

Globally Optimal Surface Mapping for Surfaces with Arbitrary Topology

Xin Li, Yunfan Bao, Xiaohu Guo, Miao Jin, Xianfeng Gu, and Hong Qin

Abstract—Computing smooth and optimal one-to-one maps between surfaces of same topology is a fundamental problem in graphics and such a method provides us a ubiquitous tool for geometric modeling and data visualization. Its vast variety of applications includes shape registration/matching, shape blending, material/data transfer, data fusion, information reuse, etc. The mapping quality is typically measured in terms of angular distortions among different shapes. This paper proposes and develops a novel quasi-conformal surface mapping framework to globally minimize the stretching energy inevitably introduced between two different shapes. The existing state-of-the-art inter-surface mapping techniques only afford local optimization either on surface patches via boundary cutting or on the simplified base domain, lacking rigorous mathematical foundation and analysis. We design and articulate an automatic variational algorithm that can reach the global distortion minimum for surface mapping between shapes of arbitrary topology, and our algorithm is solely founded upon the intrinsic geometry structure of surfaces. To our best knowledge, this is the first attempt towards rigorously and numerically computing globally optimal maps. Consequently, we demonstrate our mapping framework offers a powerful computational tool for graphics and visualization tasks such as data and texture transfer, shape morphing, and shape matching.

Index Terms—Quasi-conformal surface mapping, harmonic map, uniformization metric, surface parameterization.

I. INTRODUCTION

A. Surface Mapping

How to compute surface mappings is one of the most fundamental problems in graphics and visualization fields. It aims to find a bijective (one-to-one and onto) map from one surface to another. Numerous applications such as shape registration, matching and comparison, shape morphing, and texture/attribute/motion transfer all benefit from such a bijective correspondence between two given surfaces. Researchers usually measure the mapping quality using angular or area distortions, because such criteria dictate the end effect of the enabling applications (e.g., texture mapping). Given two surfaces with different geometry, distortions are usually inevitable; we naturally want to seek the mapping that can minimize distortions as much as possible. However, if two given surfaces are not isometric to each other, there does not exist a mapping that can eliminate the angle and area distortion simultaneously.

In this work, we choose the harmonicity (measuring angular distortion) as the criterion because it is most physically meaningful. If we assume surfaces are made of elastic materials.

When surfaces deform and are mapped to others, the stretching energy caused by the elastic distortion can be formulated as harmonic energy. Among all possible mappings, a harmonic map minimizes the stretching energy and has its direct physical meaning. Also, the harmonic map minimizes angular distortions. For example, conformal mappings are harmonic, which are free of angular distortion.

Besides the physical intuition, harmonicity and conformal mapping have other merits which are critical for real-world applications. First, the dimension of conformal mappings between two given surfaces is finite; by fixing the images of finite points, the mapping can be uniquely determined. Therefore, these kinds of mapping are easy to control. Second, conformal mappings can transform arbitrary surfaces to several canonical domains, and convert all geometric processing into these domains. This greatly simplifies the complexity of these geometric processing algorithms. Third, the theoretic foundation and algorithms of conformal mappings are relatively mature.

In contrast, area-preserving mappings lack the physical meaning, lack analytic methodology as well as practical computation algorithm, and are hard to control. Therefore, in our current work, we use the harmonic maps and try to minimize the angular distortion.

B. State-of-the-art Techniques and Challenges

In terms of finding optimal mapping between two general surfaces under some specific criteria, current state-of-the-art techniques lack mathematically rigorous discussions and analysis in principle. On the other hand, the criteria of mapping quality on angular distortions have been analyzed and optimized in the surface flattening or surface parameterization research area. Surface parameterization aims to find a bijective map between surfaces and planes (or other simple canonical domains such as spheres), thus it can be treated as a special case of surface mapping since its target surface is usually just a plane or a sphere. Parameterization arises from the purpose of texture mapping and synthesis where the angular distortion is the most critical concern to quantify the mapping quality.

Despite its earlier connection with surface parameterization over canonical domains, finding a minimally-distorted surface mapping between two general surfaces are much more technically challenging. There are three key reasons as follows.

First, there are topological differences. Surface parameterization usually “flattens” a surface onto the plane, specifically, the surface is sliced apart into a topological disk, and the parameterization refers to a map from that disk to the plane, so

it has only one topological type. For mapping between general surfaces especially with high genus, a handle of one surface being mapped to which handle of the second surface needs to be considered. This topological factor has to be explicitly determined and it gives rise to the mapping complexity for shapes with nontrivial topology. Rigorously speaking, mappings between two given surfaces are classified into infinite *homotopy* classes [1]. Two maps are isotopic to each other, i.e., belonging to a same homotopy class, only if one can deform to another smoothly. A rigorous surface mapping framework should be able to handle an arbitrarily given homotopy type. On the other hand, only topologically equivalent mappings can be compared together; mappings from different classes should be considered separately since a best mapping may exist in each class. In the following discussion, we shall consider maps that are within the same homotopy class.

Second, due to the topological inequivalence between closed surfaces and the plane, as we mentioned above, parameterizing surfaces onto planer domains cuts the surface along a boundary. The simplified target canonical shape not only leads to some well-established numerical solving techniques, but also unavoidably pushes distortions towards its cutting boundary or a collection of some singularity points. In contrast, mappings between surfaces with same topology should prohibits the cutting, and find a “seamless” result. Therefore, we are not pursuing a map from a topological disk to the plane, but a continuous map between two surfaces with complicated topology.

Third, the most important reason of lacking globally optimized surface mapping techniques is the complex geometry of the general target surfaces. The non-smoothness of the target shape actually leads to the technical obstacle in finding the global optimum among all possible mappings. A natural way is to follow ideas in surface parameterization: we can optimize the map between surfaces by simply constructing an initial map, and then locally adjust it using a variational procedure until the distortion energy is reduced to the minimum. We can call this technique “the naive method”. When the target surfaces are genus zero (e.g., parameterization onto the sphere or plane), this approach can reach a globally optimized result. However, for mapping surfaces with non-trivial topology, due to the nonexistence of canonical target domain (see Section III-C), any local optimization process will inevitably get stuck at some local minima. This is the primary reason that other state-of-the-art methods use base meshes or hierarchical structure to circumvent this problem, while giving up searching for the global optimum.

C. Our Novel Solution

In this work, we introduce a novel computational framework to tackle the aforementioned challenging problems. Our method, based on the theories of Riemannian uniformization and harmonic maps, is both theoretically rigorous and practically efficient.

Considering two general surfaces with nontrivial topology, under their induced Euclidean metric, the target shape may have complicated geometry, and the harmonic maps are

usually not globally unique. Some harmonic maps are local minima of the stretching energy.

To globally reduce the stretching distortion without getting stuck locally, we propose to use the so-called *uniformization metric*. Under their uniformization metric, surfaces with non-trivial topology have constant non-positive Gaussian curvature everywhere, so that the harmonic map becomes **globally unique** [2]. (Please refer to Section III-A and the Appendix for more theoretic details and Section VII-D for experimental demonstrations). Uniformization theory states that for all surfaces, such uniformization metric does exist; and we can compute this metric efficiently by using existing techniques.

Under surfaces’ uniformization metric, we conduct our optimization process. It is guaranteed to converge to unique global harmonicity under surfaces’ uniformization metric. Specifically, our algorithm has the following important merits.

- **Optimality.** Harmonicity under uniformization metric can be globally optimized without worrying about any local optima. The resultant map minimizes the stretching energy and distortion.
- **Uniqueness.** The global harmonic map in hyperbolic space is unique; our algorithm converges to the same result starting from arbitrarily different initial mappings, as long as they belong to the same homotopy class.
- **Conformality.** For genus-zero surfaces, arbitrary harmonic map is free of angle distortions. In genus-one case, our optimized map minimizes the angle distortion among all possible maps. Between two general surfaces, if exists an angle-distortion-free map between them, our method guarantees to find such a conformal map.
- **Efficiency.** Harmonicity relaxation under the uniformization metric is performed in 2D, which is much more efficient and robust compared with any other iterative methods directly conducted over curved surfaces.

The main contributions of this work are:

1. We propose and articulate a novel approach to compute a globally optimal map minimizing distortions between two surfaces with the same non-trivial topology. This process is fully automatic and requires no user interaction. To the best of our knowledge, this is the first attempt to compute surface mapping with globally minimized energy for arbitrary high genus ($g \geq 1$) models.
2. Using the intrinsic geometric structure of surfaces, the convergence of our algorithm is guaranteed. We quantitatively evaluate its performance, and then design toolkits to clearly visualize the mappings, as well as analyze their converging effects.
3. We use our surface mapping as a powerful tool for data and texture transfer, shape morphing, cross-surface parameterization onto canonical shape domains, shape matching, and shape comparison. Our globally optimized mapping demonstrates its great efficacy in these graphics and visualization applications, with potentials in the broad scope.

The remainder of this paper is organized as follows. We will briefly review the related literatures in Section II. Then in Section III, we introduce the theory and algorithm of our

method. Our algorithm proceeds in three main steps, as discussed in Section IV, Section V and Section VI, respectively. We then discuss our mapping performance and properties in Section VII. Finally, we demonstrate our experimental results with some applications in Section VIII and conclude our work in Section IX. In the accompanying appendix, we prove the existence, global uniqueness, and the one-to-one property of the harmonic map and we also show our algorithm will converge to such an optimized map uniquely.

II. RELATED WORK

Our current research builds upon previous work in surface mapping, conformal geometry, and non-Euclidean geometry. Earlier work on establishing a bijective mapping is mostly motivated by the need of blending two shapes. A natural approach is to use some canonical shape such as a sphere or the plane as the intermediate domain. Kent et al. [3] mapped star-shaped surfaces onto spheres, and merged them by clipping one sphere to the other. Kanai et al. [4] used harmonic map on disk to build correspondence between two genus-zero closed or open surfaces. Alexa [5] wrapped two genus-zero surfaces onto a unit sphere, and computed the mapping by minimizing some distance function. Asirvatham et al. [6] used progressive mesh and their constrained spherical parameterization to map genus-zero surfaces onto the common spherical domain. These types of techniques are usually based on spherical parameterization techniques [5], [7]–[9] or planar parameterization techniques [10].

Spheres and planar disks are natural domains for computing maps with minimized stretching energy directly. However, they can only serve as intermediate domains when the two surfaces are of genus zero. For high genus surfaces, these kinds of canonical domains can not be found. In this work, we focus on finding stretching-optimized maps between surfaces with non-trivial topology.

Approaches for surfaces with non-trivial topology are usually applied through another direction ([11]–[18]). They typically segment the meshes into subregions first. For example, in [12] and [13], a common coarse base domain mesh has to be constructed manually by the user with domain knowledge in topological surgery; in [14], [16]–[18], feature points are firstly provided by users, then some automatic subregion tracing algorithms or progressive meshes are applied for coarse base mesh generation. The advantage of these approaches is that feature correspondence can be intuitively incorporated by making the feature vertices the corners of the patches. The common drawbacks are that constructing the patch layout oftentimes involves a number of fragile heuristic algorithms. Furthermore, the mappings are generally only C^0 continuous across the patch boundaries. In applications such as building domains for splines, a global continuity is critical. The work of [19] addresses the continuity problem by taking into account linear transition functions across patch boundaries. Manifold concept in mapping is introduced in [20], which primarily focuses on topology instead of geometry, thus is difficult for designing optimization algorithms.

Conformal maps have been extensively studied in the literature of the surface parameterization field. [10], [21] provide

extensive surveys of state-of-the-art techniques in the field. We only briefly review some most related work, and refer interested readers to these surveys for details.

Angle preservation is typically addressed either from the harmonic point of view (Dirichlet energy) [22]–[24] or from the conformal point of view (Cauchy-Riemann equation) [24], [25]. Most recently, the hyperbolic structure of Riemannian surfaces has been introduced to surface parameterization. Thurston firstly introduced circle packing in [26]. An effective algorithm and implementation is presented by Stephenson in [27]. Circle packing has also been generalized to circle patterns [28] and used for surface parameterization in [29]. Hamilton first introduced Ricci flow on surfaces in [30]. Theoretical results of combinatorial Ricci flow are later generalized in [31], and applied in surface parameterization fields by [32].

III. THEORY AND ALGORITHM

A. Uniformization Metric

On a surface, a *metric*, or *Riemannian metric* is a tensor that defines inner product on the tangent plane at each point. With the metric, the length of a tangent vector can be determined, and the angle between two tangent vectors can be explicitly computed.

Suppose S is a smooth surface embedded in \mathbb{R}^3 ; it has the induced Euclidean metric \mathbf{g} . We denote the surface S together with its equipped metric \mathbf{g} as (S, \mathbf{g}) . If $\lambda : S \rightarrow \mathbb{R}$ is a scalar function defined on the surface, then $\bar{\mathbf{g}} = e^{2\lambda}\mathbf{g}$ is another metric on S . Any angles on the surface measured by \mathbf{g} equals to those measured by $\bar{\mathbf{g}}$, therefore, we say $\bar{\mathbf{g}}$ is conformal to \mathbf{g} , meaning that changing between these two metrics is angle-preserving.

Given two surfaces S_1 and S_2 , the uniqueness of the harmonic map from S_1 to S_2 , as we will discuss in the upcoming section, is determined by the distribution of the Gaussian curvature K of S_2 . It is important to note that K is fully determined by the equipped metric of the surface. The relation between the curvatures K and \bar{K} under \mathbf{g} and $\bar{\mathbf{g}}$ is $\bar{K} = e^{2\lambda}(-\Delta\lambda + K)$.

Riemann uniformization states that for an arbitrary closed surface, there exists a unique λ such that $e^{2\lambda}\mathbf{g}$ induces constant Gaussian curvature. Furthermore, the constant is one of the three choices $\{+1, 0, -1\}$ for surfaces with zero, one, and higher genus, respectively. Such kind of metric $e^{2\lambda}\mathbf{g}$ is called the *uniformization metric* of the surface. The uniformization metric can be computed using Ricci flow method (see Section V).

B. Euclidean Harmonic Map and Conformal Map

Given two surfaces embedded in \mathbb{R}^3 with the induced Euclidean metrics (S_1, \mathbf{g}_1) and (S_2, \mathbf{g}_2) , $f : S_1 \rightarrow S_2$ is a map between them, the harmonic energy (stretching energy) is defined as

$$E(f) = \int_{S_1} |\nabla f|^2 dA_1, \quad (1)$$

where ∇f is the gradient of the map. A harmonic map is a *critical point* of the harmonic energy. Harmonic maps depend on the Riemannian metrics. However, if $f : (S_1, \mathbf{g}_1) \rightarrow$

(S_2, \mathbf{g}_2) is a harmonic map, then $f : (S_1, e^{2\lambda} \mathbf{g}_1) \rightarrow (S_2, \mathbf{g}_2)$ is also a harmonic map.

If a map preserves angles, then the map is called a conformal map. Analytically, if the *pull back* metric $f^* \mathbf{g}_2$ on S_1 is conformal to \mathbf{g}_1 , $e^{2\lambda} \mathbf{g}_1 = f^* \mathbf{g}_2$, then f is conformal. A conformal map must be harmonic. For closed genus-zero surfaces, harmonic maps are also conformal. In general case, if S_1 and S_2 are with complicated topology, then there may not exist a conformal map. But there is a special map, which minimizes the maximum of the angle distortion; such a map is called the extremal quasi-conformal map. For genus-one case, such an *extremal quasi-conformal map* is the harmonic map under uniformization metric. Therefore, if the given surfaces are genus-one, our algorithm converge to the extremal quasi-conformal map.

C. Uniqueness of Harmonic Maps

The uniqueness of harmonic maps between surfaces is determined by the shape of target objects. For genus-zero surfaces, there are infinite harmonic (conformal) maps, all with zero angular-distortion. Each two of these maps differ by a möbius transformation on the sphere domain.

Harmonic maps between surfaces with non-trivial topology are also not unique if the Gaussian curvature of the target surface is positive somewhere. However, if the target surface has non-positive Gaussian curvature everywhere, then the harmonic map exists and is **unique**. For example, if the Euler number $\chi(S_i) < 0, i = 1, 2$, and we apply uniformization metric \mathbf{g}_2 on S_2 , then harmonic map f exists and is unique, with its energy $E(f)$ reaching the global minimum.

Therefore, between arbitrary two surfaces with genus ≥ 2 , there uniquely exists such a stretching-minimized harmonic map. For genus-one surfaces, $\chi = 0$, under uniformization metric, the harmonic maps are not unique, but only differ by a rigid translation on the \mathbb{R}^2 universal covering space, and we can use one feature point to uniquely determine it.

D. Poincaré Disk Model and its Harmonic Maps




Fig. 1. (a) The yellow patch represents a chart on the two-hole torus model; (b) Embed the two-hole torus model in the Poincaré disk; (c) A möbius transformation moves the chart to the center of the Poincaré disk.

If given surfaces are with higher genus, their uniformization metrics can only be embedded in hyperbolic space. We have to carry out our computation in this space, which can be modeled by the Poincaré disk as follows.

The Poincaré disk is the unit disk on the complex plane $z\bar{z} \leq 1$, with the Riemannian metric $ds^2 = \frac{4dzd\bar{z}}{(1-z\bar{z})^2}$. Our goal is to compute a map $f : (S_1, \mathbf{g}_1) \rightarrow (S_2, \mathbf{g}_2)$. We use

their uniformization metrics and compute a harmonic map $\bar{f} : (S_1, \bar{\mathbf{g}}_1) \rightarrow (S_2, \bar{\mathbf{g}}_2)$. The computational algorithm of hyperbolic harmonic maps is based on theoretic results in [33].

We denote the parameters of S_1 on the Poincaré disk as (x, y) , the parameter of S_2 as (u, v) , then the map \bar{f} is represented as $\bar{f}(x, y) = (u(x, y), v(x, y))$. The harmonic energy is

$$E(\bar{f}) = \int_{S_1} 4 \frac{|\nabla u|^2 + |\nabla v|^2}{(1 - u^2 - v^2)^2} dx dy, \quad (2)$$

where ∇u is $(\frac{\partial u}{\partial x}, \frac{\partial u}{\partial y})$ and ∇v is $(\frac{\partial v}{\partial x}, \frac{\partial v}{\partial y})$.

The harmonic energy in hyperbolic space (2) has more complicated form than harmonic energy in Euclidean space (1). We simplify the problem using the following two merits of hyperbolic harmonic energy:

1. In a small neighborhood of the origin $u^2 + v^2 < \epsilon$, since $(1 - u^2 - v^2)^{-2} \rightarrow 1$, the hyperbolic metric is close to the Euclidean metric, the hyperbolic harmonic energy is close to the Euclidean harmonic energy. We can optimize the hyperbolic energy by minimizing the Euclidean energy.
2. If ϕ is a Möbius transformation of the Poincaré disk, then the composition $\phi \circ \bar{f}$ and \bar{f} have the same hyperbolic harmonic energy. This is because Möbius transformation is the rigid motion in the hyperbolic space, harmonic energy is invariant under such isometries of the target surface.

We describe our computational methodology for hyperbolic harmonic maps as follows.

1. The surfaces are tessellated to many small triangular patches, $S_1 = \bigcup_i T_i$, where T_i is a triangular patch, then the harmonic energy is decomposed to the summation of the energy of the map restricted on these patches, the sub-maps, $E(\bar{f}) = \sum_i E(\bar{f}_i), \bar{f}_i : T_i \rightarrow \mathbb{H}^2$.
2. Each sub-map \bar{f}_i is composed with a Möbius transformation ϕ_i , such that the image $\phi_i \circ \bar{f}_i(T_i)$ is in the neighborhood of the origin.
3. If the tessellation is refined enough, T_i is small, and the diameter of its image under the corresponding sub-map is within an ϵ -threshold, the hyperbolic energy can be approximated by Euclidean harmonic energy with high accuracy.

Therefore, computing the harmonic map under hyperbolic metric, equivalent to minimizing hyperbolic harmonic energy, is now converted to optimizing a collection of Euclidean harmonic energies of sub-maps. We can use *mean value property* of the harmonic function to minimize the Euclidean harmonic energy.

E. Discrete Algorithm

We summarize our approach as the following discrete algorithm:

The inputs are the source surface S_1 and the target surface S_2 . The output is the harmonic map f under the uniformization metric of S_2 .

1. Construct an initial map $\tilde{f} : S_1 \rightarrow S_2$ (See Section IV).
2. Compute the conformal deformation (uniformization metric) of S_2 using the technique introduced in [32], then embed S_2 in the canonical domain \mathbb{C} or \mathbb{H}^2 , $\phi_2 : S_2 \rightarrow \mathbb{C}$ or $\phi_2 : S_2 \rightarrow \mathbb{H}^2$. (See Section V)
3. Compose the \tilde{f} and ϕ_2 to get $\omega : S_1 \rightarrow \mathbb{C}$ or $\omega : S_1 \rightarrow \mathbb{H}^2$, and apply heat diffusion on dynamic charts to minimize the harmonic energy:

$$\frac{d\omega}{dt} = -\Delta\omega. \quad (3)$$

(See Section VI)

4. When ω converges to the global minimum, let $f_{1H} = \omega$ and get the final map $f = \phi_2^{-1} \circ f_{1H}$.

IV. INITIAL MAPPING

We first build up an initial mapping between the given surfaces. The initial mappings determine the homotopy class of the resultant surface mapping (which will be discussed later in Section VII-B). Our pipeline for creating the initial map has two steps: (1) we unfold both surfaces to disks through a cut graph called the *system of loops* (See Section IV-A); (2) we map two surfaces via the disk domain (See Section IV-B).

A. Cutting a Surface into a Topological Disk

An orientable closed surface of genus g , ($g \geq 1$) can be cut into a single topological disk by removing a so-called *cut graph*. Computing a special case of cut graphs which pass through a common given base point, called *systems of loops*, is studied in computational geometry field. One of the state-of-the-art techniques, [34], used an efficient greedy algorithm to get an optimal cutting loop. We refine their algorithm for our surface cutting.

We first briefly describe their algorithm for computing a system of loops L on the given mesh S and the base point x :

- (1) Compute the shortest paths tree T of S from x .
- (2) For each edge $e \in S \setminus T$ (i.e. $e \notin T$), compute the shortest loop that contains e , denoted as $\sigma(e)$, which consists of 2 shortest paths from x to endpoints of e plus the e itself.
- (3) Compute the dual graph of $S \setminus T$, denoted as $(S \setminus T)^*$. Compute its maximum spanning tree T^* , where the weight of each dual edge e^* is $\sigma(e)$.
- (4) Get the set E' which contains every edge that is neither in T nor crossed by T^* .
- (5) E' has $2g$ edges e_1, e_2, \dots, e_{2g} . Compute shortest loop $\sigma'(e_i)$ containing each e_i . These loops constitute the system loop L .

In [34], they assumed that all the shortest paths from each point on the cut path to the base point only intersect at the common base point. This assumption holds in the smooth case but often fails for triangular mesh representations. Thus in step (5), shortest paths on triangle meshes may intersect each other, especially for high genus surfaces. For example, on a genus-six surface, 12 loops will go through the base point, meaning that the valence of the base point should at least be 24 to prevent paths' intersections outside the base point. Such high density connectivity is hardly satisfied in ordinary mesh data.




Fig. 2. Local Refinement on the System of Loops Computation. When the blue cut path intersect with the existing red path in V_2 , we apply a local refinement. The intersected path $[V_1, V_2, V_3]$ segment is replaced by the new green path. Yellow segments are new edges inserted during edge splits in the refinement.




Fig. 3. Refinement on the System of Loops Computation.

Therefore, a robust algorithm has to adaptively change the connectivity.

As shown in Figure 2, locally, if a cut path (blue) intersects with an existing path (red) in one point. We apply the following algorithm on the blue curve to make it bypass the red one:

Algorithm 1:

- (1) Spin around the intersected vertex V_2 , enqueue all faces between $[V_1, V_2]$ and $[V_2, V_3]$ (For example, f_1, f_2, f_3, f_4 here).
- (2) Set V_1 as the current point p .
- (3) Pop face f out from the head of queue. If $[V_2, V_3]$ is in f , add edge $[p, V_3]$ to the new path and STOP; else GOTO (4).
- (4) Split the edge opposite to the current vertex p . The new split point is denoted as q . Add the edge $[p, q]$ into the new path, move to q ; set $p := q$. GOTO (3).

After applying this algorithm, we replace the intersected path $[V_1, V_2, V_3]$ segment by the new path (as shown in green). The yellow segments are edges newly inserted onto the mesh during the edge split procedure.

In general, if the intersected parts have more than one point, we apply the above algorithm 1 iteratively on each intersected vertex. Figure 3 illustrates this. In the small picture (upper left), a cut path (blue) passes through an existed cut path (red). We apply the following algorithm on the intersected segments:

Refinement Algorithm:

- (1) Find out the point right before the intersection (V_h) and the first point right after this intersection (V_t). Push all vertices on the current path between V_h and V_t into a queue Q .




Fig. 4. Slice both meshes open and map them to a canonical $4g$ -gon; compose these two maps; get the initial mapping. Mappings of different (color-coded) regions are shown respectively with different colors.

(2) Pop each vertex in Q , and apply Algorithm 1 on it.

Intersections usually happen near the base point because cutting paths are dense in this region. The V_t in such case is the base point, and the same refinement process is applied.

B. Initial Mapping via $4g$ -gon

With the system of loops, we slice each surface onto a topological disk. For a genus g surface, the cut graph passes through the base point $4g$ times, making the disk a topological $4g$ -gon. We map two given surfaces via this $4g$ -gon, as the procedure illustrated in the Figure 4.

1. Slice each surface along its system of loops to open it up onto the $4g$ -gon.
2. Flatten each sliced surface to the canonical $4g$ -gon, using the harmonic map with fixed boundary.
3. On the canonical planar parameter domain, map S_1 to S_2 via barycentric coordinates. Unlike [1], we do *not* extract a meta-mesh by overlaying the two planar domains. Instead, we use an approximation mesh S'_2 with only the connectivity of S_1 (though its shape is like S_2), and we may later employ an adaptive remeshing procedure (Section VII-E) for mapping refinement in areas where under-sampling occurs.
4. Stitch the topological disk S'_2 along the original cutting boundary back to the closed surfaces.

By the above algorithm, we get an initial mapping from S_1 to S_2 . This initial mapping is only used to determine the homotopy type. In the following sections, we will prove and demonstrate that if two initial cuts induce two maps belonging to the same homotopy class, then the final results are identical.

V. COMPUTING UNIFORMIZATION METRIC

According to our previous discussion, given a surface $S \subset \mathbb{R}^3$ and its induced Euclidean metric (represented by its first fundamental form \mathbf{g}), let $u : S \rightarrow \mathbb{R}$ be a globally defined function on S , then $e^{2u}\mathbf{g}$ is another Riemannian metric on S , which is a *conformal metric* to the original induced Euclidean metric.

Riemann uniformization theorem [35] states that for any S , there exists a unique conformal metric, such that it induces




Fig. 5. Side-by-side comparison between Distortions of Initial Map (left) and Optimized Map (right).

constant Gaussian curvature K and zero geodesic curvature,

$$K = \begin{cases} +1 & \chi(S) > 0 \\ 0 & \chi(S) = 0 \\ -1 & \chi(S) < 0 \end{cases},$$

where χ is the Euler characteristic. Such kind of metric is called the *uniformization metric*.

We compute the uniformization metric $e^{2u}\mathbf{g}$ using the Ricci flow method [30]. Ricci flow is defined as

$$\frac{du(t)}{dt} = -2K(t), \quad (4)$$

where $K(t)$ is the Gaussian curvature induced by the metric $e^{2u(t)}\mathbf{g}$, under the area preserving constraint

$$\int_S d\sigma = \int_S e^{2u(t)} d\sigma.$$

In practice, all surfaces are represented as triangular meshes. Basically, for a triangular face ABC on the mesh with edge lengths a, b, c , we do *not* treat it as a planar triangle in the **Euclidean space**, but rather a triangle in **Hyperbolic space**. All the angles in the triangle can then be calculated using hyperbolic cosine law, and the discrete Gaussian curvature on each vertex is defined as the difference between 2π and the summation of all the corner angles surrounding the vertex.

We associate each vertex v_i with a circle of radius γ_i . Two circles centered at the end vertices of an edge e_{ij} intersect at an angle Φ_{ij} . The edge length of e_{ij} equals $l_{ij} = \sqrt{\gamma_i^2 + \gamma_j^2 + 2 \cos \Phi_{ij}}$.

Conformal maps transform infinitesimal circles to infinitesimal circles and preserve the intersection angles among the circles. Therefore, we only modify the circle radii γ_i and keep the intersection angles Φ_{ij} . Let

$$u_i = \begin{cases} \ln \gamma_i & \chi(S) = 0 \\ \ln \tanh \frac{\gamma_i}{2} & \chi(S) < 0 \end{cases},$$

The discrete Ricci flow is similar to the continuous Ricci flow in the form:

$$\frac{du_i}{dt} = -K_i,$$

where K_i is the Gaussian curvature at v_i .

The Ricci flow will converge [32], such that all discrete Gaussian curvatures are constant, and the edge lengths approximate the uniformization metric.

If the surface S is equipped with the uniformization metric, then S can be isometrically and periodically embedded in the following three canonical spaces, the unit sphere for $\chi(S) > 0$, the plane for $\chi(S) = 0$, and the hyperbolic space $\chi(S) < 0$. When $\chi(S) = 0$ the metric is called *flat metric* since curvature is zero everywhere, and when $\chi(S) < 0$, it is called *hyperbolic metric*. In Figure 1, we demonstrate the hyperbolic embedding of the two-hole torus model.

We use the Poincaré hyperbolic disk model to represent the hyperbolic space \mathbb{H}^2 . The Poincaré hyperbolic disk is a two-dimensional space defined in the unit disk $\{z \in \mathbb{C} : |z| < 1\}$ on the complex plane \mathbb{C} with hyperbolic metric. The hyperbolic metric is defined as

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

The geodesics (hyperbolic lines) in the Poincaré disk are Euclidean circular arcs perpendicular to the boundary $|z| = 1$. The rigid motions in the hyperbolic plane are the Möbius transformations $z \rightarrow w, z \in \mathbb{C}$ with the form

$$w = e^{i\theta} \frac{z - z_0}{1 - \bar{z}_0 z}, \quad (5)$$

where z_0 is an arbitrary point inside the unit disk, θ is a rotation angle. This formula rigidly transforms the hyperbolic disk so that the point z_0 is moved to the origin (Figure 1 (b) and (c)).

VI. MAP OPTIMIZATION

With the uniformization metric defined on the target mesh S_2 , we can perform the heat diffusion procedure to optimize the initial map. Because of the constant curvature distribution under the uniformization metric, our relaxation will not get stuck in local minima. An arbitrary initial map can be used as the start of our optimization procedure; it can be stretched and distorted, or even contain local flip-overs. Our optimization procedure (Section VI-B) converges to a unique bijective global optimum robustly; more discussion about this will be given in Section VII-D, and the rigorous proof is given in the Appendix.

In Figure 5, we visualize the distortion of the initial mapping from the amphora model to the vase model by texture mapping and displaying the connectivity. The checkerboard texture mapped is distorted (irregular pattern as shown in the top left image) by this initial mapping. This initial mapping, like all methods based on cutting, induces great distortions near the boundary. By relaxing each vertex on dynamic charts discussed in the following section, we alleviate the distortions all over the mesh and reach a global minimum.

A. Chart Construction

In order to smooth the mapping between S_1 and S_2 , we need to redistribute vertices of S_1 on the domain of S_2 following the heat diffusion flow. We can either embed the whole S_2 onto \mathbb{C} (genus one) or \mathbb{H}^2 (higher genus) and perform




Fig. 6. The Dynamic Covering Chart on S_2 . Given a vertex V and its one-ring on S_1 , the left figure shows a covering chart on S_2 : the vertex V is mapped to the red face; its one-ring neighbors are mapped to the yellow faces. The right figure shows the domain of the chart. The white arrow indicates the gradient direction of the harmonic energy.

the redistribution globally; or directly flow over local charts equipped with uniformization metric. In this work, we use the second method: dynamically construct a set of overlapping local charts on S_2 and perform relaxation within these charts. Compared with using one global patch, the dynamical local charts method has two important **advantages**:

- The vertices may need to flow across the cutting boundary to relax the energy. On one parameter patch domain, it is difficult to perform the relaxation across the boundary.
- Globally embedding the target mesh onto a large patch is numerically less accurate, especially for hyperbolic metric. The local embedding of small charts is more precise.

For each vertex v_1 on S_1 , we create at least one chart on S_2 that covers the images of its 1-ring on S_1 , meaning that the chart contains all faces onto which the v_1 's 1-ring are mapped.

As shown in Figure 6, to construct a covering chart for the 1-ring of a vertex on S_1 , we first map the vertices of the 1-ring to S_2 . Each vertex in this 1-ring is mapped to a face on S_2 . Given this set of faces, we find a patch on S_2 that contains these faces and is homeomorphic to a disk. We first compute an approximate geodesic distance from the ‘center’ face (red) to all other faces. Then we add the faces to the chart through Breadth First Search (BFS) while maintaining disk topology. Faces which are closer to the ‘center’ face are given higher priority during the BFS. After the chart has been constructed, we tile it in \mathbb{C} (or \mathbb{H}^2 , according to the genus of the mesh). In this way, we get a **locally constructed**, yet **globally parameterized** chart, extracted as a small subset of the continuous global parameter domain. During the relaxation, the mapping of vertices and their 1-ring can change, new charts are dynamically created when necessary; old charts which are unused for a user-specified amount of time are removed from memory on the fly.

This is Not Local Parameterization. Note that chart-based approaches have been used in local-parameterization-based remeshing [36], [37]. And our approach is fundamentally different from them in that we are not *locally parameterizing* these one-ring charts, but directly *embedding* the pre-computed uniformization metric. Local parameterization computes the flattening of charts every time separately, while we use the global metric so that a globally consistent covering is achieved.

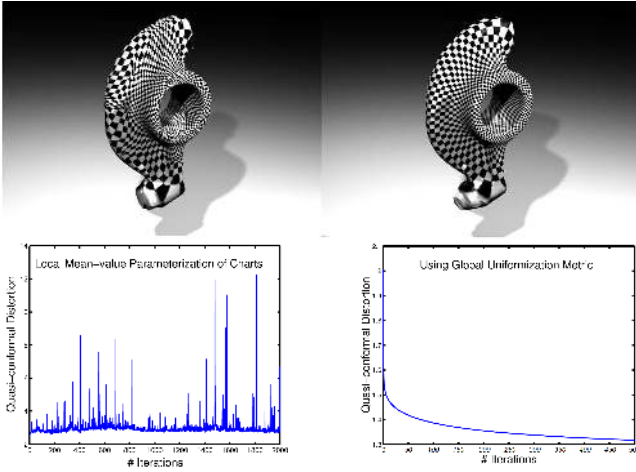


Fig. 7. Local Parameterization VS Global Uniformization Metric (Map the Torus to the Rocker Arm). Top row: side-by-side comparison between local approach and global approach. Bottom row: temporal statistics of convergence performance.

With the uniformization metric, we trivially get the flattening of each local chart by tiling it in a proper local patch domain (Figure 1). In other words, pre-computed metric already defines all the edge lengths in the mesh of the given chart; we only conduct a tiling of this triangular mesh.

The relaxation result demonstrates the key difference between local parameterization and our approach: relaxation based on local parameterization will get stuck locally, while using a globally consistent uniformization metric guarantees the global convergence. To demonstrate this, we perform experiments as shown in Figure 7. Compared with the relaxation on the uniformization metric (right), the relaxation using local parameterization¹ (left) will get stuck in some local optimum, and cannot produce the desirable result.

B. Relaxation

We let the skin ‘flow’ on the target planar domain so that the harmonic energy is minimized. This is performed via an iterative heat diffusion (relaxation) procedure. The discrete harmonic energy of a map f is defined as

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

where $|\cdot|$ is the norm with respect to Euclidean metric and w_{ij} ’s are the discrete harmonic cotangent weights. We use the gradient descent method to minimize the harmonic energy. In each single relaxation step, a vertex is moved in the domain following the gradient of harmonic energy by the Laplacian operator, which is defined as

$$\Delta f = \sum_{j \in N_i} w_{ij} (f(v_j) - f(v_i)),$$

where N_i is the index of neighboring vertices and $f = \omega$ is the composed map as given in Equation 3. Therefore, the vertex in the domain is moving towards the new position:

$$f'(v_i) = f(v_i) + \Delta f.$$

¹Local parameterization of the charts onto circular disks.

During the iteration procedure, the harmonic energy (from the source mesh to the target domain) monotonically decreases. For genus one meshes, they are embedded in \mathbb{C} under uniformization metric, these operators can be used directly. For higher genus meshes embedded in \mathbb{H}^2 , we still use this Euclidean Laplacian operator to relieve the harmonic energy after an isometric transformation in \mathbb{H}^2 , which is also called the *Möbius transformation*. The reason that we can approximate hyperbolic Laplacian operators using Euclidean Laplacian operators had been discussed in the previous Section III-D. By the Möbius transformation we rigidly transform the domain of the local chart so that the parameterization of the vertex being relaxed coincides with the center of the Poincaré disk (Figure 1). Near the origin, the hyperbolic metric $ds = \frac{2|dz|}{1-|z|^2}$ only differs by a constant factor from the Euclidean metric and thus our Euclidean Laplacian operator is a linear approximation to the Hyperbolic Laplacian operator in this relaxation region. (The local chart is usually small, so the approximation is with high precision.) For numerical issue, we change the f value on each vertex to the target using a step size 0.5, i.e. $f'(v_i) = f(v_i) + \frac{\Delta f}{2}$.

VII. DISCUSSIONS ON MAPPING PERFORMANCE AND PROPERTY

A. Mapping Quality Measurement

Harmonic energy is a natural energy to measure the stretching energy induced by the mapping. A physically meaningful mapping in reality ought to minimize the harmonic energy.

When the conformal mapping from S_1 to S_2 does not exist, the quasi-conformal mapping f maps circular regions around a local point into ellipses. The ratio of the major to the minor axis is called the *dilatation* D at this point. We use a discrete variance D' to measure the conformality of this mapping. The definition is as follows.

Given a local triangle (q_1, q_2, q_3) , $q_i = (x_i, y_i, z_i)$ of the original mesh mapped onto a triangle (p_1, p_2, p_3) , $p_i = (u_i, v_i)$ on $2D$. The interior discrete mapping $S(p) = S(u, v) = q$ is represented by

$$S(p) = (\langle p, p_2, p_3 \rangle q_1 + \langle p, p_3, p_1 \rangle q_2 + \langle p, p_1, p_2 \rangle q_3) / \langle p_1, p_2, p_3 \rangle$$

, where $\langle a, b, c \rangle$ denotes area of triangle abc . The partial derivatives of Jacobian are

$$S_u = (q_1(v_2 - v_3) + q_2(v_3 - v_1) + q_3(v_1 - v_2)) / (2\langle p_1, p_2, p_3 \rangle).$$

and

$$S_v = (q_1(u_3 - u_2) + q_2(u_1 - u_3) + q_3(u_2 - u_1)) / (2\langle p_1, p_2, p_3 \rangle).$$

The larger singular value Γ and smaller singular value γ of the Jacobian are given respectively [38]:

$$\Gamma, \gamma = \sqrt{\frac{E + G \pm \sqrt{(E - G)^2 + 4F^2}}{2}}$$

, where E, F, G are terms for the first fundamental form.

We compute D' on each triangle using $D' = \frac{\Gamma}{\gamma}$. The maximal value of D' of the mapping on the surfaces is determined by their geometry. As we mentioned above, in

genus one case, the harmonic map we get is the extremal quasi-conformal map, minimizing the angular distortion. We test our mapping performance against the theoretical bound using the following experiment. Given two torii T_1 and T_2 ; T_1 has minor and major radii 0.5 and 2 respectively, while T_2 has these radii 0.5 and 1.5. T_1 and T_2 can be conformally mapped onto two $2D$ rectangles $R_1(a_1, b_1)$ and $R_2(a_2, b_2)$, where (a_i, b_i) are the width and length of the rectangle. The extremal quasi-conformal mapping between T_1 and T_2 has the lowest theoretical bound given by the modulus ratio of T_1 and T_2 . In our setting above, these two modulus are 0.3531 and 0.5762, meaning the theoretically optimal D' bound between T_1 and T_2 is 1.632. We plot the performance of our mapping in Figure 9(c) (red curve). The x -direction shows the iteration numbers and the y -direction shows the global quasi-conformal distortions.

B. Homotopy Classes of Initial Mappings

When an arbitrary initial mapping is built up, the homotopy class of the mapping is determined. The subsequent optimization procedure (Section VI) reaches a unique optimized result in this same class. The slicing order of loops in the two systems of loops decides the homotopy class of the mapping. Usually, if we arbitrarily pick an order, an optimized result will be reached in that corresponding homotopy class; but this kind of arbitrary surface mapping may not be what we intuitively want. We naturally want handles mapped to handles consistently. To get consistent slicing orders of systems of loops, first, we can compute the canonical handle and tunnel loops using the method of [39]; second, with these handle and tunnel loops, we can decide the homotopy class of each closed loop in the system of loops, this pair loops in two systems of loops, providing the consistent slicing orders in two systems of loops. In this way, we correspond handles in the source surface with handles in the target surface. For two genus g surfaces, there will be g -factorial consistent mappings, any of them is visually reasonable.

Furthermore, in many applications, users may want more precise controls on the mapping. For example, sometimes handles of the source surface need to be mapped to some specific handles of the target surface. Also, users may require some feature points to be mapped. Both of these can be easily implemented in our framework as follows.

C. Constraints and User Controls

To assure the handle-correspondence, users only need to pick up a corresponding slicing order of two systems of loops, on the $4g$ -gon disk, users can easily set up this order once the systems of loops are computed.

In order to have constraints on the feature points, existing parameterization techniques for topological disk surfaces with constraint points, for example, MAPS [14] can be applied for the initial map. Also, many existing surface mapping framework [1], [14]–[18] all allow the feature point correspondence, and they can be applied as the initial map. In our work, since we use Carner et al.’s method [1] to generate the initial map, we also apply their method for the initial feature registration.

Starting from the initial mapping with feature points matched correspondingly, we can perform the optimization

without relaxing the feature points and prevent any neighboring vertices movement that violates the validity of the triangular mesh during the relaxation.

The insertion of landmark definitely may cause larger distortion in the neighborhood, because now the relaxation cannot be performed freely and the resultant mapping is not globally optimized. The detail discussion about feature correspondence is beyond the focus of this work. We will explore along this direction in our future work.

D. Global Convergence and Performance

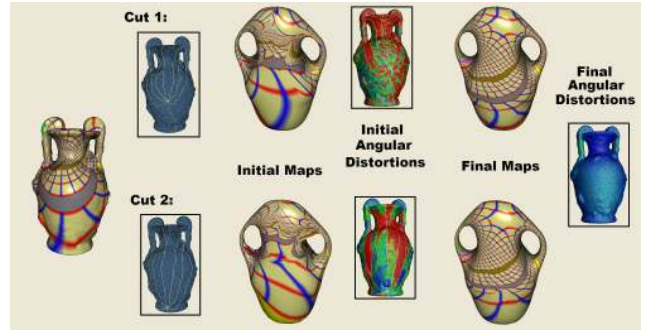


Fig. 8. Global Uniqueness of the Optimized Map in the same Homotopy Class.

Our surface mapping optimization converges robustly. Under the same homotopy class, different initial cuttings/mappings reach the same global optimized result. A rigorous proof is given in the appendix. We also perform experiments and visualize this in Figure 8: from left to right, the first column shows the original Amphora model and its texture; the second column are two different initial cut paths. In the third column, we transfer the Amphora’s textures onto the target Vase model using the corresponding initial maps. Their angular distortion distributions (average D') are color-coded in the fourth column. Transferred textures on the Vase model using the final maps are illustrated in the fifth column. Their final maps are almost the same and have the final distortion color-coded in the rightmost column.

We plot more experimental performances on computation of our mappings in Figure 9: (a) and (b) show the harmonic energy and quasi-conformal distortion convergence during the iteration, respectively. In (c), we perform experiments for genus 1 to further quantitatively test the robustness and validity of our mapping. The mapping from T_1 to T_2 discussed in the previous Section VII-A with a different initial cutting converges to the same result (green). The mapping from a torus T'_1 (different resolution with T_1) to T_2 is plotted in the blue curve. The inverse mapping ($T_2 \rightarrow T_1$, which has the same quasi-conformality bound in optimum) is plotted in brown.

E. Connectivity Refinement

Since we only use the connectivity of the source mesh S_1 , geometry loss may happen in some areas due to under sampling, most likely in high curvature (e.g. sharp feature) areas on S_2 . In order to capture such geometric details, we simply apply an adaptive remeshing algorithm similar to [17]. We locally modify the connectivity of the mesh using edge split,

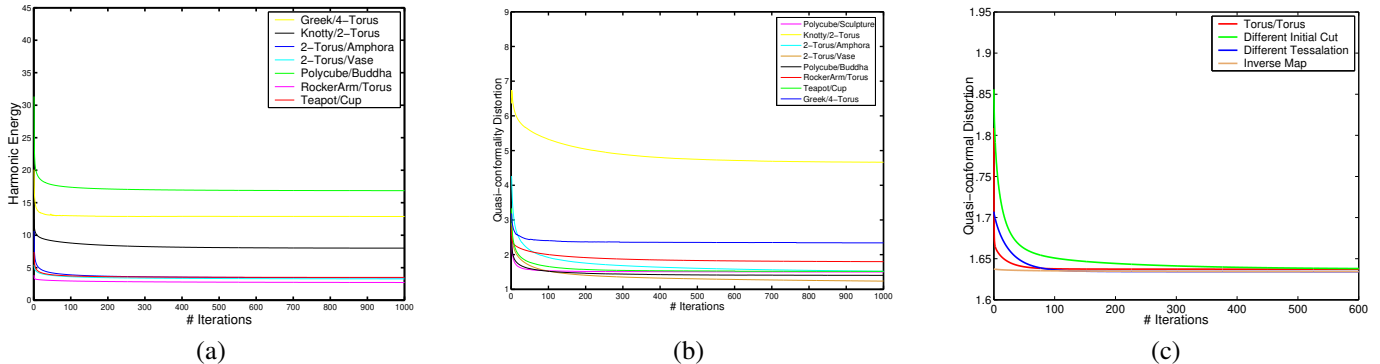


Fig. 9. Mapping Performance. (a) Harmonic energy during iterations. (b) and (c) Global quasi-conformal distortion (average of D') during iterations.

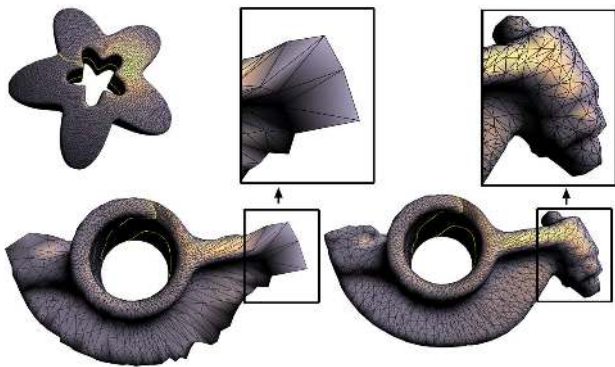


Fig. 10. Connectivity Refinement. Left column: the initial mapping from Star model to Rocker-Arm model; Right column: the refined connectivity.

guided by the following two simple error terms which capture the geometric proximity between S'_2 and S_2 : $E_{length}(e_{ij}) = |\phi(v_i) - \phi(v_j)|$, and $E_{norm}(e_{ij}) = [1 - N(v_i) \cdot N(v_j)]/2$, where $\phi : S_1 \rightarrow S_2$ and $N(v)$ is the normal of vertex v . The first term measures the length of an edge on S'_2 : longer edges are more likely to miss geometric details and we prefer splitting them early. The second term measures the normal deviation of the two vertices of the edge: a greater value implies that the edge crosses a more curved region or a region with sharp features. We iteratively split edges with large combined error. The new vertex generated by the edge split is then mapped back to the surface of the target mesh via the parameterized chart that covers this edge.

In Figure 10, we can see that the model created by mapping the ‘Star’ to the ‘Rocker Arm’, after being refined for 10 iterations, approximates the geometry of the target mesh much better: the left column is the initial mapping while the right column shows the refined connectivity. The number of vertices only increases by a fraction of 11.04%. Our simple error metric is easy to implement as we do not have to maintain the inverse map from S_2 to S_1 in this case.

VIII. RESULTS AND APPLICATIONS

A. Texture Transfer and Mapping Visualization

We need an effective way to clearly visualize a mapping between two surfaces because showing region correspondence




Fig. 11. Visualization of Surface Mapping between a Teapot Model and a Cup Model.



Fig. 12. Texture Transfer using the Global Optimized Surface Mapping.

as well as the distortion are challenging. We use a texture with the color band marks embedded in coordinate lines to aid in this visualization. The texture is first mapped onto the source model, each vertex on the source surface has its ‘UV’ coordinates. When vertices are mapped to the target surface, their ‘UV’ coordinates are carried. In this way, texture mapping on the source surface is transferred onto the target surface, the color bands on the target surface visualize the region correspondence, and the perpendicularity of the checker board or coordinate lines shows the angular distortion. Figure 11 visualizes mapping effect from genus-2 a teapot model to a cup model.

Texture as well as material transfer is straightforward as an application of our mapping. We show an example in Figure 12, which transfers the texture from the amphora model to the vase model. Since our mapping has the minimized distortions and global smoothness, such a transplant is physically natural, which potentially provides a powerful tool for reusing or

transferring information such as material properties, BRDF, etc. between models towards information integration.

B. Shape Morphing

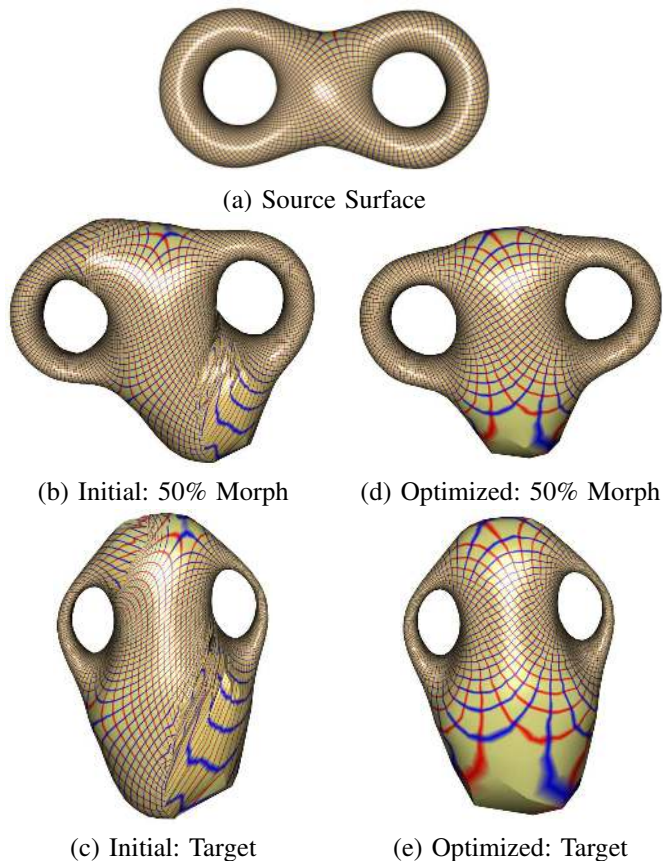


Fig. 13. Optimized Mapping for more Natural Morph. The source surface is shown in (a). If the initial map is used, the Morph generated is depicted in the left column: (b) shows the 50% morph, (c) shows the map on the target surface. When the surface map is optimized using our algorithm, the result is shown in the right column ((d) and (e)).

Another intuitive way to visualize mapping and to evaluate its distortion is via a morphing sequence. The behavior of the morph can be an intuitive visual judgement on the mapping quality. Figure 13 shows an example. The initial mapping, as indicated previously, is created by the technique of [1]. Based on initial map, we can conduct linear interpolation and generate the morph as shown in the left column. The generated sequence is obviously not attractive. We then optimize the surface map, and regenerate the morph. As shown in the right column, the new morph sequence demonstrates symmetric deformation and is visually much more smooth and pleasing.

Rigorously speaking, the morphing sequence generated by mapping with lower distortions means that the deformation sequence is closer to the ‘geodesic’ in the space of shapes, minimizes unnecessary distortion during the interpolation of shapes, thus provides better visualization results.

In graphics applications, shape morphing is widely studied as a direct application for surface mapping. Users usually want to have control on the morphing via feature or constraint points. To achieve this goal, as indicated previously in Section VII-C, we can use the existing techniques for feature

alignments during the initial mapping process; then we should keep this correspondence during the afterward optimization.

C. Canonical Mapping from Surfaces to Simplified Domains




Fig. 14. Mapping from a Polycube to the Happy Buddha Model. From left to right, we visualize the texture on the Polycube, the transferred texture on the Buddha by the initial mapping, and the transferred texture by the final mapping.




Fig. 15. Optimized Surface Mapping from the Greek Model to the 4-Torus. The left column shows the front and back of the Greek model with its texture; the right column shows the front and back of the target surface (4-Torus), respectively, with texture transferred by our mapping; the middle column shows the 50% morph from the Greek to the 4-Torus under our mapping.


Our method conveniently creates canonical mappings from arbitrary surfaces to simplified domains with globally optimized distortions. The canonical domain can be polycubes [40], so that graphics processing such as parameterization with lower distortion, polycube spline generation, etc. can be applied based on our mapping. The domain can also be some canonical N -hole tori [20], so that topologically equivalent shapes can be processed or analyzed on this smooth common domain.

In Figure 14, we visualize the polycube map for the genus-6 Buddha model. Our method successfully deforms arbitrarily built initial map with severe distortion to a global optimum. In Figure 15, we show our mapping from the genus-4 Greek model to a canonical 4-torus.

Our method has an important advantage over direct projection methods of computing polycube map such as [40] in that our method is intrinsic. Therefore it is more robust, invariant with models’ spatial positions and sizes. Furthermore, when

shapes are with complicated topology and geometry, or the source surface has great difference with the target surface (for example, Greek and Torus as shown above), direct projection method is highly error-prone but our method can robustly handle it.

D. Shape Matching and Comparison



1	2	3	4	5	6	7
2	0	3.59	22.72	20.81	59.70	19.43
3		0	21.99	21.29	59.38	18.72
4			0	10.98	39.66	19.65
5				0	44.45	16.91
6					0	32.89
7						0

Fig. 17. Shape Comparison using Conformal Representation. The first rows show all shapes to be compared. The second row and the first left column are their indices. The table has the symmetry property, and the numbers measure the distance between models in a pairwise manner.

Our optimal surface mapping creates global, low angular-distortion correspondence between two models. With such a non-rigid registration, we can easily match two shapes and clearly visualize their difference distributions for potential subsequent analysis purpose.

Conformal Representation. A natural way to characterize the matching between two surfaces is called *conformal representation* [41]. According to [41], when a surface is mapped onto a target surface, if the resultant conformal representation is fixed, the original source surface is rigidly determined. The conformal representation contains two terms: mean curvature H , and *conformal factor* λ . The conformal factor λ of a point p under a mapping f represents the local area change, i.e., the stretching of the map. Discretely, if we denote the area of one ring neighbor of p as $A(p)$, and the area of one ring neighbor of $f(p)$ on the target surface as $A(f(p))$. $\lambda(p)$ can be approximated by the ratio of $A(f(p))$ over $A(p)$. In our work, although our surface map is not fully conformal (according to Riemannian geometry, between most high genus models, these kinds of conformal maps do not exist), our global optimization aims to relieve angle distortions. Thus the (H, λ) defined on our map is a well approximated and meaningful representation.

Shape Matching. In Figure 16, we visualize our surface matching between a torus and a Rocker Arm model using the above conformal representation. (a) and (b) color-code the mean curvature distributions of Rocker Arm and Torus, respectively. We color-code the mean curvature difference in (c) and the stretching factor distribution in (d). The color-coding of two terms of conformal representation shows us where and how much the two surfaces are intrinsically different in a visually meaningful way. Since the globally integrated matching energy is smaller when the mapping is with lower stretching/distortions, our optimized surface mapping provides a great registration for the above mechanism. On the other hand, the registration by our mapping, with global smoothness and low distortion properties, can be used as a preprocessing

step for various other matching techniques. It serves as a general shape registration and visualization tool.

Shape Comparison and Retrieval. Given many shapes in database, we can match and compare them via canonical domains. This provides an efficient and geometrically meaningful way to measure their differences. Here we perform an experiment on a database containing 6 different genus-2 geometric shapes: Vase, Amphora, Teapot, Cup, Feline, and Cube. We use a two-hole torus as the canonical domain for all these genus two surfaces. We first compute mappings between these surfaces and the 2-torus domain, and then pairwise compare these surfaces via the domain using matching energy defined by the conformal representation:

$$E(S_1, S_2) = \int_{p \in T} \|\lambda_1(p) - \lambda_2(p)\|^2 + \beta \|H_1(p) - H_2(p)\|^2 dp,$$

where S_1 and S_2 are two shapes being compared, T is the canonical torus domain, λ is the conformal factor, and H is the mean curvature. In Figure 17, we can see the models in the first row. The matching energies, used as their distance, are shown in the table. Since the symmetry of the distance is obviously preserved, we only show the upper-right part of the table.

E. Algorithm Performance

Our optimization is an iterative algorithm; the total number of iteration steps is controlled by a user-defined threshold. In Figure 9, we set the threshold of quasi-conformality to be $1e - 6$; in real applications, we can use lower precisions. We perform our algorithm on a MS Windows XP PC with dual Intel Xeon 2.6GHz CPUs, 2GB RAM. The one-iteration running time for most real examples we presented in this paper are shown in the following runtime table.

Models(S_1/S_2)	Genus	Ver #	Time
2-Torus/Vase	2	3.5k/5k	0.31s
Amphora/Vase	2	10k/5k	1.4s
RockerArm/Torus	1	15k/14.4k	6.88s
Teapot/Cup	2	7.5k/10k	0.95s
Polycube/Sculpture	3	3.5k/7K	0.49s
Greek/4-Torus	4	14.8k/10K	5.07s
Polycube/Buddha	6	18k/13.3k	10.23s

F. Comparison with Existing Work

Compared with other state-of-the-art techniques, our approach has several key improvements.

First, our surface mapping framework is based on rigorous mathematical foundation and analysis, unlike most current methods that only guarantee to reach local optima, our method **globally minimizes the stretching** and converges to a **unique** result.

Second, current techniques rely upon large amount of user intervention for mapping surfaces with non-trivial topology. For example, mapping procedures need base mesh design ([12], [13]) or a large number of user-specified landmarks ([17], [18]). In contrast, our framework does not depend on user's involvement, and is fully **automatic**.




Fig. 16. Surface Registration for Matching (Torus vs Rocker Arm model). (a) Mean curvature distributions of Rocker Arm (red represents the maximum while blue represents the minimum); (2) Mean curvature distributions of Torus; (c) Mean curvature difference distributions, visualized on Torus; (d) Conformal stretching factor, visualized on Torus.

Third, since the existing methods follow the general principle to slice the surface open into subregions, the initial segmentation directly determines the mapping result. In contrast, our method, because of its global uniqueness, is not controlled by the quality of initial mappings; therefore, it is much more general than other existing techniques.

Fourth, existing work primarily focuses on low genus surfaces and few takes the homotopy types into account. Carner et al. [1] also targeted on high genus surfaces, and they studied the mapping with different homotopy classes. However, topology information is the primary information they used for mapping computation and therefore their stretching energy is not optimized. In our current work, the comparison between initial and final mapping shows the great improvement from the initial mapping generated by their method to our globally optimized result. This can be easily visualized through our optimization procedure in the accompanying video.

IX. CONCLUSION

This paper has documented our new method for computing a globally optimal map between surfaces of non-trivial topology and demonstrated many valuable applications. Based on the mathematical advances in computing the uniformization metric using intrinsic geometric structure, we can globally perform heat diffusion to alleviate the stretching and the average angle distortion of the map as much as possible. As we discussed in Section VIII-F, our algorithm have many key advantages over existing work.

Our mapping algorithm can also serve as a ubiquitous tool for a wider range of applications such as shape registration, morphing, matching, comparison, and spline surface construction over generalized domains. We would like to apply our mapping framework in more challenging research topics such as deformable model tracking, animation transfer, etc.

APPENDIX

CONVERGENCE, ONE-TO-ONE, AND UNIQUENESS

In this appendix, we will show our algorithm converges to a globally unique one-to-one map with the minimal harmonic energy under the uniformization metric. The pipeline is: (1)

we demonstrate the **existence** of the harmonic map between given surfaces S_1 and S_2 with same topology; (2) we show if the final map we get is harmonic, then it is one-to-one, and **globally unique**; (3) we will show our algorithm does **converge** to such a unique mapping with minimal harmonic energy under uniformization metric.

Existence. Given two high genus surfaces S_1 and S_2 with same non-trivial topology. The existence of the harmonic map is guaranteed by the following theorem

Theorem 1: Suppose that S_1 and S_2 are compact surfaces without boundary and that $h : S_1 \rightarrow S_2$ is a diffeomorphism. Then there exists a harmonic diffeomorphism $f : S_1 \rightarrow S_2$ isotopic to h . Furthermore, f is of least energy among all diffeomorphisms isotopic to h .

Detailed proof can be found in [35], page 176. Since our initial map is constructed as a diffeomorphism between S_1 and S_2 , the existence of harmonic map is guaranteed.

One-to-one and uniqueness. We show if the final map is harmonic, then it is a **diffeomorphism** (one-to-one and differentiable) and has the global uniqueness. We prove in the third step that we do reach a harmonic map. The following theorem guarantees the harmonic map calculated in our algorithm is a diffeomorphism.

Theorem 2: Let $f : S_1 \rightarrow S_2$ be a harmonic map between closed oriented surfaces of the same genus with degree equals ± 1 . And $K_{S_2} \leq 0$, then f is a diffeomorphism.

Detailed proof can be found in [35], page 187, or [2], page 15. In our algorithm, the initial map is constructed by matching the fundamental polygons of S_1 and S_2 . Therefore, each point on S_2 has a unique pre-image on S_1 , hence, the degree of the initial map is 1. The Gaussian curvature of the target surface is 0 (for genus-1 surfaces) or -1 (for high genus surfaces), therefore if f is harmonic, then f is one-to-one and differentiable.

The following theorem postulates the uniqueness of the map.

Corollary 1: Let u_1, u_2 be harmonic maps $M \rightarrow N$ of degree one between compact surfaces without boundaries, with genus greater than one, where $K_N \equiv -1$. If u_1 and u_2 are homotopic to each other, then $u_1 = u_2$.

The detailed proof can be found in [42], page 144 and [2], page 16. In our algorithm, the homotopy class of the map is determined by the way to match the fundamental polygons. The map is harmonic, the curvature on the target surface is -1 , therefore, the harmonic map is unique.

For genus-one surfaces, their uniformization metric is flat, which can be lifted to its universal covering space. The universal cover can be embedded on the plane isometrically. The fundamental polygons are parallelograms. A harmonic map between two genus-one surfaces with their flat uniformization metrics induces a map between their universal covering spaces, which is an affine transformation from the plane to itself. The affine transformation maps the fundamental polygon of the source surface to that of the target surface. Therefore, harmonic maps in a homotopy class only differ by a translation. Each one is the equally optimal result.

Convergence. We prove our algorithm converge to a harmonic map. Harmonic energy of a surface map is non-negative, namely, it has lower bound. Our relaxation process reduces harmonic energy monotonically; therefore, it converges to a critical point of the harmonic energy, which by definition is a harmonic map. As the aforementioned theorems show, there is no local minimum, and this critical point is globally unique. Therefore, our method converges to the global unique harmonic map, and it is one-to-one and differentiable.

For genus-one surfaces, this convergence proof also applies, and all the minima are globally equal and globally optimal. Our minimization process will converge to one of them.

REFERENCES

- [1] C. Carner, M. Jin, X. Gu, and H. Qin. Topology-driven surface mappings with robust feature alignment. In *IEEE Visualization*, pages 543–550, 2005.
- [2] R. Schoen and S. T. Yau. *Lectures on harmonic maps*. International Press, Cambridge, MA, USA, 1997.
- [3] J. Kent, W. Carlson, and R. Parent. Shape transformation for polyhedral objects. In *SIGGRAPH '92*, pages 47–54, New York, NY, USA, 1992. ACM Press.
- [4] T. Kanai, H. Suzuki, and F. Kimura. Three-dimensional geometric metamorphosis based on harmonic maps. *The Visual Comput.*, 14(4):166–176, 1998.
- [5] M. Alexa. Merging polyhedral shapes with scattered features. In *SMI '99: Proceedings of the International Conference on Shape Modeling and Applications*, pages 202–210, 1999.
- [6] A. Asirvatham, E. Praun, and H. Hoppe. Consistent spherical parameterization. In *Computer Graphics and Geometric Modeling (CGGM) 2005 Workshop*.
- [7] S. Haker, S. Angenent, A. Tannenbaum, R. Kikinis, G. Sapiro, and M. Halle. Conformal surface parameterization for texture mapping. *IEEE Trans. Vis. Comput. Graph.*, 6(2):181–189, 2000.
- [8] C. Gotsman, X. Gu, and A. Sheffer. Fundamentals of spherical parameterization for 3d meshes. *ACM Trans. Graph.*, 22(3):358–363, 2003.
- [9] E. Praun and H. Hoppe. Spherical parameterization and remeshing. In *SIGGRAPH '03*, pages 340–349, 2003.
- [10] M. S. Floater and K. Hormann. Surface parameterization: a tutorial and survey. In *Advances in Multiresolution for Geometric Modelling, Mathematics and Visualization*, pages 157–186, 2005.
- [11] R. Parent. Shape transformation by boundary representation interpolation: A recursive approach to establishing face correspondences. *The Journal of Visualization and Computer Animation*, 3:219–239, 1992.
- [12] D. DeCarlo and J. Gallier. Topological evolution of surfaces. In *GI '96: Proceedings of the conference on Graphics interface '96*, pages 194–203, Toronto, Ont., Canada, Canada, 1996. Canadian Information Processing Society.
- [13] A. Gregory, A. State, M. Lin, D. Manocha, and M. Livingston. Feature-based surface decomposition for correspondence and morphing between polyhedra. In *CA '98: Proceedings of the Computer Animation*, page 64, Washington, DC, USA, 1998. IEEE Computer Society.
- [14] A.W.F. Lee, D. Dobkin, W. Sweldens, and P. Schröder. Multiresolution mesh morphing. In *Proc. SIGGRAPH*, pages 343–350, 1999.
- [15] T. Michikawa, T. Kanai, M. Fujita, and H. Chiyokura. Multiresolution interpolation meshes. In *Proc. Pacific Graphics*, pages 60–69, 2001.
- [16] E. Praun, W. Sweldens, and P. Schröder. Consistent mesh parameterizations. In *Proc. SIGGRAPH*, pages 179–184, 2001.
- [17] V. Kraevoy and A. Sheffer. Cross-parameterization and compatible remeshing of 3d models. *ACM Trans. Graph.*, 23(3):861–869, 2004.
- [18] J. Schreiner, A. Asirvatham, E. Praun, and H. Hoppe. Inter-surface mapping. *SIGGRAPH.*, 23(3):870–877, 2004.
- [19] A. Khodakovsky, N. Litke, and P. Schröder. Globally smooth parameterizations with low distortion. *ACM Trans. Graph.*, 22(3):350–357, 2003.
- [20] C. Grimm and J. Hughes. Parameterizing n-holed tori. In *Mathematics of Surfaces X*, pages 14–29, September 2003.
- [21] K. Rose A. Sheffer, E. Praun. Mesh parameterization methods and their applications. *Foundations and Trends in Computer Graphics and Vision*, 2006.
- [22] U. Pinkall and K. Polthier. Computing discrete minimal surfaces and their conjugate. In *Experimental Mathematics*, volume 2, pages 15–36, 1993.
- [23] M. Eck, T. DeRose, T. Duchamp, H. Hoppe, M. Lounsbery, and W. Stuetzle. Multiresolution analysis of arbitrary meshes. In *SIGGRAPH*, pages 173–182, 1995.
- [24] M. Desbrun, M. Meyer, and P. Alliez. Intrinsic parameterizations of surface meshes. *Comput. Graph. Forum*, 21(3):209–218, 2002.
- [25] B. Lévy, S. Petitjean, N. Ray, and J. Maillot. Least squares conformal maps for automatic texture atlas generation. In *SIGGRAPH*, pages 362–371, 2002.
- [26] W. Thurston. *Geometry and topology of 3-manifolds*. Princeton lecture notes, 1976.
- [27] K. Stephenson. *Introduction to circle packing*. Cambridge University Press, 2005.
- [28] A. I. Bobenko and B. A. Springborn. Variational principles for circle patterns and koebe's theorem. *Trans. Amer. Math. Soc.*, 356:659–689, 2004.
- [29] L. Kharevych, B. Springborn, and P. Schröder. Discrete conformal mappings via circle patterns. *ACM Trans. Graph.*, 25(2):412–438, 2006.
- [30] R. S. Hamilton. The ricci flow on surfaces. *Contemp. Math.*, 71:237–262, 1988.
- [31] B. Chow and F. Luo. Combinatorial ricci flows on surfaces. *J. Differ. Geom.*, 63:97–129, 2003.
- [32] M. Jin, F. Luo, and X. Gu. Computing surface hyperbolic structure and real projective structure. In *SPM '06: Proc. ACM symposium on Solid and physical modeling*, pages 105–116, 2006.
- [33] R. Schoen and S-T. Yau. On univalent harmonic maps between surfaces. *Inventiones Mathematicae*, 44:265–278, 1978.
- [34] J. Erickson and K. Whittlesey. Greedy optimal homotopy and homology generators. In *SODA '05: ACM-SIAM symposium on Discrete algorithms*, pages 1038–1046, 2005.
- [35] J. Jost. *Compact Riemann Surfaces*. Springer-Verlag, 2002.
- [36] J. Vorsatz, Ch. Rössl, and H.-P. Seidel. Dynamic remeshing and applications. In *SM '03: Proceedings of the eighth ACM symposium on Solid modeling and applications*, pages 167–175, 2003.
- [37] V. Surazhsky, P. Alliez, and C. Gotsman. Isotropic remeshing of surfaces: a local parameterization approach. In *Proceedings of 12th International Meshing Roundtable*, 2003.
- [38] P. V. Sander, J. Snyder, S. J. Gortler, and H. Hoppe. Texture mapping progressive meshes. In *SIGGRAPH '01*, pages 409–416, 2001.
- [39] T. K. Dey, K. Li, and J. Sun. On computing handle and tunnel loops. In *IEEE Proc. NASAGEM 07*, page to appear.
- [40] M. Tarini, K. Hormann, P. Cignoni, and C. Montani. Polycube-maps. *ACM Trans. Graph.*, 23(3):853–860, 2004.
- [41] X. Gu, Y. Wang, T. Chan, P. Thompson, and S. T. Yau. Genus zero surface conformal mapping and its application to brain surface mapping. *IEEE Trans. Med. Imaging*, 23(8):949–958, 2004.
- [42] X. Gu and S. T. Yau. *Computational Conformal Geometry*. International Press, 2007.