

Computing a Metric on the Shape Space of Tori

YUAN LUO, University of California, Davis, USA

YANWEN LUO, University of Science and Technology Beijing, China

PATRICE KOEHL, University of California, Davis, USA

JOEL HASS, University of California, Davis, USA

This paper explores shape comparison for genus-one surfaces, namely tori, through an intrinsic metric on their shape space. While metric structures for 2-spheres are well established, extending them to tori is substantially harder, due to their nontrivial topology and the absence of conformal diffeomorphisms in general. Drawing on Teichmüller theory, we formulate an algorithm that minimizes a distortion energy over diffeomorphisms within a canonical homotopy class to compute a metric on the shape space of Riemannian tori. This produces both a pairwise distance and an optimal, landmark-free correspondence. On the computational side, we discretize this energy on triangulated tori and optimize it directly in a pair of fundamental domains. A key theoretical contribution is a convergence theorem showing that the discrete energy converges to the smooth metric energy under mesh refinement. Experiments on flat tori, tori of revolution, and real-world scanned surfaces show that the method captures natural geometric variation and provides a robust tool for genus-one shape comparison.

CCS Concepts: • **Computing methodologies** → **Shape analysis**; • **Theory of computation** → **Computational geometry**.

Additional Key Words and Phrases: Shape Comparison, Metric of Surfaces

ACM Reference Format:

Yuan Luo, Yanwen Luo, Patrice Koehl, and Joel Hass. 2026. Computing a Metric on the Shape Space of Tori. *ACM Trans. Graph.* 45, 6, Article 234 (December 2026), 16 pages. <https://doi.org/10.1145/3842520>

1 Introduction

Comparing the shapes of surfaces is a fundamental problem in computer graphics, with broad impact on shape registration, retrieval, and blending. Central to this problem is the definition of a meaningful distance or metric between shapes. A metric on a shape space assigns a real value to each pair of shapes that quantifies their geometric dissimilarity while satisfying positivity, symmetry, and the triangle inequality. Such a structure enables pairwise comparison and supports distance-based analysis methods, including clustering, low-dimensional embedding, and topological data analysis. As an intrinsic metric, it has the advantage of being able to detect similar objects that exhibit extrinsic deformation, such as flexible biomolecules or objects in different poses.

In this work, we introduce a new computational framework for generating a stable and consistent metric on the shape space of genus-one surfaces that is derived as the minimizer of an energy

Authors' Contact Information: Yuan Luo, University of California, Davis, Davis, USA, luoyuan9809@gmail.com; Yanwen Luo, University of Science and Technology Beijing, Beijing, China, yanwen.luo@outlook.com; Patrice Koehl, University of California, Davis, Davis, USA, koehl@cs.ucdavis.edu; Joel Hass, University of California, Davis, Davis, USA, jhass@ucdavis.edu.



This work is licensed under a Creative Commons Attribution 4.0 International License. © 2026 Copyright held by the owner/author(s). ACM 1557-7368/2026/12-ART234 <https://doi.org/10.1145/3842520>




Fig. 1. Our method generates a pairwise distance matrix for five tori of revolution with increasing minor radii, representing approximately a geodesic in the shape space.

function. Our method is landmark-free, avoiding the need for manually or automatically selected correspondences, which are often expensive to generate or compute and sensitive to noise. A central theoretical contribution is a convergence theorem showing that the discrete energy optimized in our algorithm converges to the smooth metric energy under mesh refinement, thereby connecting the numerical scheme to the continuum shape-space metric. When applied to a synthetic family of tori of revolution with increasing minor radii, the resulting pairwise distance matrix yields a linear embedding of the family into the real line, as seen in Figure 1. This correspondence exhibits a natural alignment consistent with geometric intuition.

A *metric* on a space X is a function $d : X \times X \rightarrow \mathbb{R}$ satisfying three properties for all $x, y, z \in X$:

- (1) **POSITIVITY**: $d(x, y) \geq 0$ with equality if and only if $x = y$;
- (2) **SYMMETRY**: $d(x, y) = d(y, x)$;
- (3) **TRIANGLE INEQUALITY**: $d(x, z) \leq d(x, y) + d(y, z)$.

Collectively, these properties make the metric structure suitable for the shape comparison problem: symmetry ensures independence from the order of measurement, positivity guarantees that shapes

with different geometries are distinguishable, and the triangle inequality provides robustness to noise, implying that two nearby objects have approximately the same distance to a third object. Each of these properties plays an important role in applications.

Hass and Koehl [2017] constructed a metric on genus-zero shapes by finding optimal conformal diffeomorphisms between two shapes. This has been applied in biological shape analysis, including measuring the roundness of proteins [Hass and Koehl 2014] and generating the evolutionary trees of primates based on the shapes of their bones [Koehl and Hass 2015].

Similar tools for higher-genus surfaces remain scarce due to theoretical and computational difficulties. There is no conformal diffeomorphism between an arbitrary pair of tori, and this forces a search for an optimal correspondence over the full group of diffeomorphisms. Cutting a torus along a pair of loops to produce a planar domain requires choosing a pair of curves that generate homology classes and intersect at one point. Different choices can lead to different correspondences and distortions, making consistent comparisons more challenging.

Nevertheless, a metric structure on the shape space of tori can be constructed using the energy functional over diffeomorphisms defined formally in Equation 1. Its two terms measure area and angle distortion. As established in [Luo 2021], minimizing it over diffeomorphisms in a fixed homotopy class defines a metric on the shape space of Riemannian tori. For flat tori, the minimizer is the Teichmüller map, and the resulting metric is the classical Teichmüller metric; see Section 4.2 for the details.

The choice of this energy functional is motivated by both theoretical and practical considerations. Somewhat surprisingly, finding an energy that satisfies the metric property is highly restrictive. It can be proven under some natural assumptions on the energy functional that if we restrict to a conformal class of metrics, the first term in Equation 1 is the unique metric structure up to scaling. Therefore, it is natural to combine it with the Teichmüller metric, which measures distance between conformal classes, extending the Teichmüller metric from the space of flat tori to arbitrary genus-one shapes. From a practical perspective, the two-term structure allows us to separately control and interpret area and angle distortions, which helps explain geometric differences between shapes.

In implementing the computation of this energy and metric we work with discretized surfaces. To address the computational challenges, we introduce a framework that directly optimizes vertex positions in a fundamental domain while maintaining a homotopy class. We discretize the energy functional by computing the singular values λ_1 and λ_2 of the differential df for each triangular face, enabling efficient evaluations of area and angle distortions. Our approach can be defined briefly as follows. Starting with two tori defined by triangulations of their surfaces, we parametrize them using the unit square $[0, 1]^2$. We optimize a map between their parameterizations using gradient descent. In Figure 2 the bottom row shows the map $\hat{f}$ between fundamental domains, and the composition $f = \phi_2 \circ \hat{f} \circ \phi_1^{-1}$ gives the map between tori. By averaging gradient vectors at boundary vertices that are identified by the quotient map,




Fig. 2. An example of a map between two tori. Two maps ϕ_1 and ϕ_2 are parameterizations of two mugs from the unit square. The red and blue loops are the two canonical homology generators. The map $\hat{f}_i$ induces the map f between the two tori.

we ensure that the consistency of boundary identifications is maintained throughout optimization. This process preserves the initial homotopy class.

Contributions. Our main contributions are as follows.

- We formulate genus-one shape comparison as the computation of a metric on shape space, rather than as a registration problem. The resulting metric has the desirable properties of positivity, symmetry, and the triangle inequality.
- We prove that under mesh refinement, the discrete energy computed by our algorithm converges to the smooth metric energy.
- We introduce a direct optimization framework in torus fundamental domains that preserves the initial homotopy class. This avoids the PDE-based machinery used in registration methods.
- We validate the method on flat tori, tori of revolution, and real-world scanned genus-one surfaces, showing stable behavior under measurement and useful discrimination between shapes.

Related Work. While the literature on shape comparison is extensive, especially for planar regions and the 2-sphere, much less is known about distances between high-genus surfaces. We therefore focus on parametrization and registration methods, which are the closest alternatives to metric-based comparison in our setting.

Computing surface parameterizations is a fundamental problem in computational geometry [Floater and Hormann 2005]. Parametrization is difficult for high genus surfaces due to their complicated topology and intrinsic geometry, often involving cutting surfaces to disks [Campen et al. 2019]. By the celebrated uniformization theorem, any torus with a Riemannian metric is conformally equivalent to a flat torus, locally parametrized by the Euclidean plane, and

any surface with genus at least two is conformally equivalent to a hyperbolic surface, locally parametrized by the hyperbolic plane. Algorithms have been developed to compute these conformal maps [Gu and Yau 2003; Jin et al. 2008]. On the other hand, discrete harmonic maps provide a more flexible method to parametrize high genus surfaces locally from the Euclidean plane [Bright et al. 2017; Gortler et al. 2006; Wang 2025] or orbifolds [Aigerman et al. 2017; Aigerman and Lipman 2015, 2016]. These results generalize the original idea of Tutte’s embedding via barycenters of interior vertices with convex boundary conditions [Tutte 1963].

A variational method [Lui and Wen 2014; Wen et al. 2015] registered high-genus surfaces by minimizing an energy involving geometric features such as mean and Gaussian curvature. It computed an optimal correspondence by solving PDEs on a covering space. The Teichmüller extremal mapping was applied to genus-one surfaces for landmark-driven torus registration [Lam et al. 2014], minimizing harmonic energy through quasi-conformal iteration. This was used for medical image analysis [Lam et al. 2015]. While it provides bounded-distortion mappings, it targets specific registration tasks that require landmarks. Dirichlet-energy minimization under uniformization metrics has been used to compute correspondences between surfaces of arbitrary topology [Li et al. 2008], with a recent structure-preserving discretization developed for closed surfaces of genus at least two [Zhu et al. 2025]. These methods return optimized harmonic maps in prescribed homotopy classes. In contrast to our energy, none of the above methods provide a distance between surfaces satisfying the symmetry and triangle inequality properties.

Gromov–Wasserstein (GW) comparison treats shapes as metric-measure spaces and computes soft probabilistic correspondences [Mémoli 2011; Solomon et al. 2016]. The method operates on within-shape distance matrices and probability weights, so the inputs need not share the same number of samples, combinatorial structure, or topology. Our experiments later show that entropic GW detects geometric changes but is sensitive to sampling and subdivision.

Several Riemannian approaches can compare shapes by minimizing deformation energy along paths [Kilian et al. 2007]. Discrete shell metrics derive this energy from membrane and bending deformation [Heeren et al. 2014, 2020]. Second-order Sobolev metrics provide a separate path-energy model on immersed surfaces [Hartman et al. 2023]. These methods involve both intrinsic stretching and extrinsic deformation, but generally require meshes of the same triangulation as the input.

We provide an open-source implementation of our method at <https://github.com/YuanL12/tori-comparison>.

2 Metrics in Tori Shape Space

2.1 The Smooth Case

We summarize our underlying mathematical theory in the setting of smooth surfaces. Our approach to comparing a pair of genus-one surfaces is based on measuring the distortion of an optimal diffeomorphism between them. Because we compare shape rather than scale, we normalize all surfaces to have unit area. The framework can be extended to shapes with scale when needed.

Given two tori with Riemannian metrics (S, g) and (S, h) , we define an energy functional that quantifies how much a diffeomorphism $f : (S, g) \rightarrow (S, h)$ distorts the geometry. The energy of a diffeomorphism $f : (S, g) \rightarrow (S, h)$ is defined to be

$$E(f) = \sqrt{\int_S (\sqrt{\lambda_1 \lambda_2} - 1)^2 dA_g} + \frac{1}{2} \left\| \log \frac{\lambda_1}{\lambda_2} \right\|_\infty. \quad (1)$$

Here λ_1 and λ_2 are the singular values of the differential df at each point. These are the stretch factors of f at each point, computed with respect to orthonormal bases in the tangent spaces equipped with Riemannian metrics.

The energy $E(f)$ consists of two terms, each measuring a different aspect of geometric distortion:

- **Area distortion term:** The first term $\sqrt{\int_S (\sqrt{\lambda_1 \lambda_2} - 1)^2 dA_g}$ measures how much the diffeomorphism distorts area. The quantity $\lambda_1 \lambda_2$ represents the Jacobian determinant, equal to the local area scaling factor: when $\lambda_1 \lambda_2 = 1$, the map is area-preserving at that point. This term integrates the squared deviation from pointwise area preservation over the entire surface, giving a global measure of area distortion.
- **Angle distortion term:** The second term $\frac{1}{2} \left\| \log \frac{\lambda_1}{\lambda_2} \right\|_\infty$ measures how much the diffeomorphism distorts angles pointwise. The ratio $\frac{\lambda_1}{\lambda_2}$ captures the local anisotropy: at points where $\frac{\lambda_1}{\lambda_2} = 1$, the map is conformal, namely, angle-preserving. The logarithm converts this ratio into an additive measure, and taking the L^∞ norm captures the largest angle distortion across the surface.

A marking records how a torus is identified with a fixed reference torus. Let X be a reference topological torus. A *marking* of a torus S is an orientation-preserving homeomorphism $\psi : X \rightarrow S$, and a marked Riemannian torus is a triple (S, ψ, g) . For two marked Riemannian tori (S_1, ψ_1, g_1) and (S_2, ψ_2, g_2) , we minimize the energy in Equation 1 over diffeomorphisms f satisfying the homotopy condition $f \circ \psi_1 \simeq \psi_2$. In particular, the markings select the homotopy class of $\psi_2 \circ \psi_1^{-1} : S_1 \rightarrow S_2$. For any torus embedded in $\mathbb{R}^3$, there exists a canonical choice of meridian and longitude. In the one-point compactification of $\mathbb{R}^3$, namely the 3-sphere S^3 , the complement of an embedded torus consists of an open solid torus K and an open manifold containing the point at infinity. The meridian bounds an embedded disk in K , and the longitude is homotopic to the core curve inside K , which generates its fundamental group. Given oriented meridian-longitude bases for two marked tori, there exists a unique homotopy class of homeomorphisms sending the oriented meridian and longitude of the first torus to those of the second.

The distance between two marked Riemannian tori is defined as the infimum of the energy over this class:

$$d((S_1, \psi_1, g_1), (S_2, \psi_2, g_2)) = \inf_{f \simeq \psi_2 \circ \psi_1^{-1}} E(f), \quad (2)$$

where the infimum is taken over diffeomorphisms $f : (S_1, g_1) \rightarrow (S_2, g_2)$ in the selected homotopy class. It was established in [Luo 2021] that this defines a metric structure on the space of marked Riemannian tori.

2.2 The Discrete Case

A triangulation of a surface is a decomposition of the surface into faces Δ_{ijk} , where each face Δ_{ijk} is a triangle with vertices v_i, v_j , and v_k . The triangulated tori T_1 and T_2 may have different numbers of vertices and need not have compatible connectivity. The homeomorphism $f : T_1 \rightarrow T_2$ is not required to be simplicial: vertices of T_1 may map to arbitrary points on T_2 , and source edges are mapped to geodesic paths connecting their mapped endpoints. The energy function defined above can then be discretized as

$$E_D(f) = \sqrt{\sum_{\Delta_{ijk} \in T_1} \left(\sqrt{\lambda_1(\Delta_{ijk})\lambda_2(\Delta_{ijk})} - 1 \right)^2 A(\Delta_{ijk})} + \frac{1}{2} \max_{\Delta_{ijk} \in T_1} \left| \log \frac{\lambda_1(\Delta_{ijk})}{\lambda_2(\Delta_{ijk})} \right|,$$

where Δ_{ijk} is a face of T_1 and $A(\Delta_{ijk})$ is its area. The surface points $f(v_i)$, $f(v_j)$, and $f(v_k)$ on T_2 are represented by their face indices and barycentric coordinates. They form a geodesic triangle $f(\Delta_{ijk})$ in T_2 that can contain vertices of T_2 . We approximate this geodesic triangle by a Euclidean triangle with the same three pairwise intrinsic geodesic distances. We then approximate $f|_{\Delta_{ijk}}$ as an affine map with singular values $\lambda_1(\Delta_{ijk})$ and $\lambda_2(\Delta_{ijk})$. This approximation allows us to express $\lambda_1(\Delta_{ijk})$, $\lambda_2(\Delta_{ijk})$, and hence the discrete energy solely in terms of edge lengths on T_1 and geodesic lengths on T_2 .

Computationally, let L_1, L_2, L_3 be the edge lengths of the domain triangle Δ_{ijk} , with opposite angles α, β, γ . Let l_1, l_2, l_3 be the corresponding pairwise geodesic distances between $f(v_1), f(v_2), f(v_3)$; these lengths determine a Euclidean triangle. There is then a unique affine map from Δ_{ijk} to this Euclidean triangle. We define $\lambda_1(\Delta_{ijk})$ and $\lambda_2(\Delta_{ijk})$ to be the singular values of this affine map. Then set $D(\Delta_{ijk}) = \lambda_1^2 + \lambda_2^2$, $R(\Delta_{ijk}) = \lambda_1/\lambda_2$ and $P(\Delta_{ijk}) = \lambda_1\lambda_2$. We have

$$D(\Delta_{ijk}) = \frac{1}{2A} (l_1^2 \cot \alpha + l_2^2 \cot \beta + l_3^2 \cot \gamma)$$

and

$$P(\Delta_{ijk}) = \lambda_1\lambda_2 = \sqrt{\frac{c(c-l_1)(c-l_2)(c-l_3)}{C(C-L_1)(C-L_2)(C-L_3)}},$$

where A is the area of Δ_{ijk} , C is its semiperimeter, and c is the semiperimeter of the image triangle. The formula for D is given in Lemma 5 of [Pinkall and Polthier 1993], and P is the area ratio of the two triangles computed by Heron's formula. Then we can solve R by

$$\left(R + \frac{1}{R} \right) P = D \quad \Rightarrow \quad R = \frac{\frac{D}{P} + \sqrt{\left(\frac{D}{P} \right)^2 - 4}}{2}. \quad (3)$$

In summary, this gives a discrete energy depending only on edge lengths:

$$E_D(f) = \sqrt{\sum_{\Delta_{ijk} \in T_1} \left(\sqrt{P(\Delta_{ijk})} - 1 \right)^2 A(\Delta_{ijk})} + \frac{1}{2} \max_{\Delta_{ijk} \in T_1} \left| \log R(\Delta_{ijk}) \right|. \quad (4)$$

We now state the convergence theorem for the discrete energy and defer its proof to Section 5. The statement and proof do not rely on the genus-one assumption: they hold for any connected closed orientable Riemannian surface (S, g) .

Theorem 2.1. Let (S, g) and (S, h) be two Riemannian surfaces, and let $f : (S, g) \rightarrow (S, h)$ be a diffeomorphism. For any $\epsilon > 0$, there exists a $\delta > 0$ such that the induced discrete energy $E_D(f)$ above, computed from any geodesic triangulation on (S, g) with maximal edge length smaller than δ and minimum angle at least $\pi/5$, satisfies

$$|E(f) - E_D(f)| < \epsilon.$$

2.3 The Symmetrized Discrete Distance

In practice, we observe a small discrepancy between the minimized discrete energies $\min_{f:T_1 \rightarrow T_2} E_D(f)$ and $\min_{g:T_2 \rightarrow T_1} E_D(g)$. We do not expect the exact symmetry $\min_{f:T_1 \rightarrow T_2} E_D(f) = \min_{g:T_2 \rightarrow T_1} E_D(g)$. This discrepancy is not solely a floating-point effect: reversing the direction changes the number of vertices and faces as well as the area weights in E_D , so the two directions give different finite discretizations of the corresponding smooth maps. Real meshes may also violate the assumptions on angles and edge lengths in Theorem 2.1, while approximating geodesic triangles on triangular meshes by Euclidean comparison triangles introduces further discrepancies. The convergence theorem implies approximate symmetry under dense, shape-regular refinement of triangulations rather than exact symmetry at a fixed resolution. Section 4.5 examines this behavior experimentally. Thus, we define the symmetrized discrete distance between two triangulated meshes T_1 and T_2 as

$$d_E(T_1, T_2) = \frac{\min_{f:T_1 \rightarrow T_2} E_D(f) + \min_{g:T_2 \rightarrow T_1} E_D(g)}{2}. \quad (5)$$

The importance of symmetrization was elaborated in [Abulnaga et al. 2023].

3 Algorithm and Implementation Details

To minimize the discrete energy function $E_D(f)$ defined in Equation 4, we optimize the map $f : T_1 \rightarrow T_2$ by working on the fundamental domains rather than directly on the tori. Specifically, we parameterize both tori T_1 and T_2 to their respective fundamental domains F_1 and F_2 , then optimize a map $\hat{f} : F_1 \rightarrow F_2$ between these domains. The composition $\phi_2 \circ \hat{f} \circ \phi_1^{-1}$ gives us the desired map f between the tori, as illustrated in Figure 2.

Our optimization framework in Algorithm 1 can be summarized as follows:

- (1) Find piecewise linear parameterizations $\phi_1 : F_1 \rightarrow T_1$ and $\phi_2 : F_2 \rightarrow T_2$ for both tori, where F_1 and F_2 are fundamental domains and the parameterizations are chosen to give the desired homotopy class. See an example in Figure 2.
- (2) Initialize the map $\hat{f}_0 : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ as the identity map.
- (3) Iteratively optimize by computing the energy $E_D(f_i)$ and its gradient, then updating vertex positions via gradient descent while ensuring boundary consistency and injectivity constraints.

We now discuss these three key steps in detail.

3.1 Parameterization

For a torus T embedded in $\mathbb{R}^3$, we first find two canonical homology generators, the oriented longitude and meridian, using [Dey et al. 2013]. See the red and blue curves in Figure 2. These canonical choices provide a natural marking on any torus embedded in $\mathbb{R}^3$. We

Algorithm 1 Gradient Descent Optimization on Fundamental Domain**Input:** Two triangulated tori T_1 and T_2 **Output:** Optimized energy $E_D(f)$ with its minimizer $f = \phi_2 \circ \hat{f} \circ \phi_1^{-1}$

- 1: Initialize two piecewise linear parameterizations $\phi_1 : F_1 \rightarrow T_1$ and $\phi_2 : F_2 \rightarrow T_2$ from the unit square to the two tori T_1 and T_2 with square edges corresponding to canonical generators.
- 2: Store the attaching mapping $\mathcal{O}$ as a collection of vertex index sets, where each set $S \in \mathcal{O}$ contains either 4 or 2 vertices (the corner or the other boundary vertices)
- 3: Initialize $\hat{f}_0 = \text{Id}$
- 4: **for** iteration $i = 0, 1, 2, \dots$ until convergence **do**
- 5: Compute the energy $E_D = E_D(f_i)$ for the map $f_i = \phi_2 \circ \hat{f}_i \circ \phi_1^{-1}$ {see Equation 4}
- 6: Compute the face flip penalty by the signed area: $L_{\text{flip}} \leftarrow \sum_{\text{faces } \Delta_{ijk}} \max(0, -\text{signed area}(\Delta_{ijk}))$
- 7: Augment energy: $E_D \leftarrow E_D + \lambda L_{\text{flip}}$ {large penalty weight λ }
- 8: Compute energy gradient $\nabla E_D(v)$ for each vertex $v \in \hat{f}_i(F_1)$
- 9: $\widetilde{\nabla E_D}(v) \leftarrow \nabla E_D(v)$ for all v
- 10: **for** each vertex set $S \in \mathcal{O}$ **do** {Average gradient vectors on all attaching points}
- 11: $\widetilde{\nabla E_D}(v) \leftarrow \frac{1}{|S|} \sum_{v \in S} \nabla E_D(v)$ for all $v \in S$
- 12: **end for**
- 13: Update vertex positions as $\hat{f}_{i+1}(F_1): v \leftarrow v - \eta \widetilde{\nabla E_D}(v)$
- 14: **end for**

can also prescribe a homotopy class manually: Section 4.3 reports an experiment using a different homotopy class. Other cut systems, including greedy homotopy or homology generators [Erickson and Whittlesey 2005], are admissible when they yield a fundamental polygon with the required side identifications. However, the greedy basis alone does not guarantee that its loops are disjoint away from a single common base point, so further processing may be needed to obtain the standard fundamental polygon used by our parameterization.

Next, we cut the mesh along those edges into a topological disk while storing the attaching map. We evenly distribute the vertices of the longitude loop along the bottom and top edges of the unit square, and those of the meridian loop along the left and right edges. The four copies of the unique intersection point are pinned at the four corners of the unit square.

To compute planar vertex coordinates, we use these boundary points as boundary conditions for Tutte's embedding of the cut disk with uniform edge weights [Tutte 1963]. An example is shown in Figure 2. This embedding induces a piecewise linear parameterization $\phi : F \rightarrow T$, such that any point on F can be mapped by locating the face containing it and using barycentric coordinates to map it to the corresponding face on T . The algorithm for finding barycentric coordinates is described in Algorithm 2 of [Koehl and Hass 2013].

Using the computed linear parameterization $\phi : F \rightarrow T$, we can map any point in the fundamental domain F to the torus T . If a point lies outside the unit square, ϕ first reduces its coordinates modulo the unit square and then maps the resulting point to T .

3.2 Computing the Energy

As discussed in Equation 4, the energy of the map $f : T_1 \rightarrow T_2$ depends only on the domain edge lengths L_i on T_1 and the image edge lengths l_i on T_2 . For an edge $e = (u, v)$ on T_1 , its image is the piecewise linear path between the two points $f(u)$ and $f(v)$ on T_2 . In our formulation, edge lengths on T_2 are given by the geodesic distance computed with [Sharp and Crane 2020].

3.3 Minimizing the Energy: Gradient Descent

We minimize $E_D(f)$ through $f = \phi_2 \circ \hat{f} \circ \phi_1^{-1}$, where we store and update the 2D vertex positions of $\hat{f}(F_1)$ on the fundamental domain F_2 and then map them back to T_2 via ϕ_2 . The gradients ∇E_D are taken with respect to these vertex coordinates, so each vertex receives a 2D gradient vector through the chain rule. This process requires addressing the following issues.

Boundary consistency. The first concern is ensuring that $\hat{f}(F_1)$ remains a valid fundamental domain, which requires that the boundary edges are equivariant and can be properly identified with each other via integer translation. To achieve this, for each boundary point in $\hat{f}(F_1)$, we take the average of the gradient vectors at all corresponding points. Specifically, except for the corner vertex that has 4 copies, all other boundary vertices have two gradients directly from ∇E_D . We assign the average gradient vectors to the shared vertices, ensuring that identified boundary vertices move in the same direction with the same step size. Moreover, by maintaining the boundary constraints and moving boundary vertices consistently, our gradient descent optimizes the energy within the fixed initial homotopy class.

Injectivity. We need to guarantee that the map $f_i = \phi_2 \circ \hat{f}_i \circ \phi_1^{-1}$ is injective. Non-injectivity can arise when faces in $\hat{f}(F_1)$ cross each other, causing face flips and preventing f from being a homeomorphism. We add a large penalty term to forbid such face flips. Specifically, for an oriented face Δ_{ijk} in $\hat{f}(F_1)$, we compute its signed area $\text{SArea}(\Delta_{ijk}) = \frac{1}{2}(\mathbf{x}_j - \mathbf{x}_i) \times (\mathbf{x}_k - \mathbf{x}_i)$. A negative value indicates an orientation reversal. To handle numerical precision, we apply the penalty only when the signed area falls below a threshold of $\epsilon = 10^{-10}$, with a small regularization term to prevent gradient vanishing. Thus, each flipped face contributes a penalty to the energy by $E_D \leftarrow E_D + 10^8 \cdot (|\text{SArea}(\Delta_{ijk})| + 10^{-10})^2$. In our experiments, this penalty is triggered only rarely and mainly acts as a safeguard during optimization.

Maximum and subgradients. The maximum function in Equation 4 is non-smooth. PyTorch handles this issue by using a subgradient, where only the face with maximum angle-distortion contributes to the gradient. We also tested LogSumExp as a smooth approximation. However, our experiments show that LogSumExp accelerates energy convergence only on flat tori and does not consistently improve computations on 3D tori.

Stopping criteria. The gradient descent stops when the energy converges. Our method attains stable convergence using plateau-based stop criteria and adaptive step-size control. In practice, we stop when the relative improvement in the energy stays below 10^{-5} for 5000 iterations. If the energy fails to decrease by at least 10^{-3} in relative terms for 3000 iterations, we halve the step size.

Nonconvexity. We have not established that the objective is convex, and it is likely nonconvex, so the optimization may converge to a local minimum. In practice, we can compare the optimized directional values $E_D(f : T_1 \rightarrow T_2)$ and $E_D(g : T_2 \rightarrow T_1)$. Since the smooth global minimum is symmetric, similar minimized values of $E_D(f : T_1 \rightarrow T_2)$ and $E_D(g : T_2 \rightarrow T_1)$ provide evidence that we have reached a global minimum, as shown in Section 4.5.

Existence of a Minimizer. For the discrete energy $E_D(f)$, a minimizer exists. Once we fix the triangulation T_1 , $E_D(f)$ is a continuous function of the finitely many vertex positions in $V(T_1)$. Since the domain and codomain are compact tori, we can conclude that a minimizer of $E_D(f)$ exists. However, it has not been proved that the minimizer is a homeomorphism. In the smooth setting, the space of diffeomorphisms is infinite-dimensional, and the existence question is more complicated, which can be found in Theorem 4.4 of [Luo 2021].

Finally, the distance results reported in the experiments are computed using the exact energy in Equation 4, without the face-flip penalty or any smooth approximation of the maximum function.

4 Experiments

All experiments were run on a personal Linux machine with a 12th Gen Intel Core i7-12700K CPU and 31 GB RAM. We use GeometryCentral [Sharp et al. 2019] for geometric processing, including construction of linear parameterizations, and PyTorch [Paszke et al. 2019] for automatic differentiation during optimization.

4.1 Datasets

We evaluate our framework on three types of datasets:

- **Flat tori:** Triangulated parallelograms with opposite sides identified to verify the convergence to Teichmüller maps.
- **Tori of revolution:** Synthetic tori generated by the standard parametrization $X(u, v) = ((R+r \cos v) \cos u, (R+r \cos v) \sin u, r \sin v)$, with the major radius R and the minor radius r . This dataset lets us test the distance on curved surfaces with known geometry. We also minimize the Dirichlet energy within the same framework, showing that it does not produce a metric on the shape space.
- **Thing10K:** We use eight models from [Zhou and Jacobson 2016] to demonstrate the usefulness of our framework for practical shape comparison tasks.

We first discuss the symmetrized discrete distance defined in Equation 5 on these datasets and then examine the symmetry and triangle inequality of the unsymmetrized minimized energy.

4.2 Flat Tori

A flat torus is a parallelogram in $\mathbb{R}^2$ with opposite sides identified via translations. With one spanning vector fixed at $(1, 0)$, every flat torus can be represented by a point (a, b) in the upper half-plane $\mathbb{H}^2 = \{z \in \mathbb{C} \mid \text{Im}(z) > 0\}$. In the smooth setting, by classical Teichmüller theory, the minimizer of our energy on two equal-area flat tori is the Teichmüller map, the linear transformation A between the two parallelograms. The distance d_E equals half the logarithm of the ratio of the two singular values σ_1 and σ_2 of A .

For example, in Figure 3, we compare two flat tori T_1 (Figure 3a) and T_2 (Figure 3b) with parameterizations $\phi_1(x, y) = (x, y)$ and $\phi_2(x, y) = (x, 0.5x + y)$, respectively. We initialize with $\hat{f}_0(x, y) = (1 - (1 - x)^2, 1 - (1 - y)^2)$ on the unit square. After 1000 iterations, the image $\hat{f}_0(T_1) = \phi_2 \circ \hat{f}_0 \circ \phi_1^{-1}(T_1)$ (Figure 3c) converges to the linear map, giving $\hat{f}_{1000}$ (Figure 3d). Figure 5 visualizes the evolution of the vertices $\hat{f}_i(F_1)$ over iterations. Figure 4 shows that the area-distortion term converges to 0, while the angle-distortion term converges to the Teichmüller distance $\frac{1}{2} \log(\sigma_1/\sigma_2) \approx 0.24$. The orange curves are obtained by replacing the maximum in the shape-stretching term ($\|\log \frac{\sigma_1}{\sigma_2}\|_\infty$) with its smooth approximation, LogSumExp, and show a faster convergence rate. However, in later experiments on 3D surfaces, we find that LogSumExp does not improve the convergence rate of the discrete energy.




Fig. 3. Comparing two flat tori T_1 and T_2 by minimizing the energy $E_D(f)$. After 1000 iterations, f converges to the Teichmüller map.




Fig. 4. The two terms in the energy $E_D(f_i)$ over iterations. The area-distortion term converges to zero, and the angle-distortion term converges to the Teichmüller distance (red dashed line).

When two flat tori represented by $(a, b), (c, d) \in \mathbb{H}^2$ have different areas, we can find the theoretical value of the energy of the Teichmüller map $E(A)$. The two terms are $\sqrt{\int_{T_1} (\sqrt{\lambda_1 \lambda_2} - 1)^2 dA} = \sqrt{(\sqrt{\frac{d}{b}} - 1)^2 b} = |\sqrt{b} - \sqrt{d}|$ and $\frac{1}{2} \log \frac{\sigma_1}{\sigma_2}$, where σ_1 and σ_2 are the singular values of the linear transformation between two flat tori. We test this case using $\phi_1(x, y) = (x, 3x + 4y)$ and $\phi_2(x, y) = (x, 5x + 9y)$, with the same initialization $\hat{f}_0$ as above. In Figure 6, both terms




Fig. 5. The images of the vertices $\hat{f}_i(F_1)$ at different iterations. The map converges to the identity in the end, and therefore $\phi_2 \circ \hat{f} \circ \phi_1^{-1}$ converges to the Teichmüller map.

of $E(f_i)$ converge to the theoretical values $|\sqrt{4} - \sqrt{9}| = 1$ and $\frac{1}{2} \log \frac{\sigma_1}{\sigma_2} \approx 0.4349$, shown as red dashed lines.




Fig. 6. The two terms in the energy $E_D(f_i)$ for the unequal-area experiment over 2000 iterations, shown on logarithmic vertical axes. *Left*: area distortion. *Right*: angle distortion. The red dashed lines indicate the theoretical values from Teichmüller theory.

We further evaluate the metric on five flat tori represented by $(0, 1)$, $(1, 10)$, $(1, 0.1)$, $(-1, 10)$, and $(-1, 0.1)$ to assess numerical stability and robustness to triangulation changes. For each torus, we use two triangulations: a coarse mesh with 72 faces and a subdivided mesh with 432 faces, giving ten meshes in total. Their pairwise distances are shown in Table 1. The index system is organized as pairs: indices $(2k, 2k + 1)$ correspond to the same underlying torus, where $2k$ is the 72-face mesh and $2k + 1$ is the 432-face mesh. The entries in each odd row $2k + 1$ are close to those in the even row above $2k$, indicating that the distance is robust to triangulation changes. The largest absolute deviation from the theoretical value occurs at $(3, 7)$, where $|d_E - d_{\text{theory}}| = 1.58 \times 10^{-3}$ and $\frac{|d_E - d_{\text{theory}}|}{d_{\text{theory}}} \approx 1.58\%$. This is on the same scale as the largest diagonal entry $d_E(9, 9) = 1.50 \times 10^{-3}$ that can be treated as numerical tolerance.

Table 1. Distance matrix d_E for flat tori. **Bold** marks the entry with the largest deviation from the theoretical value.

$T_1 \backslash T_2$	0	1	2	3	4	5	6	7	8	9
0	0.000020	0.000036	3.318595	3.318595	2.182890	2.182890	3.318595	3.318595	2.182890	2.182890
1		0.000052	3.318595	3.318595	2.182890	2.182890	3.318595	3.318595	2.182890	2.182890
2			0.000101	0.000237	5.148635	5.148635	0.099924	0.100580	5.168247	5.168247
3				0.000374	5.148635	5.148635	0.100759	0.101415	5.168247	5.168247
4					0.000150	0.000500	5.168247	5.168247	2.998316	2.998323
5						0.000850	5.168247	5.168247	2.998304	2.998311
6							0.000621	0.000496	5.148635	5.148635
7								0.000371	5.148635	5.148635
8									0.000316	0.000908
9										0.001500

4.3 Tori of revolution

For a standard torus of revolution with major radius R and minor radius r , the parametrization is

$$X(u, v) = ((R + r \cos v) \cos u, (R + r \cos v) \sin u, r \sin v). \quad (6)$$

We build a dataset of five shapes by fixing $R = 1$ and varying $r \in \{0.1, 0.2, 0.5, 0.8, 0.9\}$. For each shape, we use two sampling resolutions on the same smooth surface: r_i denotes a mesh with 20×10 vertices (20 along longitude and 10 along meridian), and r_i^d denotes a denser mesh with 20×30 vertices. Their distances are listed in Table 2. The heatmap for r_i^d in Figure 1 follows the expected pattern: nearby minor radii (e.g., 0.1 vs. 0.2 and 0.8 vs. 0.9) have smaller distances and appear darker.

We observe that these tori approximately lie on a geodesic in shape space under the induced path metric. In Figure 1, for example, the three pairwise distances among $r = 0.1, 0.2, 0.5$ numerically satisfy triangle equality: $0.43 + 0.763 \approx 1.196$. This agrees with the geometric intuition that the sequence follows a near-optimal deformation from $r = 0.1$ to $r = 0.9$.

We next use seven variants of the torus $(R, r) = (1, 0.5)$ in Figure 7 to compare discretization effects with genuine shape changes. Variants (a)–(e) are different triangulations or samplings of the same smooth torus, while (f) and (g) introduce substantial radial perturbations. Figure 8 shows that our distance keeps same-shape variants close, including (a) and its subdivision (b), as well as the denser samplings (d) and (e). This indicates that denser discretizations better approximate the same underlying smooth surface and remain separated from the geometrically perturbed variants.




Fig. 7. Variants of tori of revolution with $R = 1$ and $r = 0.5$ in front, slanted, and side views: (a) 20×10 samples; (b) barycentric subdivision of (a); (c) 10×10 samples; (d) 30×30 samples; (e) 40×40 samples; (f) minor radius varied along longitude; (g) minor radius varied along meridian.





Fig. 8. Distance matrix of the seven tori, ordered as in Figure 7. Each entry (i, j) is the distance between the i -th and j -th meshes.

Entropic Gromov–Wasserstein baseline. For comparison, we computed an entropic GW (EGW) [Solomon et al. 2016] distance matrix for the seven variants in Figure 7. We sampled 100 points by farthest-point sampling, used normalized area weights and exact surface-geodesic distance matrices, and ran the entropy-regularized optimal transport solver with $\epsilon = 0.005$. Figure 9 shows that EGW can distinguish the perturbed torus (g) from the other variants, with distances around 0.052, but these values are still on the same scale as the remaining off-diagonal entries. Compared to our distance in Figure 8, EGW is less informative about discretization effects, as the distance between (a) and its barycentric subdivision (b) is 4.177×10^{-2} , which is close to the other entries.

Dirichlet energy. As an independent check of the optimization framework, we minimized the Dirichlet energy on the five dense tori of revolution and compared the optimized values in Table 3 with the theoretical Dirichlet energies of the rotationally equivariant harmonic maps reported in Table 4. The largest relative discrepancy from the theoretical value is about 0.6%. This validates the numerical optimization pipeline. Moreover, it demonstrates that the Dirichlet energy does not produce a metric on the shape space. Indeed, the directed Dirichlet values are asymmetric (9.09 versus 17.57 for $r_1^d \leftrightarrow r_9^d$) and violate the triangle inequality. For example,

Table 2. Distance matrix d_E for tori of revolution.

$T_1 \setminus T_2$	r_1^d	r_2	r_2^d	r_5	r_5^d	r_8	r_8^d	r_9	r_9^d
r_1	0.028442	0.428164	0.442515	1.179618	1.205273	1.872820	1.930528	2.230565	2.283417
r_1^d		0.422830	0.430210	1.177199	1.195657	1.897293	1.928980	2.274788	2.295306
r_2			0.024874	0.752508	0.777550	1.467255	1.508461	1.836771	1.873121
r_2^d				0.751724	0.763383	1.472170	1.497210	1.850317	1.870128
r_5					0.030035	0.714675	0.760718	1.076982	1.144695
r_5^d						0.707069	0.728298	1.083630	1.101268
r_8							0.043846	0.361732	0.419839
r_8^d								0.340981	0.371688
r_9									0.076267

Fig. 9. Entropic Gromov-Wasserstein distance matrix for the seven torus variants in Figure 7. In-cell annotations are shown in units of 10^{-2} .

a triangle-inequality violation is

$$E_{\text{Dirichlet}}(r_8^d \rightarrow r_1^d) = 11.66$$

$$> E_{\text{Dirichlet}}(r_8^d \rightarrow r_5^d) + E_{\text{Dirichlet}}(r_5^d \rightarrow r_1^d) = 7.77.$$

Table 3. Optimized Dirichlet energy between five dense tori of revolution. Each entry is the minimized Dirichlet energy for the map from the source torus (row) to the target torus (column).

$T_1 \setminus T_2$	r_1^d	r_2^d	r_5^d	r_8^d	r_9^d
r_1^d	—	2.484162	5.190113	8.122672	9.093769
r_2^d	2.499404	—	2.891102	4.245006	4.707621
r_5^d	5.475123	2.947741	—	2.214375	2.347642
r_8^d	11.661931	5.374198	2.293539	—	1.991618
r_9^d	17.573420	7.741899	2.649806	2.000032	—

Choices of homotopy classes. Apart from the canonical homotopy class, we also tested another homotopy class that maps the meridian of one torus to the longitude of the other and vice versa, using the $r = 0.1$ and $r = 0.5$ tori in Figure 1. The optimized energy was 1.832, compared with 1.196 for the class used in the main experiment. It shows that different homotopy classes define different admissible map families and can produce different minimized values for the same pair of shapes.

Table 4. Theoretical Dirichlet energy of the degree-one rotationally equivariant harmonic map between five dense tori of revolution.

$T_1 \setminus T_2$	r_1^d	r_2^d	r_5^d	r_8^d	r_9^d
r_1^d	—	2.5000	5.2000	8.1251	9.1112
r_2^d	2.5006	—	2.9003	4.2505	4.7229
r_5^d	5.4682	2.9494	—	2.2292	2.3617
r_8^d	11.5916	5.3663	2.3039	—	2.0173
r_9^d	17.5160	7.7068	2.6690	2.0235	—

4.4 Thingi10K Dataset

Pairwise distances are often used for clustering. We use eight genus-one meshes from Thingi10K [Zhou and Jacobson 2016], grouped into four shape categories: two rings, two rollers, two spheres, and two circular nuts. Each mesh is remeshed to 600 vertices for uniform resolution. Figure 10 shows the models (front, slanted, and side views), and Table 5 reports the distance matrix.




Fig. 10. Eight Thingi10K-derived meshes used in this experiment.

For each mesh pair, we compute (i) our discrete distance and (ii) Euclidean Hausdorff distance between vertex sets after scaling each mesh to unit area. We run average-linkage hierarchical clustering on the two 8×8 matrices, summarized in Figure 11 and Figure 12.

Table 5. Distance matrix d_E between the eight Thingi10K-derived meshes.

	ring ₁	roll ₀	roll ₁	sph ₀	sph ₁	nut ₀	nut ₁
ring ₀	0.505489	1.008749	1.143026	0.693118	0.718523	0.651501	0.612866
ring ₁		1.086253	1.062546	0.844025	0.736091	0.522911	0.498400
roll ₀			0.515594	0.831689	0.837155	0.917694	0.814220
roll ₁				0.937884	0.944770	0.751489	0.751060
sph ₀					0.568940	0.842467	0.773149
sph ₁						0.740319	0.709386
nut ₀							0.227391




Fig. 11. Clustering eight meshes with our distance d_E based on energy minimization. Left: Heatmap of d_E ; Right: Average-linkage dendrogram for d_E .

The Hausdorff distance is symmetric and satisfies the triangle inequality, but it is often a weak choice for shape comparison. It depends on the embedding in $\mathbb{R}^3$ rather than intrinsic geometry, is sensitive to outliers and noise, and does not provide meaningful correspondences. As shown in Figures 11 and 12, our distance d_E separates the four intended pairs at the leaves of the dendrogram. In contrast, Hausdorff distance in $\mathbb{R}^3$ is dominated by the large geometric gap between the rings and the other six meshes. The remaining structure is relatively flat, so the dendrogram does not recover the four groups. Therefore, our distance outperforms the Hausdorff distance.




Fig. 12. Clustering eight meshes with the Hausdorff distance d_H . Left: Heatmap of d_H ; Right: Average-linkage dendrogram for d_H .

4.5 Symmetry of the Optimized Discrete Distance

The smooth energy is symmetric, that is, $E(f) = E(f^{-1})$, but this equality is not guaranteed for the discrete energy E_D computed on a pair of discrete meshes. Mesh resolution and regularity, geodesic-distance approximation, and optimization error can all contribute to directional discrepancies. The following experiments on flat tori and tori of revolution examine how directional discrepancies in the minimized unsymmetrized energies behave under refinement. We use $E_{i,j}$ to denote the directional minimized energy $E_D(f)$ for the map $f: T_i \rightarrow T_j$.

4.5.1 Flat Tori. Using the same ten-mesh flat-torus dataset, we report directional minimized energies among the five 72-face meshes and among the five 432-face meshes in Tables 7 and 8, respectively. We also list the theoretical values of the energy of the Teichmüller map between the two parallelograms in Table 6 for reference. In

each table, **bold** marks entries that realize the maximum asymmetry gap $|E_{ij} - E_{ji}|$, and *italic* marks entries with the maximum deviation from the theoretical values.

Table 6. Theoretical values of the energy of the linear transformation (Teichmüller map) between the five flat tori with parameters $(0, 1)$, $(1, 10)$, $(1, 0.1)$, $(-1, 10)$, and $(-1, 0.1)$.

	$(0, 1)$	$(1, 10)$	$(1, 0.1)$	$(-1, 10)$	$(-1, 0.1)$
$(0, 1)$	0.000000	3.318595	2.182890	3.318595	2.182890
$(1, 10)$	3.318595	0.000000	5.148635	0.099834	5.168247
$(1, 0.1)$	2.182890	5.148635	0.000000	5.168247	2.998223
$(-1, 10)$	3.318595	0.099834	5.168247	0.000000	5.148635
$(-1, 0.1)$	2.182890	5.168247	2.998223	5.148635	0.000000

Table 7. Directional minimized energies among the 72-face flat tori.

	0	2	4	6	8
0	0.000020	3.318595	2.182890	3.318595	2.182890
2	3.318595	0.000101	5.148635	0.099990	5.168247
4	2.182890	5.148635	0.000150	5.168247	2.998318
6	3.318595	0.099857	5.168247	<i>0.000621</i>	5.148635
8	2.182890	5.168247	2.998314	5.148635	0.000316

Table 8. Directional minimized energies among the 432-face flat tori.

	1	3	5	7	9
1	0.000052	3.318595	2.182890	3.318595	2.182890
3	3.318595	0.000374	5.148635	0.101660	5.168247
5	2.182890	5.148636	0.000850	5.168247	2.998295
7	3.318595	0.101170	5.168247	0.000371	5.148635
9	2.182890	5.168247	2.998328	5.148635	0.001500

The highlighted entries show where the largest numerical differences occur. In Table 7, the maximum asymmetry is $|E_{2,6} - E_{6,2}| = 1.33 \times 10^{-4}$, and the maximum deviation from the theoretical values is 6.21×10^{-4} at $(6, 6)$. In Table 8, these maxima are $|E_{3,7} - E_{7,3}| = 4.90 \times 10^{-4}$ and 1.83×10^{-3} at $(3, 7)$. In both resolutions, the maximum asymmetry is smaller than the largest diagonal self-distance, suggesting that the directional discrepancies are comparable to numerical error.

4.5.2 Tori of Revolution. Using the same five-shape dataset with $R = 1$ and $r \in \{0.1, 0.2, 0.5, 0.8, 0.9\}$, we report directional minimized energies among the five sparse meshes (20×10 vertices) and among the five dense meshes (20×30 vertices) in Tables 9 and 10, respectively.

From Table 9, the largest asymmetry is $|E_{r_5, r_9} - E_{r_9, r_5}| = |0.990132 - 1.163831| = 0.173699$. From Table 10, the largest asymmetry is $|E_{r_8, r_9}^d - E_{r_9, r_8}^d| = |0.361814 - 0.381562| = 0.019748$. The smaller value in the dense case shows that denser sampling of the same smooth surfaces reduces asymmetry. The largest discrepancies occur near $r = 0.9$: when $R = 1$ and $r = 0.9$, triangles on the inner side of the torus become very skinny (near-degenerate). This can violate the angle-regularity assumption in our convergence theorem

Table 9. Directional minimized-energy matrix for sparse samplings of rotational tori (r_i : 20×10 vertices).

$T_1 \setminus T_2$	r_1	r_2	r_5	r_8	r_9
r_1	–	0.427665	1.171629	1.818593	2.189619
r_2	0.428663	–	0.740754	1.428485	1.792835
r_5	1.187607	0.764262	–	0.675854	0.990132
r_8	1.927047	1.506025	0.753496	–	0.304562
r_9	2.271512	1.880707	1.163831	0.418902	–

Table 10. Directional minimized-energy matrix for dense samplings of rotational tori (r_i^d : 20×30 vertices).

$T_1 \setminus T_2$	r_1^d	r_2^d	r_5^d	r_8^d	r_9^d
r_1^d	–	0.430051	1.196057	1.928269	2.285559
r_2^d	0.430368	–	0.760756	1.494182	1.862984
r_5^d	1.195257	0.766010	–	0.719315	1.092578
r_8^d	1.929690	1.500239	0.737282	–	0.361814
r_9^d	2.305053	1.877272	1.109959	0.381562	–

(Theorem 2.1) and make the optimization less stable. In addition, distances on a flat torus are exactly Euclidean, whereas for a torus of revolution, we use geodesic distances approximated on a polyhedral mesh, which can further lead to asymmetry.

4.6 Triangle Inequality of the Optimized Discrete Energy

The triangle inequality is another basic criterion for determining whether a matrix of pairwise scores behaves like a metric. For a discrete distance matrix M , we quantify violations as follows. For pairwise distinct indices p, q, r , we define the triangle discrepancy as $M_{pr} - M_{pq} - M_{qr}$. We define the *maximum triangle violation* by

$$\mathcal{V}(M) = \max_{p,q,r \text{ distinct}} \max\{0, M_{pr} - M_{pq} - M_{qr}\}.$$

Then $\mathcal{V}(M) = 0$ means that M satisfies the triangle inequality for every distinct triple. When $\mathcal{V}(M)$ is positive but small, we compare it with the scale of the matrix and with the numerical self-distances.

For flat tori in Table 1, we obtain $\mathcal{V} \approx 5.98 \times 10^{-4}$, which is below the numerical tolerance suggested by self-distances. For tori of revolution in Table 2, the five minor radii appear to follow an almost one-dimensional path in shape space, so some triples have nearly additive distances and small optimization errors can show up as triangle violations. For the sparse subset $\{r_i\}$ and dense subset $\{r_i^d\}$ of Table 2, we obtain $\mathcal{V} \approx 7.78 \times 10^{-3}$ and $\mathcal{V} \approx 5.53 \times 10^{-3}$, respectively; both are smaller than the smallest nonzero entries in the corresponding subsets. For the Thingi10K matrix in Table 5, $\mathcal{V} = 0$.

4.7 Running Times

Runtime is dominated by geodesic computations. Experiments with flat tori finish in minutes per pair, while high-resolution 3D pairs require about 15 hours. This gap mainly comes from implementation-level geodesic cost. Replacing edge-wise geodesic calls with distance-field approximations (e.g., heat-method variants [Crane et al. 2017]) and restructuring exact geodesic evaluation into reusable single-source propagation should provide substantial speedups.

5 Proof of the convergence theorem

In this section, we prove that the discretized energy $E_D(f)$ of a diffeomorphism f converges to the smooth energy $E(f)$ as the surfaces are approximated by sufficiently dense triangulations. We also recall the quantities in the discrete energy defined in Section 2.2.

Recall the following definitions in Section 2. Given two Riemannian metrics g and h on S , we define an energy functional that measures the distortion induced by a diffeomorphism $f : (S, g) \rightarrow (S, h)$. Its energy is

$$E(f) = \sqrt{\int_S (\sqrt{\lambda_1 \lambda_2} - 1)^2 dA_g} + \frac{1}{2} \left\| \log \frac{\lambda_1}{\lambda_2} \right\|_\infty. \quad (7)$$

Here, $\lambda_1(p)$ and $\lambda_2(p)$ are the stretch factors of f at $p \in S$: they are the singular values of df_p , computed with respect to orthonormal bases in the source and target tangent spaces equipped with g and h , respectively. Set

$$R(p) = \frac{\lambda_1(p)}{\lambda_2(p)}, \quad D(p) = \lambda_1^2(p) + \lambda_2^2(p), \quad P(p) = \lambda_1(p)\lambda_2(p). \quad (8)$$

Then $(R + R^{-1})P = D$, $D(p) = \|df_p\|_F^2$, where $\|\cdot\|_F$ denotes the Frobenius norm, and all these functions are smooth on S . The energy $E(f)$ can be expressed as

$$E(f) = \sqrt{\int_S (\sqrt{P} - 1)^2 dA_g} + \frac{1}{2} \left\| \log R \right\|_\infty. \quad (9)$$

In practice, a surface is described by a triangulation T and a piecewise Euclidean metric determined by an edge-length function l ; we denote it by (S, T, l) . A triangulation of a surface is a decomposition of the surface into triangular faces Δ_{ijk} , where each face Δ_{ijk} is a triangle with vertices v_i, v_j , and v_k .

For a homeomorphism $f : (S, T_1, l_1) \rightarrow (S, T_2, l_2)$ between triangulated tori, the discrete energy is

$$E_D(f) = \sqrt{\sum_{\Delta_{ijk} \in T_1} (1 - \sqrt{\lambda_1(\Delta_{ijk})\lambda_2(\Delta_{ijk})})^2 A(\Delta_{ijk})} \quad (10)$$

$$+ \frac{1}{2} \max_{\Delta_{ijk} \in T_1} \left| \log \frac{\lambda_1(\Delta_{ijk})}{\lambda_2(\Delta_{ijk})} \right|. \quad (11)$$

where Δ_{ijk} is a face of T_1 , $A(\Delta_{ijk})$ is its Euclidean area, and $\lambda_1(\Delta_{ijk})$ and $\lambda_2(\Delta_{ijk})$ are the two singular values of the differential df at Δ_{ijk} determined below.

In the remainder of the proof, we show that the first and second terms of Equation (4) converge to the corresponding terms of Equation (9).

5.1 Setup of the theorem

To compare $E(f)$ and $E_D(f)$, we choose a sufficiently dense geodesic triangulation on (S, g) and decompose the integral in $E(f)$ over its geodesic triangles. For any $\delta > 0$ and any Riemannian surface (S, g) , we can choose a geodesic triangulation T whose maximal edge length is smaller than δ and whose minimum angle is at least $\pi/5$. The existence of such dense and regular geodesic triangulations has been established in [Colin de Verdiere and Alexis 1990], and the choice of $\pi/5$ here is for convenience and can be relaxed. Without

loss of generality, we normalize the area of (S, g) such that $A(S, g) = 1$.

We fix the following notation for the rest of the paper. For a triangulation T , let $V(T)$, $E(T)$, and $F(T)$ denote its vertices, edges, and faces. We write vertices as v_i , edges as e_{ij} , and faces as triangles Δ_{ijk} . Let $l : E(T) \rightarrow \mathbb{R}$ be the edge length function of a geodesic triangulation T . Let $A_g(\Delta_{ijk})$ denote the area of a geodesic triangle with respect to the metric g . The edge length l and the triangulation T determine a unique piecewise Euclidean surface (T_E, l_E) . Here T_E contains Euclidean triangles Δ_{ijk}^E with the same edge lengths as the corresponding geodesic triangles Δ_{ijk} in (T, l) . Notice that $l = l_E$ and $T_E = T$ as combinatorial triangulations. However, Δ_{ijk}^E and Δ_{ijk} have different geometries. Let $A_E(\Delta_{ijk}^E) = A_E(\Delta_{ijk})$ denote the area of a Euclidean triangle in (T_E, l_E) . By Lemma 5.3 in [Wu and Zhu 2024], we can assume that the angles of Δ_{ijk}^E have a lower bound $\pi/6$ when δ is sufficiently small. Again, the specific choice of $\pi/6$ is just for convenience.

Similarly, on the target surface (S, h) , let $f(\Delta_{ijk})$ be the image of the geodesic triangle Δ_{ijk} . This image is not necessarily a geodesic triangle. Let $A_h(f(\Delta_{ijk}))$ be its area on (S, h) . When Δ_{ijk} is sufficiently small, the vertices $f(v_i), f(v_j), f(v_k)$ determine a unique geodesic triangle Δ_{ijk}^h on (S, h) . Let $l^h(e_{ij})$ be the length of the geodesic connecting $f(v_i)$ and $f(v_j)$ that is homotopic to the image of $f(e_{ij})$. Let $A_h(\Delta_{ijk}^h) = A_h(\Delta_{ijk})$ be the area of this geodesic triangle on (S, h) . Let $A_E(\Delta_{ijk}^h)$ be the area of the corresponding Euclidean triangle with the same edge lengths.

The proof uses the following fact several times. Because (S, g) and (S, h) are closed surfaces with smooth Riemannian metrics and $f : (S, g) \rightarrow (S, h)$ is a diffeomorphism, the metric tensors g_{ij} and h_{ij} , their derivatives (including their Christoffel symbols), and the first and second derivatives of f are smooth and uniformly bounded by a constant $C = C(S, g, h, f)$. The value of C may differ between estimates below, but it depends only on S, g, h , and f .

We also use the fact that the map f is Lipschitz with some constant C , namely $d_h(f(x), f(y)) \leq C d_g(x, y)$ for any $x, y \in (S, g)$, where d_h and d_g are the distance functions on (S, h) and (S, g) respectively. For any $\epsilon > 0$, we can choose δ sufficiently small so that every geodesic triangle Δ_{ijk} lies in a strongly convex normal neighborhood of (S, g) with radius of order $O(\epsilon)$. Since f is Lipschitz, the diameter of $f(\Delta_{ijk})$ is also of order $O(\epsilon)$. Hence, we can assume that every triangle $f(\Delta_{ijk})$ also lies in a strongly convex normal neighborhood of (S, h) with radius of order $O(\epsilon)$.

5.2 Auxiliary lemmas about lengths and areas

Recall that for a point p in (S, g) , the normal coordinates centered at p are given by the chart $y \mapsto \exp_p(y)$ from a neighborhood of 0 in $T_p S$ to a neighborhood of p in S . Points near p are identified with vectors $y \in T_p S$, where the center p corresponds to $y = 0$. In these coordinates, the metric tensor g has components $g_{ij}(y)$. At the center, $g_{ij}(0) = \delta_{ij}$ and $\partial_k g_{ij}(0) = 0$, so geodesics through p are straight lines through the origin in the y -coordinates. Moreover, by the Taylor expansion in y , $g_{ij}(y) = \delta_{ij} + O(|y|^2) = \delta_{ij} + O(\epsilon^2)$. It leads to a classical fact in Riemannian geometry that the distance in the normal coordinate ball is preserved up to $O(\epsilon^3)$ after mapping

to the surface. Specifically, let $B_y(r)$ be a ball of radius r centered at y in $T_p S$. Then

Lemma 5.1. Let $p = \exp_x(y)$ and $q = \exp_x(z)$ with $y, z \in B_0(\epsilon) \subset T_x S$, i.e., p, q are in the normal coordinate ball $U_\epsilon = \exp_x(B_0(\epsilon))$. Then $d_g(p, q) = |y - z| + O(\epsilon^3)$.

Let q_0, q_1 be two points on (S, g) in a normal ball of radius $O(\epsilon)$ centered at $p \in S$. Let $\gamma(t)$ be a smooth curve parametrized proportionally with $\gamma(0) = q_0$ and $\gamma(1) = q_1$. Let $y = \exp_p^{-1}(q_0)$, $z = \exp_p^{-1}(q_1)$, and $\gamma(t) = \exp_p(\sigma(t))$, where $\sigma(t)$ is a curve in $T_p S$. We show that $\sigma(t)$ lies on an $O(\epsilon)$ neighborhood of the line segment between y and z in $T_p S$ if the second derivatives of $\gamma(t)$ are bounded.

Lemma 5.2. Let $\eta(t) = (1 - t)y + tz$ be the line segment and $\gamma(t), \sigma(t)$ defined above. If $|\gamma''(t)| = O(\epsilon)$ for $t \in [0, 1]$, then $|\sigma(t) - \eta(t)| = O(\epsilon)$.

PROOF. Define $\phi(t) = \sigma(t) - \eta(t)$ on $[0, 1]$ and $|\phi''(t)| = |\sigma''(t)|$. By the property of metric tensors in normal neighborhoods, we have $|\sigma''(t)| = O(\epsilon)$ from $|\gamma''(t)| = O(\epsilon)$. Then we have a function ϕ on $[0, 1]$ such that $\phi(0) = \phi(1) = 0$ and $|\phi''(t)| = O(\epsilon)$. Rolle's Theorem implies that there exists a $\xi \in (0, 1)$ such that $\phi'(\xi) = 0$. Then $|\phi'(t)| = O(\epsilon)$ and $|\phi(t)| = O(\epsilon)$ on $[0, 1]$ as well. $\square$

Now we are ready to establish area estimates of various triangles.

Lemma 5.3. If Δ_{ijk} lies in a normal coordinate ball centered at $p \in (S, g)$ of radius $O(\epsilon)$, then

$$|A_g(\Delta_{ijk}) - A_E(\Delta_{ijk})| = O(\epsilon^4),$$

$$|A_h(f(\Delta_{ijk})) - A_E(\Delta_{ijk}^h)| = O(\epsilon^3).$$

PROOF. We will write $\Delta = \Delta_{ijk}$ in this proof. To establish the first inequality, let $\tilde{\Delta} = \exp^{-1}(\Delta) \subset B_0(\epsilon)$. Then from $g_{ij}(y) = \delta_{ij} + O(\epsilon^2)$, we have

$$A_g(\Delta) = \int_{\tilde{\Delta}} \sqrt{\det g(y)} dy = \int_{\tilde{\Delta}} 1 + O(\epsilon^2) dy = A_E(\tilde{\Delta}) + O(\epsilon^4).$$

Let x, y, z be three vertices of $\tilde{\Delta}$ in $T_p S$. Let $\gamma(t)$ be the geodesic parametrized with constant speed connecting $\exp(y) = q_0$ and $\exp(z) = q_1$ on (S, g) , $\sigma(t) = \exp^{-1}(\gamma(t))$ and $\eta(t)$ be the line segment parametrized with constant speed such that $\eta(0) = \sigma(0)$ and $\eta(1) = \sigma(1)$. Define $\phi(t) = \sigma(t) - \eta(t)$ as Lemma 5.2. In normal coordinates, the geodesic equation of σ is $\ddot{y}^k = -\Gamma_{ij}^k(y) \dot{y}^i \dot{y}^j$ with $\Gamma_{ij}^k(y) = O(|y|) = O(\epsilon)$ and $|\dot{y}| = O(\epsilon)$, so $|\sigma''(t)| = O(\epsilon^3)$ and Lemma 5.2 implies that $|\phi(t)| = O(\epsilon^3)$. Geometrically, it means that the curve $\sigma(t)$ lies in an $O(\epsilon^3)$ neighborhood of $\eta(t)$, which has the area $O(\epsilon^3|y - z|) = O(\epsilon^4)$. Therefore, let Δ_T be the Euclidean triangle determined by x, y, z in $T_p S$, then the $O(\epsilon^4)$ neighborhood of Δ_T contains $\tilde{\Delta}$, and

$$|A_g(\Delta) - A_E(\Delta_T)| \leq |A_g(\Delta) - A_E(\tilde{\Delta})| + |A_E(\tilde{\Delta}) - A_E(\Delta_T)| = O(\epsilon^4).$$

On the other hand, Δ_T does not have the same edge lengths as the geodesic triangle $\Delta = \Delta_{ijk}$, but their lengths differ by $O(\epsilon^3)$ by Lemma 5.1. Specifically, let (a, b, c) be the side lengths of Δ_T , so $a, b, c = O(\epsilon)$. The Euclidean triangle Δ_{ijk}^E has side lengths

$$L_{12}, L_{23}, L_{31} = (a, b, c) + (O(\epsilon^3), O(\epsilon^3), O(\epsilon^3))$$

We then compare the two Euclidean areas by Heron's formula:

$$A(a, b, c) = \sqrt{s(s-a)(s-b)(s-c)}, \quad s = (a+b+c)/2.$$

By the mean value theorem,

$$|A(L_{12}, L_{23}, L_{31}) - A(a, b, c)| \leq \sup |\nabla A| O(\epsilon^3).$$

One has $\frac{\partial A}{\partial a} = \frac{a}{2} \cot \alpha$ (and similarly for b, c), where α is the angle opposite to a . By the minimum angle bound $\alpha, \beta, \gamma \geq \pi/6$, $\cot \alpha$ is bounded and $|\nabla A| = O(\epsilon)$. Hence,

$$|A_E(\Delta) - A_E(\Delta_T)| \leq O(\epsilon) O(\epsilon^3) = O(\epsilon^4),$$

and therefore $|A_g(\Delta_{ijk}) - A_E(\Delta_{ijk})| = O(\epsilon^4)$.

The idea of the proof of the second inequality above is similar. The only difference is the definition of $\gamma(t)$. Let $\xi(t) : [0, 1] \rightarrow (S, g)$ be the geodesic connecting q_0 and q_1 above and $\gamma(t) = f(\xi(t))$ connecting two vertices $f(q_0)$ and $f(q_1)$. Notice that the bound on the derivatives of f implies that the length of $\gamma(t)$ is of order $O(\epsilon)$, so $|\gamma'(t)| = O(\epsilon)$ and $|\gamma''(t)| = O(\epsilon^2)$. Then repeat the computation above with Lemma 5.2 and we have $|A_h(f(\Delta)) - A_E(\Delta_T)| = O(\epsilon^3)$, where Δ_T is the Euclidean triangle in $T_{f(p)}S$ and $f(\Delta)$ is contained in some normal neighborhood of $f(p)$. $\square$

5.3 Convergence of the first term

In this section, we prove

Proposition 5.4. Let (S, g) and (S, h) be two Riemannian surfaces, and let $f : (S, g) \rightarrow (S, h)$ be a diffeomorphism. For every $\eta > 0$, there exists $\delta > 0$ such that, for every geodesic triangulation T of (S, g) with maximal edge length smaller than δ and minimum angle at least $\pi/5$,

$$\left| \int_S (\sqrt{P(x)} - 1)^2 dA_g - \sum_{\Delta_{ijk}} (1 - \sqrt{P(\Delta_{ijk}^E)})^2 A_E(\Delta_{ijk}) \right| < \eta.$$

Here, Δ_{ijk} is a geodesic triangle in T , $P(\Delta_{ijk}^E) = A_E(\Delta_{ijk}^h)/A_E(\Delta_{ijk})$, and $P = P(x) = \det(df_x)$ is the Jacobian.

To this end, we first estimate the pointwise difference between $P(x)$ and $P(\Delta_{ijk}^E)$ by the area estimation.

Lemma 5.5. Define

$$T_{12}(\Delta) = \frac{A_E(\Delta_{ijk}^h)}{A_E(\Delta_{ijk})} - \frac{A_h(f(\Delta_{ijk}))}{A_g(\Delta_{ijk})}.$$

Then $|T_{12}(\Delta_{ijk})| = O(\epsilon)$ for every Δ_{ijk} in T .

PROOF. Notice that

$$T_{12}(\Delta) = \frac{A_E(\Delta_{ijk}^h) - A_h(f(\Delta_{ijk}))}{A_E(\Delta_{ijk})} + \frac{A_h(f(\Delta_{ijk}))(A_g(\Delta_{ijk}) - A_E(\Delta_{ijk}))}{A_E(\Delta_{ijk})A_g(\Delta_{ijk})}.$$

By Lemma 5.3, $|A_E(\Delta_{ijk}^h) - A_h(f(\Delta_{ijk}))| \leq C\epsilon^3$ and $|A_g(\Delta_{ijk}) - A_E(\Delta_{ijk})| \leq C\epsilon^4$. Using the minimum angle bound of Δ_{ijk} , we have $A_E(\Delta_{ijk}) \geq C\epsilon^2$ and $A_g(\Delta_{ijk}) \geq C\epsilon^2$ for small ϵ . Then the first term is of order $C\epsilon^3/C\epsilon^2 = O(\epsilon)$. The second term is $O(\epsilon^2) \cdot \frac{C\epsilon^4}{C\epsilon^2 C\epsilon^2} = O(\epsilon^2)$. Hence $|T_{12}(\Delta)| = O(\epsilon)$. $\square$

Lemma 5.6. Define $T_3(\Delta_{ijk}) := \frac{A_h(f(\Delta_{ijk}))}{A_g(\Delta_{ijk})} - P(x_\Delta)$. If x_Δ is chosen as the barycenter of Δ_{ijk} with respect to dA_g , then $|T_3(\Delta)| = O(\epsilon^2)$ in Δ_{ijk} in the triangulation T .

PROOF. We write $\Delta = \Delta_{ijk}$ in this proof. Then $A_h(f(\Delta)) = \int_\Delta P dA_g$. Its Taylor expansion gives $P(x) = P(x_\Delta) + \nabla P(x_\Delta) \cdot (x - x_\Delta) + R_\Delta(x)$ with $|R_\Delta(x)| \leq C\epsilon^2$. Integrating over Δ , we have

$$A_h(f(\Delta_{ijk})) = \int_\Delta P dA_g = P(x_\Delta) A_g(\Delta) + \int_\Delta \nabla P(x_\Delta) \cdot (x - x_\Delta) dA_g + O(\epsilon^2) A_g(\Delta).$$

With the barycenter choice, $\int_\Delta (x - x_\Delta) dA_g = 0$, so the linear term vanishes and the result follows. $\square$

Corollary 5.7. There exists $C > 0$ such that for all $\Delta_{ijk} \in T$, $|P(\Delta_{ijk}^E) - P(x_\Delta)| \leq C\epsilon$.

PROOF. It is straightforward since $P(\Delta_{ijk}^E) - P(x_\Delta) = T_{12}(\Delta) + T_3(\Delta) = O(\epsilon)$. $\square$

Now we prove Proposition 5.4.

PROOF. Fix $\eta > 0$ and for each Δ , take $x_\Delta \in \Delta_{ijk}$ to be the barycenter. Notice that

$$\begin{aligned} \int_S (\sqrt{P(x)} - 1)^2 dA_g &= \sum_{\Delta_{ijk}} \int_{\Delta_{ijk}} (\sqrt{P(x)} - 1)^2 dA_g \\ &= \sum_{\Delta_{ijk}} \int_{\Delta_{ijk}} (\sqrt{P(x)} - 1)^2 - (\sqrt{P(x_\Delta)} - 1)^2 + (\sqrt{P(x_\Delta)} - 1)^2 dA_g. \end{aligned}$$

Since $P > 0$ is a smooth function on S , it is bounded away from zero and above, and $|P(x) - P(x_\Delta)| = O(\epsilon)$ for any $x \in \Delta_{ijk}$ contained in a normal ball of radius ϵ . Since $\Phi(t) = (1 - \sqrt{t})^2$ is Lipschitz on compact intervals in the positive real line, $(\sqrt{P(x)} - 1)^2 - (\sqrt{P(x_\Delta)} - 1)^2 = O(\epsilon)$. Therefore, we have

$$\begin{aligned} \left| \int_S (\sqrt{P(x)} - 1)^2 dA_g - \sum_{\Delta_{ijk}} \int_{\Delta_{ijk}} (\sqrt{P(x_\Delta)} - 1)^2 dA_g \right| \\ = O(\epsilon) \cdot \text{Area}_g(S) = O(\epsilon). \end{aligned}$$

Hence, we only need to estimate

$$\begin{aligned} \sum_{\Delta_{ijk}} (1 - \sqrt{P(\Delta_{ijk}^E)})^2 A_E(\Delta_{ijk}) - \sum_{\Delta_{ijk}} \int_{\Delta_{ijk}} (\sqrt{P(x_\Delta)} - 1)^2 dA_g \\ = \sum_{\Delta_{ijk}} (1 - \sqrt{P(\Delta_{ijk}^E)})^2 A_E(\Delta_{ijk}) - (\sqrt{P(x_\Delta)} - 1)^2 A_g(\Delta_{ijk}). \end{aligned}$$

For every term in the sum, from Corollary 5.7, we have

$$(1 - \sqrt{P(\Delta_{ijk}^E)})^2 A_E(\Delta_{ijk}) - (\sqrt{P(x_\Delta)} - 1)^2 A_g(\Delta_{ijk}) \quad (12)$$

$$= \left((1 - \sqrt{P(\Delta_{ijk}^E)})^2 - (\sqrt{P(x_\Delta)} - 1)^2 \right) A_E(\Delta_{ijk}) \quad (13)$$

$$+ (\sqrt{P(x_\Delta)} - 1)^2 (A_E(\Delta_{ijk}) - A_g(\Delta_{ijk})) \quad (14)$$

$$= O(\epsilon) O(\epsilon^2) + O(\epsilon^4) = O(\epsilon^3). \quad (15)$$

Take the sum over triangles and notice that the number of triangles in the triangulation is of order $O(\epsilon^{-2})$; we have

$$\begin{aligned} & \left| \sum_{\Delta_{ijk}} (1 - \sqrt{P(\Delta_{ijk}^E)})^2 A_E(\Delta_{ijk}) - \sum_{\Delta_{ijk}} \int_{\Delta_{ijk}} (\sqrt{P(x_\Delta)} - 1)^2 dA_g \right| \\ &= \sum_{\Delta_{ijk}} O(\epsilon^3) = O(\epsilon). \end{aligned}$$

Therefore,

$$\left| \int_S (\sqrt{P(x)} - 1)^2 dA_g - \sum_{\Delta_{ijk}} (1 - \sqrt{P(\Delta_{ijk}^E)})^2 A_E(\Delta_{ijk}) \right| = O(\epsilon).$$

This completes the proof. $\square$

5.4 Convergence of the second term

In this section, we prove

Proposition 5.8. Let (S, g) and (S, h) be two Riemannian surfaces, and let $f : (S, g) \rightarrow (S, h)$ be a diffeomorphism. For every $\eta > 0$, there exists $\delta > 0$ such that, for every geodesic triangulation T of (S, g) with maximal edge length smaller than δ and minimum angle at least $\pi/5$,

$$\left| \max_{\Delta_{ijk} \in T} |\log R(\Delta_{ijk})| - \sup_{x \in S} |\log R(x)| \right| < \eta.$$

Here, $R(\Delta_{ijk})$ is defined in Equation (3) and $R(x) = \lambda_1(x)/\lambda_2(x)$ is the ratio of singular values of df_x .

We have studied $P(x)$ and its discrete counterpart in the previous section, and now we discretize df_x and $D(x)$ using a dense geodesic triangulation. For any triangle Δ_{ijk} in a geodesic triangulation T of (S, g) , we approximate df_x in Δ_{ijk} by the differential at a point $x_\Delta \in \Delta_{ijk}$. Here x_Δ is not necessarily a barycenter.

Assume that every triangle of T is contained in some normal ball of radius $O(\epsilon)$. In normal coordinates at x_Δ , let ϕ_g be the exponential map at x_Δ on (S, g) . Let ϕ_h be the exponential map at $f(x_\Delta)$ on (S, h) . Set $\tilde{f} = \phi_h^{-1} \circ f \circ \phi_g$. Let y_i, y_j, y_k be the coordinates of the vertices of Δ_{ijk} via ϕ_g . The tangent triangle Δ_T has edges $u_1 = y_j - y_i$ and $u_2 = y_k - y_i$. Its image edges are $\tilde{u}_1 = \tilde{f}(y_j) - \tilde{f}(y_i)$ and $\tilde{u}_2 = \tilde{f}(y_k) - \tilde{f}(y_i)$. In matrix notation, the affine map F_Δ with $dF_\Delta[u_1 u_2] = [\tilde{u}_1 \tilde{u}_2]$ satisfies $dF_\Delta = G_T E_T^{-1} = J_T$. Here $G_T = [\tilde{u}_1 \tilde{u}_2]$ and $E_T = [u_1 u_2]$. The matrix J_T is the Jacobian, and the subscript T means that these quantities are on the tangent space $T_{x_\Delta} S$.

On the other hand, $R(\Delta_{ijk})$ is determined by an affine map between two Euclidean triangles with the same edge lengths of geodesic triangles Δ_{ijk} and Δ_{ijk}^h respectively. As in the smooth case in Equation (8), it can be computed from $D(\Delta_{ijk})$ and $P(\Delta_{ijk})$ by

$$Q(\Delta_{ijk}) = R(\Delta_{ijk}) + R^{-1}(\Delta_{ijk}) = \frac{D(\Delta_{ijk})}{P(\Delta_{ijk})}.$$

For this affine map between Euclidean triangles, the Dirichlet energy $D(\Delta_{ijk})$ [Pinkall and Polthier 1993] is

$$D(\Delta_{ijk}) = \frac{1}{2A_E(\Delta_{ijk})} (\cot \alpha l_1^2 + \cot \beta l_2^2 + \cot \gamma l_3^2),$$

where $A_E(\Delta)$ is the Euclidean area of the domain triangle with side geodesic lengths L_i , α, β, γ are the angles of the comparison triangle with side lengths L_i , and l_i are the geodesic lengths of image edges.

In our case, two Euclidean comparison triangles with side lengths L_i in Δ_{ijk} and l_i in Δ_{ijk}^h are embedded in $\mathbb{R}^2$. They determine another affine map given by the edge matrices $E_E = [v_1 v_2]$ and $G_E = [\tilde{v}_1 \tilde{v}_2]$ and its Jacobian $J_E = G_E E_E^{-1}$. Then $D(\Delta_{ijk}) = \|J_E\|_F^2$. Here J_E is the Jacobian, and the subscript E denotes quantities associated with the Euclidean comparison triangles. It remains to estimate the difference among $D(\Delta_{ijk}) = \|J_E\|_F^2$, $D(x) = \|df_x\|_F^2$, $\|dF_\Delta\|_F^2 = \|J_T\|_F^2$.

Lemma 5.9. Let Δ_{ijk} be a geodesic triangle with an edge length ϵ . In the notation above, $\|E_E - E_T\|_F = O(\epsilon^3)$, $\|E_E^{-1}\|_F = O(1/\epsilon)$ and $\|E_T^{-1}\|_F = O(1/\epsilon)$. Similar results hold for G_E and G_T as well.

PROOF. Recall that $E_E = [v_1 v_2]$ and $E_T = [u_1 u_2]$. By Lemma 5.1, we have $|u_1| = |v_1| + O(\epsilon^3)$, $|u_2| = |v_2| + O(\epsilon^3)$ and $|u_1 - u_2| = |v_1 - v_2| + O(\epsilon^3)$. Thus $\|E_E - E_T\|_F^2 = |u_1 - v_1|^2 + |u_2 - v_2|^2 = O(\epsilon^6)$, so $\|E_E - E_T\|_F = O(\epsilon^3)$. To bound the inverses, notice that

$$E_T = \begin{pmatrix} a & b \\ c & d \end{pmatrix}, \text{ and } E_T^{-1} = \frac{1}{\det E_T} \begin{pmatrix} d & -b \\ -c & a \end{pmatrix}.$$

Then

$$\|E_T^{-1}\|_F^2 = \frac{1}{|\det E_T|^2} (a^2 + b^2 + c^2 + d^2),$$

By the uniform lower bound on the angles of the triangle Δ_{ijk} , we have $|u_1| \geq C\epsilon$, $|u_2| \geq C\epsilon$, and $\sin \theta > C > 0$. Therefore,

$$|u_1|^2 = a^2 + c^2 = O(\epsilon^2) = |u_2|^2, \text{ and } |\det E_T| = |u_1||u_2| \sin \theta \geq C\epsilon^2,$$

and

$$\|E_T^{-1}\|_F^2 = O(1/\epsilon^2), \text{ and } \|E_E^{-1}\|_F^2 = O(1/\epsilon^2).$$

$\square$

Lemma 5.10. In the notation above, $\|J_E - J_T\|_F = O(\epsilon^2)$ and $\|dF_\Delta\|_F - \|df_{x_\Delta}\|_F = O(\epsilon)$ for any Δ_{ijk} in T .

PROOF. We have

$$J_E - J_T = G_E E_E^{-1} - G_T E_T^{-1} = G_E (E_E^{-1} - E_T^{-1}) + (G_E - G_T) E_T^{-1}.$$

By Lemma 5.9, $\|E_E - E_T\|_F = O(\epsilon^3)$, $\|E_T^{-1}\|_F = O(1/\epsilon)$, and $\|E_E^{-1}\|_F = O(1/\epsilon)$. Then

$$\begin{aligned} \|E_E^{-1} - E_T^{-1}\|_F &= \|E_E^{-1}(E_T - E_E)E_T^{-1}\|_F \\ &\leq \|E_E^{-1}\|_F \|E_T - E_E\|_F \|E_T^{-1}\|_F = O(1/\epsilon) \cdot O(\epsilon^3) \cdot O(1/\epsilon) = O(\epsilon). \end{aligned}$$

Therefore, $\|G_E(E_E^{-1} - E_T^{-1})\|_F = O(\epsilon) \cdot O(\epsilon) = O(\epsilon^2)$ and $\|(G_E - G_T)E_T^{-1}\|_F = O(\epsilon^3) \cdot O(1/\epsilon) = O(\epsilon^2)$. We have $\|J_E - J_T\|_F = O(\epsilon^2)$.

For the second estimate, the Taylor expansion in normal coordinates gives $\tilde{f}(y) = \tilde{f}(0) + d\tilde{f}_0 y + O(|y|^2)$. Apply it to y_i, y_j , we get $\tilde{u}_1 = \tilde{f}(y_j) - \tilde{f}(y_i) = d\tilde{f}_0(y_j - y_i) + O(\epsilon^2)$. Similarly, $\tilde{u}_2 = d\tilde{f}_{x_\Delta}(y_k - y_i) + O(\epsilon^2)$. Thus the matrix $G_T = [\tilde{u}_1 \tilde{u}_2]$ satisfies $G_T = d\tilde{f}_0(E_T) + O(\epsilon^2)$, hence $J_T = G_T E_T^{-1} = d\tilde{f}_{x_\Delta} + O(\epsilon^2)E_T^{-1}$. As shown in Lemma 5.9, $\|E_T^{-1}\|_F = O(1/\epsilon)$, so $\|J_T - d\tilde{f}_{x_\Delta}\|_F = O(\epsilon)$. Therefore $\|J_T\|_F - \|d\tilde{f}_{x_\Delta}\|_F \leq \|J_T - d\tilde{f}_{x_\Delta}\|_F = O(\epsilon)$. $\square$

Proposition 5.11. $|D(\Delta_{ijk}) - D(x_\Delta)| = O(\epsilon)$ and $|Q(\Delta_{ijk}) - Q(x_\Delta)| = O(\epsilon)$ for any triangle Δ_{ijk} in T , where $D(x_\Delta) = \|df_{x_\Delta}\|_F^2$ and $Q(x_\Delta) = D(x_\Delta)/P(x_\Delta) = R(x_\Delta) + R(x_\Delta)^{-1}$.

PROOF. Since $D(\Delta_{ijk}) = \|J_E\|_F^2$ and $D(x_\Delta) = \|df_{x_\Delta}\|_F^2$, we have $\|J_E\|_F^2 - \|df_{x_\Delta}\|_F^2 = (\|J_E\|_F - \|df_{x_\Delta}\|_F)(\|J_E\|_F + \|df_{x_\Delta}\|_F)$. By Lemma 5.10, $\|J_E - J_T\|_F = O(\epsilon^2)$, so $\|J_E\|_F = \|J_T + O(\epsilon^2)\|_F = \|J_T\|_F + O(\epsilon^2)$. Similarly, $\|J_T\|_F = \|df_{x_\Delta}\|_F + O(\epsilon)$. Thus $\|J_E\|_F - \|df_{x_\Delta}\|_F = O(\epsilon)$. Because f is a smooth function on S , $\|df_{x_\Delta}\|_F = O(1)$ and so $\|J_E\|_F + \|df_{x_\Delta}\|_F = O(1)$. Thus $|D(\Delta_{ijk}) - D(x_\Delta)| = O(\epsilon)$.

Meanwhile,

$$\begin{aligned} |Q(\Delta_{ijk}) - Q(x_\Delta)| &= \left| \frac{D(\Delta_{ijk})}{P(\Delta_{ijk})} - \frac{D(x_\Delta)}{P(x_\Delta)} \right| \\ &\leq \frac{|D(\Delta_{ijk}) - D(x_\Delta)| P(x_\Delta) + D(x_\Delta) |P(\Delta_{ijk}) - P(x_\Delta)|}{P(x_\Delta) P(\Delta_{ijk})}. \end{aligned}$$

By Corollary 5.7, $|P(\Delta_{ijk}) - P(x_\Delta)| = O(\epsilon^2)$. On S , P and D are bounded from below and above. Hence, $|Q(\Delta_{ijk}) - Q(x_\Delta)| = O(\epsilon)$. $\square$

Proposition 5.12. $|R(\Delta_{ijk}) - R(x_\Delta)| = O(\sqrt{\epsilon})$ and $|\log R(\Delta_{ijk}) - \log R(x_\Delta)| = O(\sqrt{\epsilon})$ for any triangle Δ_{ijk} in T .

PROOF. For $R \geq 1$, $Q = R + R^{-1} \geq 2$ with equality if and only if $R = 1$. Also, notice that $R(p)$ is bounded above by some $M > 0$ on S . The inverse function is $R = \phi(Q) = \frac{Q + \sqrt{Q^2 - 4}}{2}$ for $Q \geq 2$. We claim that if $y_1, y_2 \in [1, M]$ with $y_1 = \phi(x_1)$ and $y_2 = \phi(x_2)$, we have $|y_1 - y_2| \leq M\sqrt{|x_1 - x_2|}$. Then $|Q(\Delta_{ijk}) - Q(x_\Delta)| = O(\epsilon)$ in Proposition 5.11 implies that $|R(\Delta_{ijk}) - R(x_\Delta)| = O(\sqrt{\epsilon})$. Furthermore, for $a, b \geq 1$, $|\log a - \log b| \leq |a - b|$ by the mean value theorem, so $|\log R(\Delta_{ijk}) - \log R(x_\Delta)| = O(\sqrt{\epsilon})$.

It remains to show $|y_1 - y_2| \leq M\sqrt{|x_1 - x_2|}$. Assume without loss of generality that $y_2 \geq y_1$, then

$$x_2 - x_1 = (y_2 - y_1) \frac{y_1 y_2 - 1}{y_1 y_2}.$$

Since $y_1 \geq 1$, we have $y_1 y_2 - 1 \geq y_2 - y_1$, and since $y_1, y_2 \leq M$, we have $y_1 y_2 \leq M^2$. Therefore,

$$x_2 - x_1 \geq \frac{(y_2 - y_1)^2}{M^2}, \text{ and } y_2 - y_1 \leq M\sqrt{x_2 - x_1}.$$

$\square$

Now we are ready to prove Proposition 5.8.

PROOF. Let $\phi(x) := |\log R(x)|$ and $\phi(\Delta_{ijk}) = |\log R(\Delta_{ijk})|$. By Proposition 5.12, $|\phi(\Delta_{ijk}) - \phi(x_\Delta)| \leq |\log R(\Delta_{ijk}) - \log R(x_\Delta)| = O(\sqrt{\epsilon})$ for any Δ_{ijk} in T . Assume that $\phi(x)$ achieves its maximum at $x_0 \in \Delta_{ijk}$ for some Δ_{ijk} in T . Set $x_\Delta = x_0$ and $|\phi(\Delta_{ijk}) - \phi(x_\Delta)| = |\log R(\Delta_{ijk}) - \log R(x_\Delta)| = O(\sqrt{\epsilon})$. Meanwhile, the maximum over the finite set of triangles satisfies

$$\max_{\Delta} \phi(\Delta_{ijk}) = \max_{\Delta} \phi(x_\Delta) + O(\sqrt{\epsilon}) \leq \phi(x_0) + O(\sqrt{\epsilon}).$$

Hence, $|\max_{\Delta} \phi(\Delta_{ijk}) - \sup_{x \in S} \phi(x)| = O(\sqrt{\epsilon})$. $\square$

6 Limitations and Future Work

Theory of metrics on spaces of discrete shapes. It remains unclear whether the theory of metric structures by minimizing energies over diffeomorphisms for smooth surfaces has a direct discrete counterpart. A natural starting point is to define a discrete shape as a geometric realization of a fixed abstract triangulation. In this setting, $E_D(f)$ is well-defined for piecewise linear homeomorphisms f that

preserve the triangulation, and the three metric properties follow directly. This assumption on triangulations, however, is too restrictive for practical shape comparison, where shapes often have different triangulations or connectivity. If different meshes are allowed, the minimized discrete energies in opposite directions need not agree at finite resolution because the source triangulation determines the per-face sampling and area weights. Mesh quality, nonconvex optimization, approximate geodesics, and floating-point error can further disrupt symmetry and the triangle inequality. Developing a comprehensive theory of discrete shapes, maps, and metrics in this setting, including exact symmetry and the triangle inequality, therefore remains an open problem.

Applications to Complex Real-World Data. Our experiments focus on synthetic tori and scanned objects with clean genus-one topology. More complicated data, such as vertebrate bones [Bennani et al. 2016], introduce noise, self-intersections, incomplete surfaces, and topological issues that complicate preprocessing and comparison. In addition, meaningful comparison in such domains depends on biological or application-specific criteria beyond geometry. Extending the framework to these settings will require both more robust geometric processing and richer problem-dependent metrics.

Extension to Higher-Genus Surfaces. It is natural to extend the current framework to compare higher-genus surfaces, although the topology and geometry of surfaces with genus at least two introduce several difficulties. First, for a higher-genus surface, the construction of Dey et al. [2013] does not necessarily produce a system of loops with a single common base point; the homology generators in the resulting cut graph can instead have multiple base points. The chosen cut system determines the boundary and side pairings of the fundamental polygon and therefore changes the identification constraints that the optimization must preserve. Finding homology generators is not the principal obstacle, but a cut graph that works readily with our framework is not always available. Moreover, whereas our genus-one method uses a Euclidean square whose opposite sides are identified by translations, a natural construction for genus $g \geq 2$ uses a hyperbolic $4g$ -gon with paired sides related by Möbius transformations. Constructing this polygon, maintaining its side identifications, and optimizing consistently in the hyperbolic disk are the main computational complications. The smooth theory is established in [Luo 2021], but an efficient discrete implementation remains future work.

Extension to energies with extrinsic geometry. An interesting direction for future research is to add terms from extrinsic geometry, such as the second fundamental form and mean curvature, to our energy to produce a metric structure on the space of surface embeddings. However, ensuring that all three metric axioms hold is generally difficult. This extension requires quantifying how maps distort the surface normal at each point of a surface in $\mathbb{R}^3$. We would therefore need to consider more complicated maps, such as immersions of surfaces into $\mathbb{R}^3$, and their energies instead of diffeomorphisms between surfaces.

Computational Efficiency. The current implementation is effective but not yet competitive for large-scale implementation. Beyond the geodesic bottleneck (Section 4.7), possible improvements include GPU parallelization, faster spatial data structures for point location in the fundamental domain, and multiresolution optimization.

References

- S Mazdak Abulnaga, Oded Stein, Polina Golland, and Justin Solomon. 2023. Symmetric volume maps: Order-invariant volumetric mesh correspondence with free boundary. *ACM Transactions on Graphics* 42, 3 (2023), 1–20.
- Noam Aigerman, Shahar Z Kovalsky, and Yaron Lipman. 2017. Spherical orbifold Tutte embeddings. *ACM Trans. Graph.* 36, 4 (2017), 90.
- Noam Aigerman and Yaron Lipman. 2015. Orbifold tutte embeddings. *ACM Trans. Graph.* 34, 6 (2015), 190–1.
- Noam Aigerman and Yaron Lipman. 2016. Hyperbolic orbifold tutte embeddings. *ACM Trans. Graph.* 35, 6 (2016), 217–1.
- Hamza Bennani, Brendan McCane, and Jon Cornwall. 2016. Three dimensional (3D) lumbar vertebrae data set. *Data Science Journal* 15 (2016), 9–9.
- Alon Bright, Edward Chien, and Ofir Weber. 2017. Harmonic global parametrization with rational holonomy. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–15.
- Marcel Campen, Hanxiao Shen, Jiaran Zhou, and Denis Zorin. 2019. Seamless parametrization with arbitrary cones for arbitrary genus. *ACM Transactions on Graphics (TOG)* 39, 1 (2019), 1–19.
- Yves Colin de Verdiere and Marin Alexits. 1990. Triangulations presque équilatérales des surfaces. *Journal of Differential Geometry* 32.1 (1990), 199–207.
- Keenan Crane, Clarisse Weischedel, and Max Wardetzky. 2017. The heat method for distance computation. *Commun. ACM* 60, 11 (2017), 90–99.
- Tamal K Dey, Fengtao Fan, and Yusu Wang. 2013. An efficient computation of handle and tunnel loops via Reeb graphs. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 1–10.
- Jeff Erickson and Kim Whittlesey. 2005. Greedy optimal homotopy and homology generators. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms* (Vancouver, British Columbia) (SODA '05). Society for Industrial and Applied Mathematics, USA, 1038–1046.
- Michael S Floater and Kai Hormann. 2005. Surface parameterization: a tutorial and survey. *Advances in multiresolution for geometric modelling* (2005), 157–186.
- Steven J Gortler, Craig Gotsman, and Dylan Thurston. 2006. Discrete one-forms on meshes and applications to 3D mesh parameterization. *Computer Aided Geometric Design* 23, 2 (2006), 83–112.
- Xianfeng Gu and Shing-Tung Yau. 2003. Global conformal surface parameterization. In *Proceedings of the 2003 Eurographics/ACM SIGGRAPH symposium on Geometry processing*. 127–137.
- Emmanuel Hartman, Yashil Sukurdeep, Eric Klassen, Nicolas Charon, and Martin Bauer. 2023. Elastic shape analysis of surfaces with second-order Sobolev metrics: a comprehensive numerical framework. *International Journal of Computer Vision* 131, 5 (2023), 1183–1209.
- Joel Hass and Patrice Koehl. 2014. How round is a protein? Exploring protein structures for globularity using conformal mapping. *Frontiers in molecular biosciences* 1 (2014), 26.
- Joel Hass and Patrice Koehl. 2017. Comparing shapes of genus-zero surfaces. *Journal of Applied and Computational Topology* 1, 1 (2017), 57–87.
- Behrend Heeren, Martin Rumpf, Peter Schröder, Max Wardetzky, and Benedikt Wirth. 2014. Exploring the Geometry of the Space of Shells. *Computer Graphics Forum* 33, 5 (2014), 247–256. doi:10.1111/cgf.12450
- Behrend Heeren, Martin Rumpf, Max Wardetzky, and Benedikt Wirth. 2020. Discrete Riemannian Calculus on Shell Space. In *Handbook of Numerical Analysis*. Vol. 21. Elsevier, 621–679. doi:10.1016/bs.hna.2019.06.005
- Miao Jin, Junho Kim, Feng Luo, and Xianfeng Gu. 2008. Discrete surface Ricci flow. *IEEE Transactions on Visualization and Computer Graphics* 14, 5 (2008), 1030–1043.
- Martin Kilian, Niloy J. Mitra, and Helmut Pottmann. 2007. Geometric modeling in shape space. *ACM Trans. Graph.* 26, 3, Article 64 (July 2007), 8 pages. doi:10.1145/1276377.1276457
- Patrice Koehl and Joel Hass. 2013. Automatic alignment of genus-zero surfaces. *IEEE transactions on pattern analysis and machine intelligence* 36, 3 (2013), 466–478.
- Patrice Koehl and Joel Hass. 2015. Landmark-free geometric methods in biological shape analysis. *Journal of The Royal Society Interface* 12, 113 (2015), 20150795.
- Ka Chun Lam, Xianfeng Gu, and Lok Ming Lui. 2014. *Genus-One Surface Registration via Teichmüller Extremal Mapping*. Springer International Publishing, 25–32. doi:10.1007/978-3-319-10443-0_4
- Ka Chun Lam, Xianfeng Gu, and Lok Ming Lui. 2015. Landmark constrained genus-one surface Teichmüller map applied to surface registration in medical imaging. *Medical Image Analysis* 25, 1 (Oct. 2015), 45–55. doi:10.1016/j.media.2015.04.006
- Xin Li, Yunfan Bao, Xiaohu Guo, Miao Jin, Xianfeng Gu, and Hong Qin. 2008. Globally Optimal Surface Mapping for Surfaces with Arbitrary Topology. *IEEE Transactions on Visualization and Computer Graphics* 14, 4 (July 2008), 805–819. doi:10.1109/tvcg.2008.32
- Lok Ming Lui and Chengfeng Wen. 2014. Geometric Registration of High-Genus Surfaces. *SIAM Journal on Imaging Sciences* 7, 1 (Jan. 2014), 337–365. doi:10.1137/130932053
- Yanwen Luo. 2021. Comparing shapes of high genus surfaces. *Asian Journal of Mathematics* 25, 4 (2021), 579–596. https://intlpress.com/JDetail/1806251628829315074
- Facundo Mémoli. 2011. Gromov–Wasserstein distances and the metric approach to object matching. *Foundations of computational mathematics* 11, 4 (2011), 417–487.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems* 32. Curran Associates, Inc., 8024–8035. http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf
- Ulrich Pinkall and Konrad Polthier. 1993. Computing Discrete Minimal Surfaces and Their Conjugates. *Exp. Math.* 2 (1993), 15–36. https://api.semanticscholar.org/CorpusID:10068306
- Nicholas Sharp and Keenan Crane. 2020. You can find geodesic paths in triangle meshes by just flipping edges. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–15.
- Nicholas Sharp, Keenan Crane, et al. 2019. GeometryCentral: A modern C++ library of data structures and algorithms for geometry processing. https://geometry-central.net/.
- Justin Solomon, Gabriel Peyré, Vladimir G Kim, and Suvrit Sra. 2016. Entropic metric alignment for correspondence problems. *ACM Transactions on Graphics (ToG)* 35, 4 (2016), 1–13.
- William Thomas Tutte. 1963. How to draw a graph. *Proceedings of the London Mathematical Society* 3, 1 (1963), 743–767.
- Yuan-guang Wang. 2025. High genus surface parameterization using the Euclidean Ricci flow method. *Scientific Reports* 15, 1 (2025), 17784.
- Chengfeng Wen, Defeng Wang, Lin Shi, Winnie CW Chu, Jack CY Cheng, and Lok Ming Lui. 2015. Landmark constrained registration of high-genus surfaces applied to vestibular system morphometry. *Computerized Medical Imaging and Graphics* 44 (2015), 1–12.
- Tianqi Wu and Xiaoping Zhu. 2024. The convergence of discrete uniformizations for closed surfaces. *Journal of Differential Geometry* 127, 3 (2024), 1305–1343.
- Qingnan Zhou and Alec Jacobson. 2016. Thingi10K: A Dataset of 10,000 3D-Printing Models. arXiv:1605.04797
- Zhipeng Zhu, Wai Yeung Lam, and Lok Ming Lui. 2025. A Structure-Preserving Numerical Method for Harmonic Maps Between High-genus Surfaces. arXiv:2509.01256