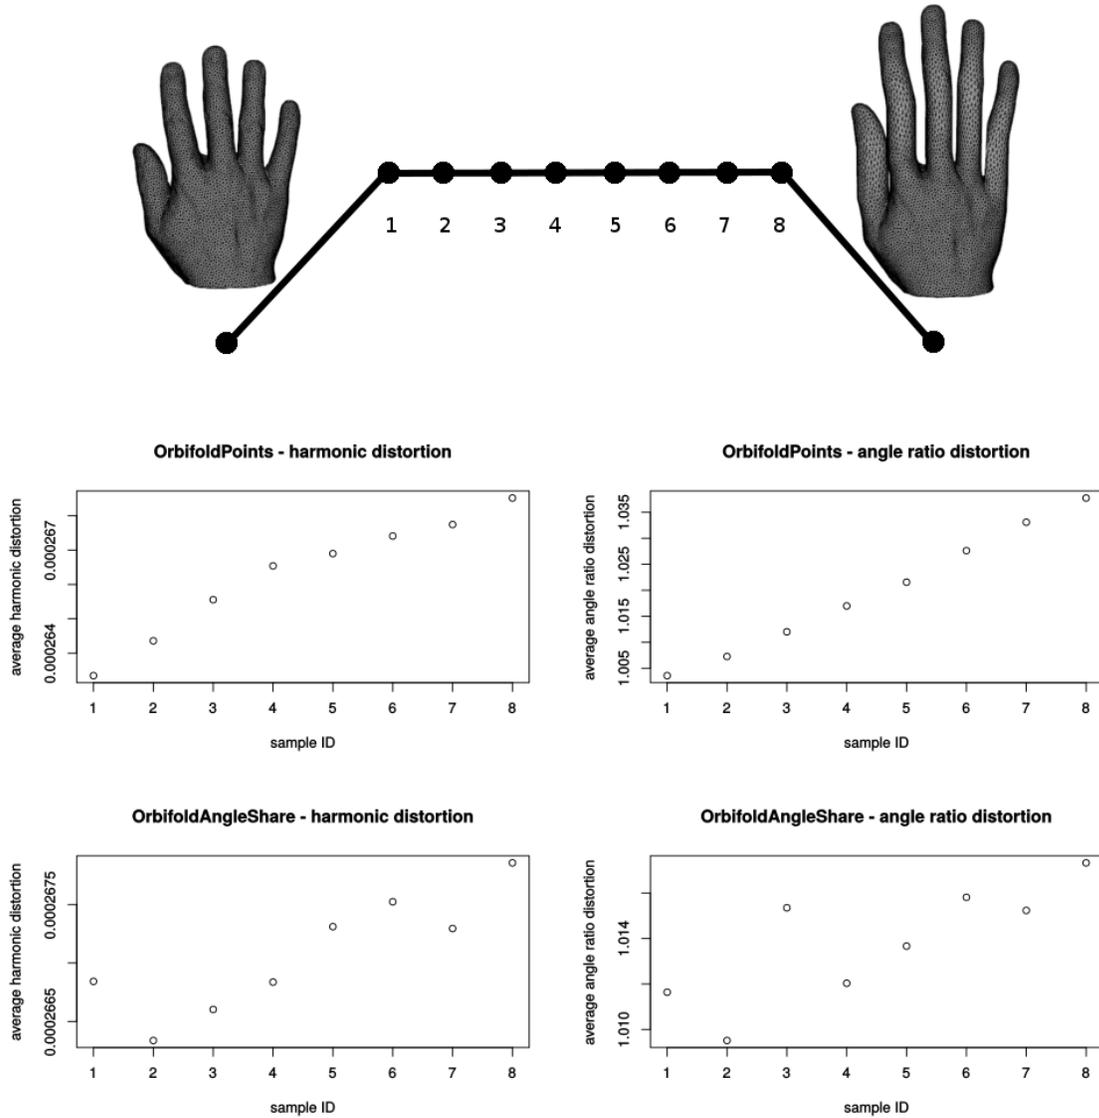


Figure 5.3.10: Comparing sensitivity of length and angle distortion in **OrbifoldPoints** and in **OrbifoldAngleShare**. Note that the range of angle ratio distortion is twice as large in **OrbifoldPoints** than in **OrbifoldAngleShare**. We first interpolate the original hand with the hand with long fingers, which introduces length distortion in one direction. The harmonic and angle distortion for the sequence of maps generated by **OrbifoldPoints** intuitively correlates with each interpolated mesh. The **OrbifoldAngleShare** map distortion is less well ordered, which suggests (1) intrinsically, the interpolated hands are somehow still similar with respect to angle, and (2) the method is susceptible to local minima.

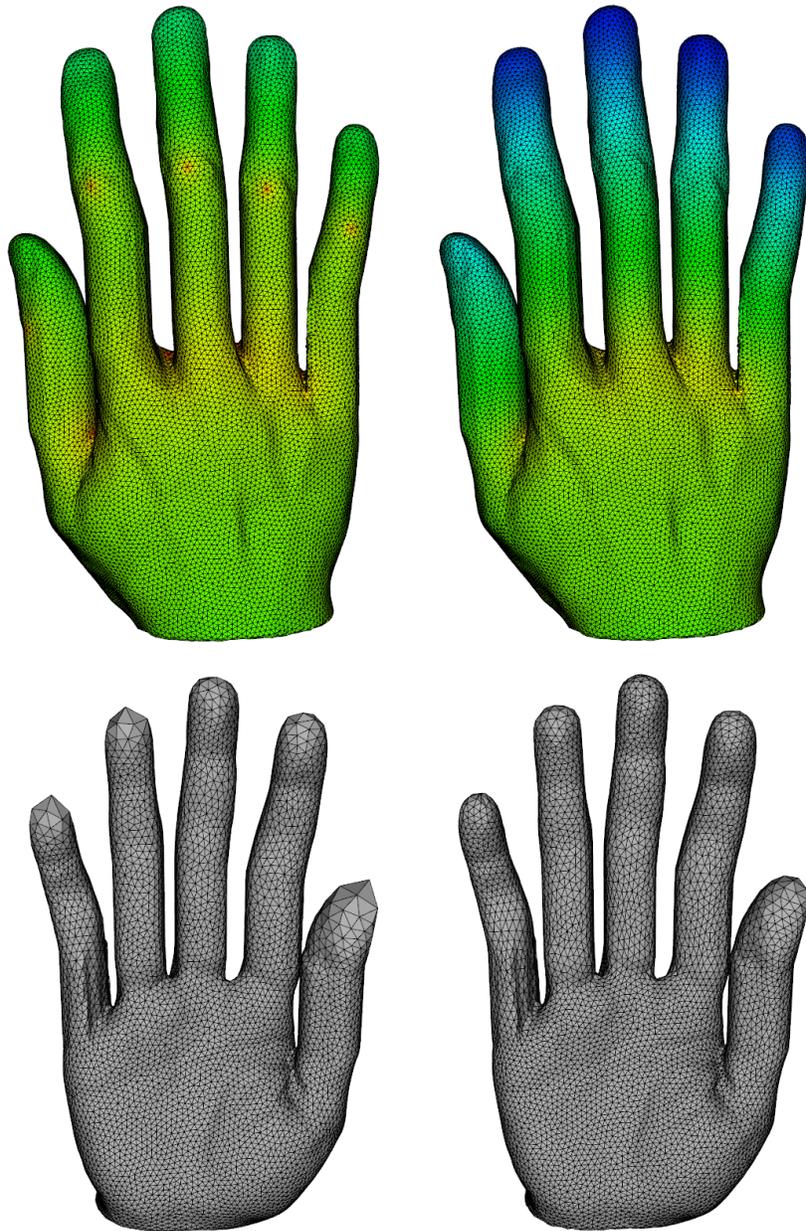


Figure 5.3.11: Effect of uniformity of conformal factors on angle relaxation. In the top left, the cone points defining the orbifold metric are at the fingertips and at the wrist, resulting in a metric that has small conformal factors at the fingertips. The result of angle optimization using a parameterization based on these conformal factors is in the lower left. Notice that there is loss of detail at the fingers where the conformal factors are more extreme. In the top right, additional cone points along the finger (visible as orange shaded vertices) allow the factors to be more uniform overall over the surface. In the bottom right, the effect of this is a more balanced placement of vertices.

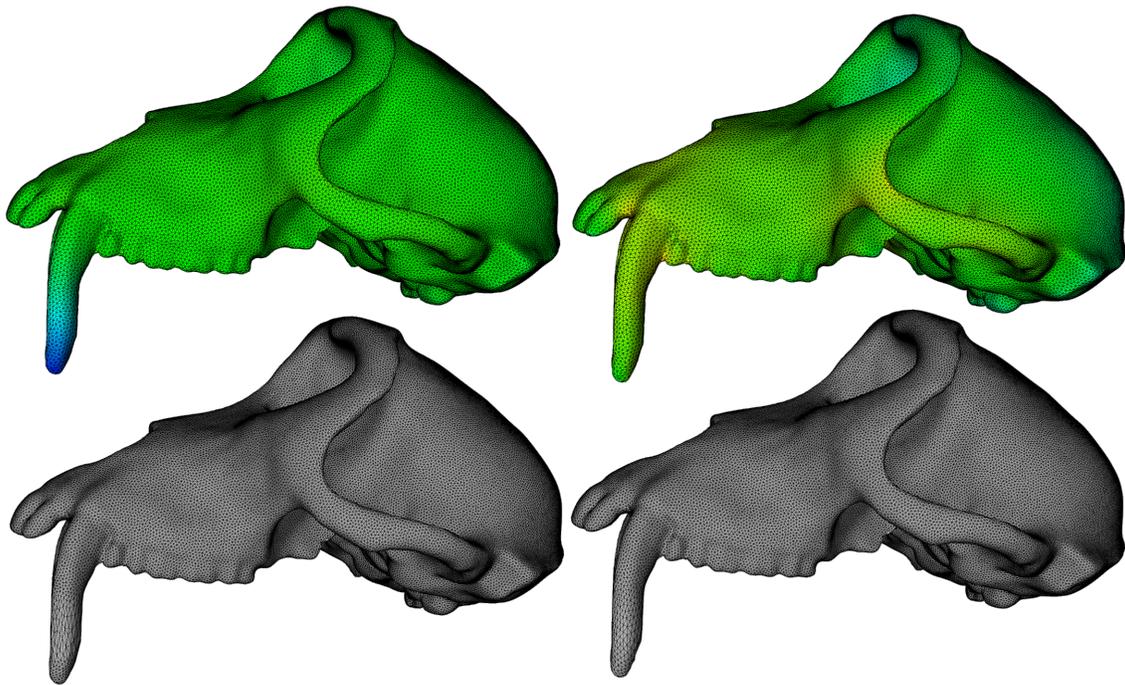


Figure 5.3.12: Uniformity of conformal factors on harmonic relaxation. The top skulls show different conformal factors representing hyperbolic orbifold metrics defined by different sets of cone points, with the right one being clearly more uniform. The bottom skulls show the resulting harmonic map onto the skull above. Notice the exaggerated stretch on bottom left in the canine is much less pronounced on the right.

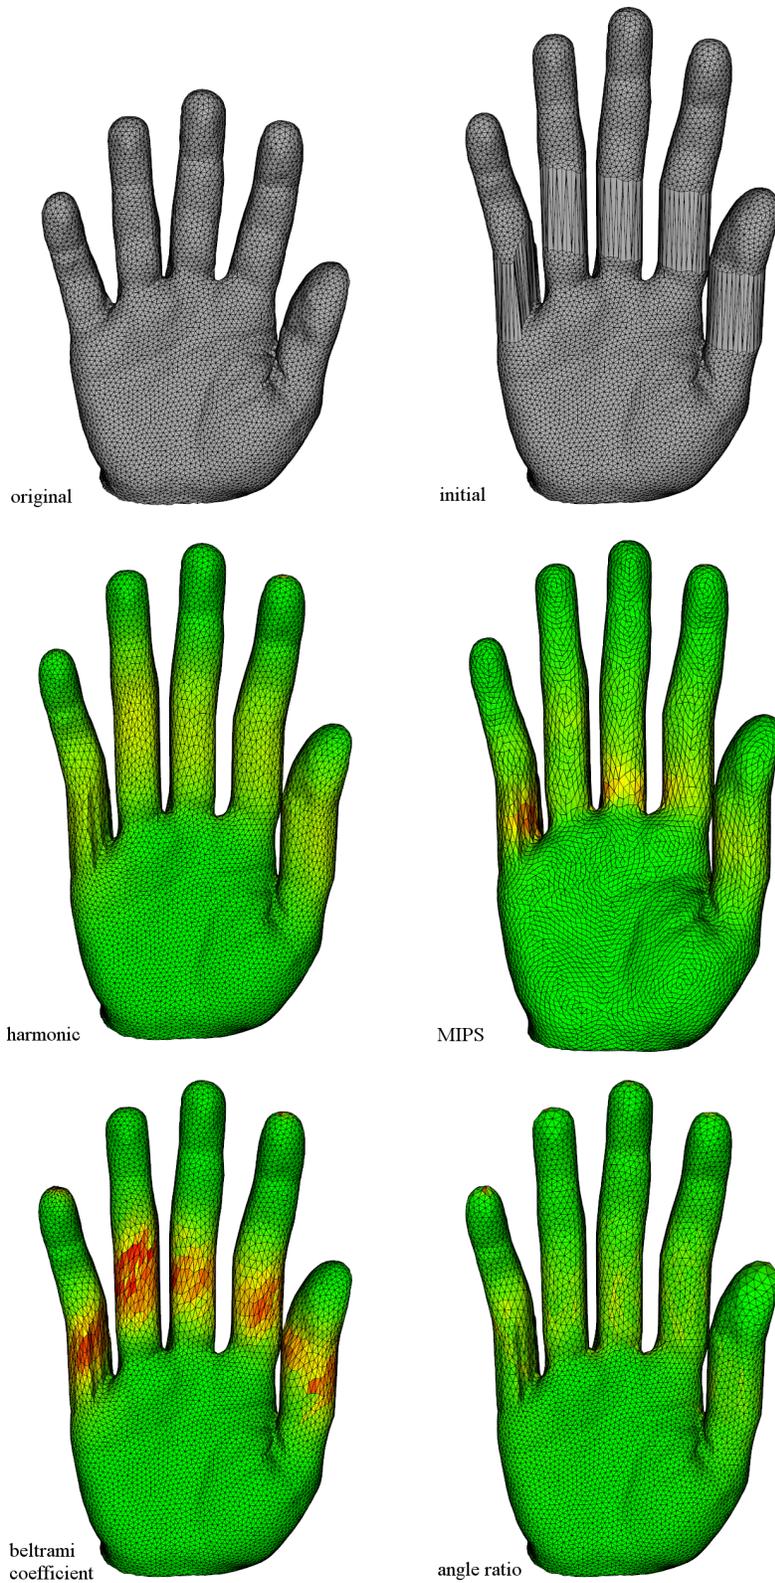


Figure 5.3.13: Hand dilatation comparison after conformal factor adjustment (see Section 5.3.5). The original hand mesh (upper left) is mapped onto a version of itself where the fingers are extended (see upper right for initial correspondence). After relaxation, we compare the final correspondences between the harmonic map (lower left) and dilatation map (lower right) colored by dilatation.

Appendix A

Technical Overview Of Mapping

Software

In this chapter, we will explain the details of the software implementation of the surface mapping algorithm described in the previous chapters. The software itself, in addition to detailed documentation, will be distributed from <https://github.com/geomlab-ucd>, where the source code can be accessed and compiled. The software is organized into three parts:

1. A library defining the common data structures and algorithms supporting the surface mapping procedures.
2. A set of command-line tools that carry out the individual steps of the surface mapping algorithm.
3. A set of graphical examples for visualizing results of the mapping algorithm and for inspecting different aspects of a surface map produced by the algorithm.

The purpose of this chapter is to present the layout of the software project to the end-user. The next purpose is to explain the architecture and design of the various components in the project so that a potential developer might be able to maintain and perform maintenance of the existing codebase, identify and implement optimizations, or extend the library with new

functionality.

The surface mapping library is a C++ project that is set up with a CMake build system. At the time of this writing, the project was developed and built on a Linux workstation running the Ubuntu 14.04 operating system. The project depends on a number of third-party libraries for certain functionality

1. Computational Geometry Algorithms Library (CGAL) - At the time of this writing, we use CGAL 4.x. The project provides a stock implementation of some common utility algorithms in computational geometry. In particular, we make use of the halfedge-based polyhedral data structure to represent a triangulated surface. We also use the 2D parameterization package to compute Euclidean harmonic maps of topological disks.
2. Qt - At the time of this writing, Qt version 5 is available.
3. Visualization Toolkit (VTK) - We use VTK 5.10.1. This library provides geometry processing to a limited extent, but primarily provides widgets for 3D surface visualization and user interaction.
4. Cairo - This library provides 2D drawing.
5. Googletest - This library provides unit testing.
6. Boost - We use the version of Boost bundled with Ubuntu 14.04. This is a library that provides C++ extensions and utilities that are not available in the C++ standard template library, such as multithreading support, tuples, program options, and cross-platform filesystem handling support.

A.1 Prerequisites

A.2 Conventions

The surface is represented as a halfedge-based polyhedron data structure. Consider a surface triangulation is determined by a set of vertices and facets describing locally connected subsets of these vertices. The geometry revolves around the edges. Each edge corresponds to two halfedges. The halfedges are oriented in such a way that they point in the counterclockwise direction around any facet. A halfedge has a reference to the next halfedge in the counterclockwise walk around the facet to which it belongs, or if it does not belong to a facet (i.e. on the boundary of a surface), then it is the next halfedge along the boundary. A halfedge is incident to the vertex that it points to and to the facet that it circulates.

A.3 Usage

`computeMetric` computes the hyperbolic metric. We use an isotropic remesher to uniformly and densely resample a mesh. We use our graphical user interface to select cone points and trace paths connecting the cone points into a tree. `generateInitialMap` takes as input two meshes with corresponding marking tree and generates an initial map by overlaying the harmonic map of each surface.

Appendix B

List of Mapping Software

There is a plethora of surface mapping techniques in the literature, and fortunately, a number of them also have publically available software implementations provided to the research community. Surface mapping has many applications and, indeed, not every surface mapping method that is developed is tailored to be generally applicable to all surface types or be interchangeable with another method for some specific application. Nevertheless, knowledge of available mapping software, whether they be commercial or open source, is useful, and we give a cursory overview of some available software at the time of this writing in this appendix.

Freesurfer

<https://surfer.nmr.mgh.harvard.edu/fswiki/DownloadAndInstall> An end-to-end medical image analysis software package tailored for processing brain MR images, it includes a set of command-line tools that can operate on topologically-spherical surfaces to map them to a sphere in a metric-minimizing fashion and then onto each other.

Spherical Demons

<https://sites.google.com/site/yeoyeo02/software/sphericaldemonsrelease> Spherical Demons performs fast diffeomorphic landmark-free surface registration of 2D closed surfaces of spherical topology, as described in [75]. Generally, point-based features are used to drive the alignment, and in the case of cortical surface alignment, this can include such measures as mean curvature and sulcal depth.

BrainVoyager

<http://www.brainvoyager.com/RainerGoebel.html> (Commercial) Another brain MR image analysis software suite that, in addition to data management, statistical analysis, and visualization, also provides facilities to perform spherical parameterization and flattening.

BrainVISA

<http://brainvisa.info/web/index.html> Another medical image analysis software package with an emphasis on head MRI, it includes a module where a surface can be parameterized onto the plane and landmarks can be used to bring surfaces into closer alignment.

Meshlab

<https://sourceforge.net/projects/meshlab/> Meshlab is a mesh and general 3D data visualization tool that also includes many built-in geometric algorithms, including a module to perform rigid registration. The graphical user interface allows the user to specify a rough initial alignment by corresponding points followed by an iterative rigid alignment.

Spherical harmonic map

<http://www3.cs.stonybrook.edu/~manifold/book/> An implementation from Gu that approximates the conformal map to a unit sphere by computing a harmonic map. General

spherical surface mapping may be achieved by using this software as an initial step, reducing the problem to spherical alignment as in [41].

Blended intrinsic map

<http://vovakim.com/projects/CorrsBlended/> Software implementation of blended intrinsic maps method proposed by Kim, Lipman, and Funkhouser [29]. The software provides a pipeline for mapping surfaces that are topologically spherical by blending candidate, nearly-isometric maps into an optimal map.

CirclePack

<http://www.math.utk.edu/~kens/CirclePack/> A Java-based implementation of the method to compute the circle packing metric proposed by Hurdal and Stephenson [25]. This serves mainly as a visualization of the planar parameterization of surfaces by circle packings, which may form the basis of a general mapping pipeline. Sample packings in the sphere and the Euclidean plane are also made available.

CGAL - Computational Geometry Algorithms Library

<http://cgal.org/download/> While the goal of this C++ library is to be a general computational geometry library and does not provide an end-to-end surface mapping solution out of the box, the subpackages provide most of the tools required to create the necessary pipeline as we have in our own mapping framework. CGAL provides implementations to a variety of well-known 2D planar parameterizations to the unit disk and square. It also provides low-level mesh data structures and mesh topology-modifying operators.

Schwarz-Christoffel Toolbox for MATLAB

<http://www.math.udel.edu/~driscoll/SC/index.html> The Schwarz-Christoffel mapping is a conformal transformation of the upper half plane onto the interior of a simple polygon. Polygons are homeomorphic to a disk and, as a corollary to the Riemann mapping theorem, can be conformally mapped to the unit disk, which can subsequently be identified to the upper half plane by the Möbius transformation $w = (z - 1)/(z + 1)$. This toolbox provides a MATLAB toolbox that allows for constructing maps between arbitrary planar polygonal regions and “canonical” regions such as a disk, square, half space, or more.

IDAV Landmark Editor

<http://www.idav.ucdavis.edu/research/EvoMorph> Developed in the work by Wiley et al. [72] to assist in annotating archaeological models with dense landmark sets, the software is capable of utilizing those landmarks to define thin-plate spline warps to align surfaces.

Stratovan Checkpoint

(Commercial) <https://www.stratovan.com/products/checkpoint> Stratovan is a commercial venture that was spun off from the work by Wiley, Ghosh, and colleagues with Landmark to create software solutions for 3D geometry-based archaeological morphometry and biomedical analytics. Some mesh manipulation features are shared with Landmark Editor, and surfaces can be aligned for general Procrustes analysis; however, the emphasis of the supported features extends far beyond the scope of surface mapping, providing surface extraction capability for domain-specific file formats, landmark collection, management, and statistical analysis tools, and so on.

Isotropic Surface Remeshing

<http://www.gis.informatik.tu-darmstadt.de/~sfuhrman/remesher.html> This is not mapping software but rather remeshing software implementing techniques described in [19] that proves to be useful in our conformal factors-based approach.

Chapter 6

References

- [1] L. Ahlfors and L. Sario. Riemann surfaces. 1960.
- [2] N. Aigerman and Y. Lipman. Orbifold tutte embeddings. *ACM Transactions on Graphics*, 34:90.1–90.12, 2015.
- [3] N. Aigerman, R. Poranne, and Y. Lipman. Lifted bijections for low distortion surface mappings. *ACM Transactions on Graphics*, 33(4):1–12, July 2014.
- [4] Gunnar Aronsson, Michael G Crandall, and Petri Juutinen. A Tour of the Theory of Absolutely Minimizing Functions. *Bulletin of the American Mathematical Society*, 41(4):439–505, 2004.
- [5] G. Auzias, J. Lefèvre, A. Le Troter, C. Fischer, M. Perrot, J. Régis, and O. Coulon. Model-driven harmonic parameterization of the cortical surface. MICCAI, pages 310–317, 2011.
- [6] F Bernardini, J Mittleman, H Rushmeier, C Silva, and G Taubin. The ball-pivoting algorithm for surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics*, 1999.

- [7] L. Bers. Uniformization, moduli, and kleinian groups. *Bull. London Math. Soc.*, 4:257–300, 1972.
- [8] K.R. Beutner, G. Prasad, E. Fletcher, C. DeCarli, and O.T. Carmichael. Estimating uncertainty in brain region delineations. In *IPMI*, volume 5636 of *Lecture Notes in Computer Science*, pages 479–490. Springer, 2009.
- [9] A. Bobenko, U. Pinkall, and B. Springborn. Discrete conformal maps and ideal hyperbolic polyhedra. pages 1–49, 2010.
- [10] K. Brodmann. *Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues*. Barth, Leipzig, Germany, 1909.
- [11] A.M. Bronstein, M.M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. 103:1168–1172, 2006.
- [12] B. Chow and F. Luo. Combinatorial ricci flows on surfaces. *J. Differential Geom.*, pages 97–129, 2003.
- [13] S. Dettle, A. Beiser, C. DeCarli, R. Au, J.J. Himali, M. Kelly-Hayes, J.R. Romero, C.S. Kase, P.A. Wolf, and S. Seshadri. Association of MRI Markers of Vascular Brain Injury With Incident Stroke, Mild Cognitive Impairment, Dementia, and Mortality The Framingham Offspring Study. *Stroke*, 41(4):600–606, 2010.
- [14] M. P. Do Carmo. *Differential geometry of curves and surfaces*. Prentice-Hall, 1976.
- [15] J. Eells and J.H. Sampson. Harmonic mappings of riemannian manifolds. *American journal of mathematics*, 86:109–160, 1964.
- [16] I Fáy. On straight line representation of planar graphs. *Acta. Sci. Math.*, pages 229 – 233, 1948.

- [17] E. Fletcher, B. Singh, D. Harvey, O. Carmichael, and C. DeCarli. Adaptive image segmentation for robust measurement of longitudinal brain tissue change. EMBC, pages 5319–5322. IEEE, 2012.
- [18] Face Recognition Grand Challenge. <http://face.nist.gov/frgc>.
- [19] S. Fuhrmann, J. Ackermann, T. Kalbe, and M. Goesele. Direct resampling for isotropic surface remeshing. In *VMV*, pages 9–16, 2010.
- [20] J. Glaunes, A. Qiu, M.I. Miller, and L. Younes. Large deformation diffeomorphic metric curve mapping. *Int'l journal of computer vision*, 80(3):317–336, 2008.
- [21] R.S. Hamilton. The ricci flow on surfaces, mathematics and general relativity. *Contemp. Math*, 71:237–262, 1988.
- [22] P. Hartman. On homotopic harmonic maps. *Canadian journal of mathematics*, 19:673–687, 1967.
- [23] J. Hass and P. Koehl. A metric for genus-zero surfaces. pages 1–33, 2015.
- [24] K Hormann. *Theory and Applications of Parameterizing Triangulations*. PhD thesis, Department of Computer Science, University of Erlangen, 2001.
- [25] M.K. Hurdal and K. Stephenson. Discrete conformal methods for cortical brain flattening. *NeuroImage*, 45:S86–S98, 2009.
- [26] L. Ju, M.K. Hurdal, J. Stern, K. Rehm, K. Schaper, and D. Rottenberg. Quantitative evaluation of three cortical surface flattening methods. *NeuroImage*, 28(4):869–80, December 2005.
- [27] Chiu Yen Kao, Michael Hofer, Guillermo Sapiro, Josh Stern, Kelly Rehm, and David a. Rottenberg. A geometric method for automatic extraction of sulcal fundi. *IEEE Trans. Med. Imaging*, 26:530–540, 2007.

- [28] L. Kharevych, B. Springborn, and P. Schröder. Discrete conformal mappings via circle patterns. *ACM Trans. Graph.*, 25:412–438, 2006.
- [29] Vladimir G Kim, Yaron Lipman, and Thomas Funkhouser. Blended intrinsic maps. *ACM Transactions on Graphics*, 30(4):1, 2011.
- [30] P. Kochunov, W. Rogers, J.-F. Mangin, and J. Lancaster. A library of cortical morphology analysis tools to study development, aging and genetics of cerebral cortex. *Neuroinformatics*, 10:81–96, 2012.
- [31] P. Koehl. Protein structure classification. In K.B. Lipkowitz, T.R. Cundari, V.J. Gillet, and B. Boyd, editors, *Reviews in Computational Chemistry*, volume 22, pages 1–56. John Wiley & Sons, Hoboken, NJ, USA, 2006.
- [32] P. Koehl and J. Hass. Automatic alignment of genus-zero surfaces. *IEEE transactions on pattern analysis and machine intelligence*, 36(3):466–78, March 2013.
- [33] P. Koehl and J. Hass. Automatic alignment of genus-zero surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(3):466–478, 2014.
- [34] P. Koehl and J. Hass. Landmark-free geometric methods in biological shape analysis. *J. R. Soc. Interface*, 12:1–11, 2015.
- [35] X. Li, Y. Bao, X. Guo, M. Jin, X. Gu, and H. Qin. Globally optimal surface mapping for surfaces with arbitrary topology. *IEEE Trans. Vis. Comput. Graphics*, 14:805–19, 2008.
- [36] C. Lin and J. Moré. Newton’s method for large bound-constrained optimization problems. *SIAM Journal on Optimization*, 9:1100–1127, 1999.
- [37] Y. Lipman and I. Daubechies. Conformal Wasserstein distances: Comparing surfaces in polynomial time. *Advances in Mathematics*, 227(3):1047–1077, 2011.

- [38] H. Lombaert, J. Sporring, and K. Siddiqi. Diffeomorphic spectral matching of cortical surfaces. *IPMI*, pages 376–389, 2013.
- [39] LM Lui, X Gu, and ST Yau. Convergence of an iterative algorithm for Teichmuller maps via harmonic energy optimization. *Mathematics of Computation*, page 20, 2015.
- [40] LM Lui, KC Lam, ST Yau, and X Gu. Teichmuller Mapping (T-Map) and Its Applications to Landmark Matching. *SIAM Journal on Imaging Sciences*, 7(1):391–426, 2014.
- [41] M. Lui, Y. Wang, T.F. Chan, and P. Thompson. Landmark constrained genus zero surface conformal mapping and its application to brain mapping research. *Appl. Numer. Math.*, 57:847–858, 2007.
- [42] F Luo. Combinatorial yamabe flow on surfaces. *Commun. Contemp. Math.*, pages 765–780, 2004.
- [43] R.H. MacNeal. *The solution of partial differential equations by means of electrical networks*. PhD thesis, California Institute of Technology, 1949.
- [44] J.F. Mangin, V. Frouin, I. Bloch, J. Régis, and J. López-Krahe. From 3d magnetic resonance images to structural representations of the cortex topography using topology preserving deformations. *J. Math. Imaging Vision*, 5(4):297–318, 1995.
- [45] F Memoli, G Sapiro, and P Thompson. Brain surface warping via minimizing lipschitz extensions. In *in Proc. Workshop Math. Foundations Computational Anatomy (MFCA)*, pages 58–67, 2006.
- [46] AM Oberman. A convergent difference scheme for the infinity laplacian: construction of absolutely minimizing lipschitz extensions. *Mathematics of Computation*, 74(251):1217–1230, 2005.

- [47] D. Pantazis, A. Joshi, J. Jiang, D.W. Shattuck, L.E. Bernstein, H. Damasio, and R.M. Leahy. Comparison of landmark-based and automatic methods for cortical surface registration. *NeuroImage*, 49:2479–2493, 2010.
- [48] U. Pinkall and K. Polthier. Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics*, 2:15–36, 1993.
- [49] B. Rodin and D. Sullivan. The convergence of circle packings to the riemann mapping. *J. Differential Geom.*, pages 349–360, 1987.
- [50] D.H. Salat, R.L. Buckner, A.Z. Snyder, D.N. Greve, R.S.R. Desikan, E. Busa, J.C. Morris, A.M. Dale, and B. Fischl. Thinning of the cerebral cortex in aging. *Cerebral Cortex*, 14(7):721–730, 2004.
- [51] P.V. Sander, J. Snyder, S.J. Gortler, and H. Hoppe. Texture mapping progressive meshes. *SIGGRAPH*, pages 409–416, 2001.
- [52] A Sheffer, B Lévy, M Mogilnitsky, and A Bogomyakov. ABF++: fast and robust angle based flattening. *ACM Transactions on Graphics*, 24(2):311–330, 2005.
- [53] J. Shi, P.M. Thompson, B. Gutman, and Y. Wang. Surface fluid registration of conformal representation: application to detect disease burden and genetic influence on hippocampus. *NeuroImage*, 78:111–34, 2013.
- [54] R. Shi, W. Zeng, Z. Su, Y. Wang, H. Damasio, Z. Lu, S.T. Yau, and X. Gu. Hyperbolic harmonic brain surface registration with curvature-based landmark matching. In *IPMI*, pages 159–170, 2013.
- [55] Y. Shi, R. Lai, D.J.J. Wang, D. Pelletier, D. Mohr, N. Sicotte, and A.W. Toga. Metric optimization for surface analysis in the Laplace-Beltrami embedding space. *IEEE transactions on medical imaging*, 33(7):1447–63, July 2014.

- [56] E.R. Sowell, P.M. Thompson, D. Rex, D. Kornsand, K.D. Tessner, T.L. Jernigan, and A.W. Toga. Mapping sulcal pattern asymmetry and local cortical surface gray matter distribution in vivo: maturation in perisylvian cortices. *Cerebral cortex*, 12:17–26, 2002.
- [57] B. Springborn, P. Schröder, and U. Pinkall. Conformal equivalence of triangle meshes. *ACM Transactions on Graphics*, 27(3), 2008.
- [58] Kenneth Stephenson. Circle Packing : A Mathematical Tale. *Notices Amer. Math. Soc*, 50:1376–1388, 2003.
- [59] I Stewart. *The Problems of Mathematics*. Oxford University Press, 1992.
- [60] D. J. Struik. Outline of a history of differential geometry. *j-ISIS*, pages 92–120.
- [61] Z. Su, W. Zeng, Y. Wang, Z. Lu, and X. Gu. Shape classification using Wasserstein distance for brain morphometry analysis. *Information Processing in Medical Imaging*, 17:411–423, 2015.
- [62] J. Sun, T. Wu, X. Gu, and F. Luo. Discrete conformal deformation : Algorithm and experiments. pages 1–34.
- [63] V. Surazhsky and C. Gotsman. Explicit surface remeshing. SGP, pages 20–30, 2003.
- [64] Gary K L Tam, Zhi Quan Cheng, Yu Kun Lai, Frank C. Langbein, Yonghuai Liu, David Marshall, Ralph R. Martin, Xian Fang Sun, and Paul L. Rosin. Registration of 3d point clouds and meshes: A survey from rigid to nonrigid. *IEEE Transactions on Visualization and Computer Graphics*, 19(7):1199–1217, 2013.
- [65] W.P. Thurston. Three dimensional manifolds, kleinian groups and hyperbolic geometry. *Bull. Amer. Math. Soc.*, 6:357–382, 1982.
- [66] Y. Tong, P. Alliez, D. Cohen-Steiner, and M. Desbrun. Designing quadrangulations with discrete harmonic forms. SGP, pages 201–210, 2006.

- [67] A. Tsui, D. Fenton, P. Vuong, J. Hass, P. Koehl, N. Amenta, D. Coeurjolly, C. DeCarli, and O. Carmichael. Globally optimal cortical surface matching with exact landmark correspondence. In *IPMI*, pages 487–498, 2013.
- [68] M. Vaillant and J. Glaunès. Surface matching via currents. *IPMI*, pages 381–92, 2005.
- [69] Y. Wang, W. Dai, Y.-Y. Chou, X. Gu, T.F. Chan, A.W. Toga, and P.M. Thompson. Studying brain morphometry using conformal equivalence class. *ICCV*, pages 2365–2372, 2009.
- [70] Y. Wang, X. Gu, T.F. Chan, P.M. Thompson, and S.-T. Yau. Conformal slit mapping and its applications to brain surface parameterization. In *MICCAI*, pages 585–593, 2008.
- [71] O. Weber, A. Myles, and D. Zorin. Computing extremal quasiconformal maps. *Computer Graphics Forum*, 31:1679–1689, 2012.
- [72] D.F. Wiley, N. Amenta, D. Ghosh, E. Delson, W. Harcourt-Smith, K. St.John, F.J. Rohlf, and B. Hamann. Evolutionary morphing. *IEEE VIS*, pages 431–8, 2005.
- [73] K.J. Worsley, S. Marrett, P. Neelin, A.C. Vandal, K.J. Friston, A.C. Evans, et al. A unified statistical approach for determining significant signals in images of cerebral activation. *Human brain mapping*, 4(1):58–73, 1996.
- [74] X.F. Yang, A. Goh, S.H. Chen, and A. Qiu. Evolution of hippocampal shapes across the human lifespan. *Human Brain Mapping*, pages 3075–3085, 2013.
- [75] B. T Thomas Yeo, Mert R. Sabuncu, Tom Vercauteren, Nicholas Ayache, Bruce Fischl, and Polina Golland. Spherical demons: Fast diffeomorphic landmark-free surface registration. *IEEE Transactions on Medical Imaging*, 29(3):650–668, 2010.
- [76] L. Younes, J.T. Ratnanather, T. Brown, E. Aylward, P. Nopoulos, H. Johnson, V.A. Magnotta, J.S. Paulsen, R.L. Margolis, R.L. Albin, M.I. Miller, and C.A. Ross. Region-

- ally selective atrophy of subcortical structures in prodromal HD as revealed by statistical shape analysis. *Human brain mapping*, 35(3):792–809, March 2014.
- [77] M. Zhang, F. Li, X. Wang, Z. Wu, S.Q. Xin, L.M. Lui, L. Shi, D. Wang, and Y. He. Automatic registration of vestibular systems with exact landmark correspondence. *Graphical Models*, 76(5):532–541, September 2014.
- [78] Tian Zhou and Kenji Shimada. An angle-based approach to two-dimensional mesh smoothing. *Proceedings of the 9th International Meshing Roundtable*, (412):373–384, 2000.