

Dense Point-to-Point Correspondences Between Genus-Zero Shapes

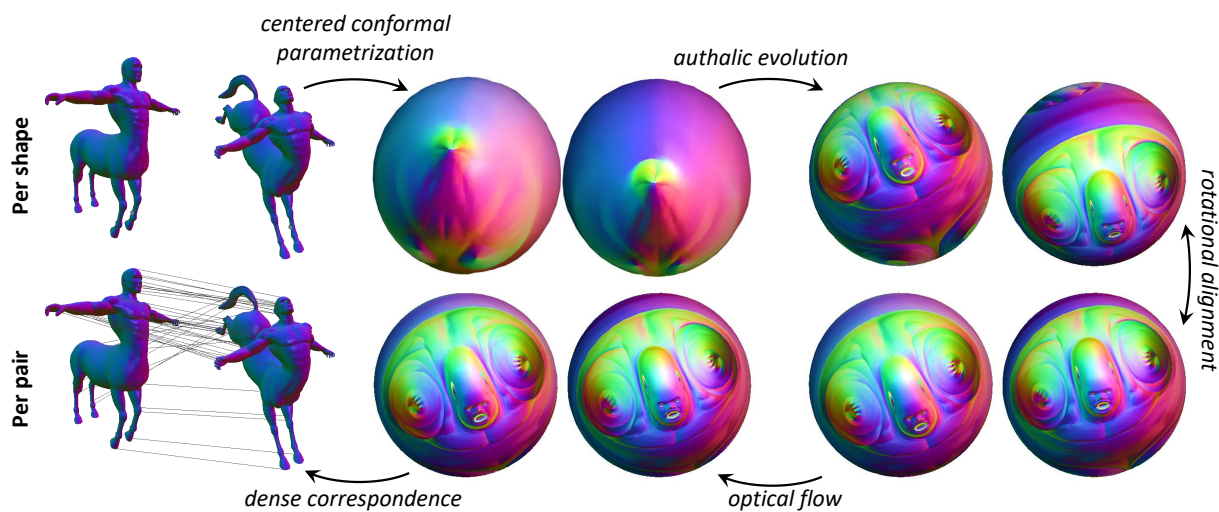
Sing Chun Lee¹ and Misha Kazhdan¹¹Johns Hopkins University, USA

Figure 1: Computing dense point-to-point correspondences: In a pre-processing step (top row), each input shape is conformally parametrized over the unit sphere, centered to remove ambiguity with respect to Möbius inversions, and evolved so the scale factors become more uniform. At run-time (bottom row), each pair of shapes is rotationally aligned using the Fast Fourier Transform and the registration is refined using optical flow. Correspondences (bottom left) are obtained by composing the inverse map from the source to the sphere, the spherical registration, and the forward map from the sphere to the target. Colors are set using normal orientation and are tracked with the parametrization.

Abstract

We describe a novel approach that addresses the problem of establishing correspondences between non-rigidly deformed shapes by performing the registration over the unit sphere. In a pre-processing step, each shape is conformally parametrized over the sphere, centered to remove Möbius inversion ambiguity, and authalically evolved to expand regions that are excessively compressed by the conformal parametrization. Then, for each pair of shapes, we perform fast $SO(3)$ correlation to find the optimal rotational alignment and refine the registration using optical flow. We evaluate our approach on the TOSCA dataset, demonstrating that our approach compares favorably to state-of-the-art methods.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems

1. Introduction

Non-rigid shape correspondence is a fundamental problem in geometry processing with applications in numerous areas of computer graphics. Such correspondences enable the transfer of animation, segmentation, texture, and other signals, between instances of the same object in different poses, and even between different objects.

In this work we introduce a novel approach for computing dense

point-to-point correspondences between genus-zero surfaces. We do this by parametrizing surfaces over the sphere, registering the spherical parametrizations, and defining the correspondences through the composition of the three maps – the inverse map taking the source to the sphere, the registration taking the sphere to itself, and the forward map taking the sphere to the target.

To do this we design a system that leverages and adapts ear-

lier work in shape matching, and image- and geometry-processing. In a pre-processing step we conformally map each shape to the sphere [KSBC12], center it so that the scale-factor-weighted center of mass is at the origin [BCK18], and evolve the parametrization to reduce the extreme area scaling at extremities [ZHG11] (Figure 1, top). Then, at run-time, we process pairs of shapes by first solving for the optimal rotation that registers the two spherical parametrizations [BCK18] and then refining the rigid registration by computing the optical flow advecting the parametrization towards each other [PKC*16] (Figure 1, bottom).

Combined, this gives an automatic, efficient, and robust method for mapping shapes to each other.

Outline We begin by reviewing related work (§2). Next, we describe our system (§3) and the details of its implementation (§4). We evaluate our approach in applications including non-rigid shape correspondence and intrinsic bilateral symmetry detection (§5). We conclude by summarizing our approach and describing possible directions for future work (§6).

2. Related Work

Shape correspondence has been extensively studied in geometry processing and computer vision. While earlier work focuses on rigid registration, more recent work focuses on establishing correspondences when shapes deform non-rigidly, and even non-isometrically. We briefly review related work. For more details, we refer the reader to recent surveys [vKZHC011, TCL*13, BCBB16].

Isometry invariant embedding One technique for establishing correspondences between surfaces deforming isometrically is to define an isometry-invariant embedding of the surface into a higher dimensional space. Then, rigid registration techniques like the Iterative Closest Points algorithm [BM92, RL01] can be used to define a mapping between the shapes.

Examples of such techniques include Generalized Multi-Dimensional Scaling [BBK06], which seeks an embedding with the property that Euclidean distances in the embedding space approximate geodesic distances on the surface, and methods like global point signatures [Rus07] and functional maps [OBCS*12, OCB*17], which leverage the (isometry-invariant) spectral decomposition of the Laplace-Beltrami operator to define the embedding function. More recently, the Functional Maps technique has been extended to partial matching using non-rigid puzzles [LRB*16], it has been extended to non-isometric matching using isospectralization [CPR*18], and the extracted mapping has been improved to preserve orientation and be approximately continuous [RPWO18].

Map refinement Another approach is to start with partial correspondence information, either in the form of sparse constraints or soft maps, and to transform these into dense point-to-point correspondences. Examples of such approaches include generalizing Euclidean affine combinations to surfaces [PBDSH13], optimizing the transport plan between surfaces [MCSK*17], and minimizing an energy that tries to preserve the given correspondences while penalizing the (approximate) Dirichlet energies of both the forward and inverse maps [ESBC19].

Shape Collections In addition to work focusing on establishing correspondences between pairs of shapes, there has been significant interest in establishing correspondences within a collection of shapes by leveraging algebraic structure like cycle-consistency to constrain the derived correspondences [NBCW*, KLM*12, HZG*12, HG13].

Machine learning Given the rapid development of GPUs in the last decade and the availability of copious ground-truth data, learning techniques like Deep functional maps [LRR*17] and 3D-CODED [GFK*18] have been proposed. Though these approaches do not directly provide geometric insight into the problem, they have been shown to be quite effective in practice, achieving state-of-the-art results on benchmarks like the FAUST dataset [BRLB14].

Conformal mapping Using the fact that isometric maps are conformal, recent techniques have addressed the correspondence problem by conformally parametrizing surfaces over a canonical base domain (e.g. the plane or the sphere) and searching for Möbius transformations of the base domain that register the parametrizations. In Möbius Voting [LF09] Lipman and Funkhouser use a Hough-style algorithm to vote for correspondences consistent with the assumption of isometry. In Blended Intrinsic Maps [KLF11], Kim and colleagues average together conformal deformations, using spatially varying weights that give precedence to a conformal map in regions where it is area-preserving. In both approaches, candidate Möbius transformations are obtained by choosing three point-pair correspondences.

An alternate approach for finding the Möbius transformation registering two conformal parametrizations was recently proposed by Baden *et al.* [BCK18]. Factoring the group of Möbius transformations into inversions and rotations, the authors show that parametrizations can be canonically centered with respect to inversions while the optimal rotation aligning two spherical parametrizations can be found using the Fast Fourier Transform on $SO(3)$.

Optical flow Though initially proposed for aligning images [LK81], optical flow has been extended to the sphere [YSV*10] and more general surfaces [PKC*16]. These extensions formulate the brightness constancy constraint in terms of local tangent spaces and iteratively update the deformation by using temporally varying tangent vectors fields and exponentiation to parametrize the change in deformation.

3. Our Approach

Given genus-zero surfaces, S and T , our goal is to compute dense point-to-point correspondences between the surfaces by defining a map $\Phi_{S \rightarrow T} : S \rightarrow T$. To this end, we first compute parametrizations, $\phi_S : S^2 \rightarrow S$ and $\phi_T : S^2 \rightarrow T$, of the two surfaces over the unit sphere and then compute a map $\psi_{S \rightarrow T} : S^2 \rightarrow S^2$ that registers the parametrizations. Using these, we define the map between the surfaces as the composition:

$$\Phi_{S \rightarrow T} = \phi_T \circ \psi_{S \rightarrow T} \circ \phi_S^{-1} : S \rightarrow T.$$

3.1. Parametrizing the surfaces

To parametrize a genus-zero surface over the unit sphere, we use conformalized Mean Curvature Flow [KSBC12]. This method evolves the surface using a simple non-linear PDE that is a hybrid of heat flow and Mean Curvature Flow. Empirically, the method quickly and robustly converges to a conformal parametrization of the surface over the sphere. We further normalize the parametrization for inversion using Möbius Centering [BCK18].




Figure 2: Correspondences between two models from TOSCA's Horse category after Möbius Registration: Using a conformal parametrization, the images of extremities like the head and legs are excessively compressed on the sphere (top). As the mapping between surfaces is not isometric, corresponding regions in the parameterizations become misaligned, resulting in poor registration, particularly for compressed extremities (bottom).

3.2. Registering the parametrizations

In theory, if the two surfaces are perfectly isometric, the conformal maps ϕ_S and ϕ_T will only differ by a Möbius transformation. Then the correspondence problem reduces to the problem of Möbius Registration [BCK18].

In practice however, the meshes we would like to register differ by more than an isometry. Combined with the fact that conformal parameterizations severely shrink extremities when mapping to the sphere (Figure 2, top), this results in poor point-to-point correspondences. Unfortunately, this problem persists even when the deformation between the surfaces is only slightly non-isometric because even small misalignments of the sphere result in significant misregistration of the compressed extremities. For example, note how points on the head of the horse on the left are not mapped well to the horse on the right in Figure 2 (bottom), despite the fact that the two horses are in similar poses.

We address this problem in four steps. First, we evolve the parametrizations to be more authalic (i.e. area-preserving), thereby decompressing the mapping of extremities to the sphere. Next, we define signals on the sphere by computing signals on the surfaces

and pulling them back via the parametrizations. We solve for the rotation aligning the signals and apply it to the spherical parametrization. Then, we refine the registration by computing the optical flow field between two (rotationally aligned) signals and advect the vertices of the spherical parametrization along it.

Authalic evolution

In order to resolve the artifacts arising due to excessive compression of extremities under the conformal mapping to the sphere, we follow the approach of Zou *et al.* [ZHG11] which advects a parametrization along a vector field in a manner that makes area more uniformly distributed.

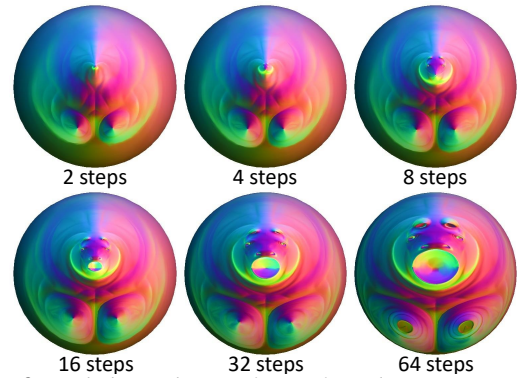


Figure 3: Authalic evolution of a conformal parametrization: As the evolution progresses, the parametrization becomes more area-preserving and extremities that were initially compressed by the conformal map become revealed.

Letting $\lambda : S^2 \rightarrow \mathbb{R}^{\geq 0}$ denote the scale factors of the parametrization, i.e. the ratio of the area of a local patch on the surface to the area of its image on the sphere, our goal is to find a vector field \vec{V} such that advection of the parametrization along \vec{V} results in diffusion of the scale factors. We do this by making two observations. On the one hand, the PDE for diffusion is:

$$\frac{d\lambda}{dt} = \Delta\lambda.$$

On the other, the infinitesimal change in λ resulting from advection along \vec{V} is given by the Lie-derivative on 2-forms:

$$\mathcal{L}_{\vec{V}}\lambda = -\nabla \cdot (\lambda\vec{V}).$$

Thus, the vector field whose advection results in the diffusion of the scale factors is defined by the linear system:

$$\Delta\lambda = -\nabla \cdot (\lambda\vec{V}).$$

Using the fact that the Laplacian is the divergence of the gradient, $\Delta = \nabla \cdot \nabla$, and leveraging the chain rule, this gives:

$$\vec{V} = -\frac{\nabla\lambda}{\lambda} = -\nabla \log \lambda.$$

That is, to diffuse the scale factors, we simply flow the parametrization away from regions on the sphere where area concentrates. The appearance of the logarithm is natural as it takes us from the multiplicative space of scale factors to the additive space of diffusion.

Figure 3 shows a visualization of authalic evolution for the conformal parametrization of the left horse in Figure 2.

Comparison to Zou et al. In the work of Zou *et al.*, the authors propose advecting the parametrization along a vector field in a way that makes the distribution of scale factors uniformly equal to one. To that end, they evolve the parametrization so the corresponding change in scale factors is proportional to the difference between the target and current scale factors, $\mathcal{L}_{\vec{v}}\lambda = 1 - \lambda$. Solving, this gives:

$$\vec{v} = -\frac{\nabla(\Delta^{-1}(1-\lambda))}{\lambda}.$$

This is a more complicated flow field and, in our implementation, we have found that trying to advect with this flow results in more unstable behavior, particularly for long timescale evolutions.

Surface signals

After authalic evolution, we can no longer use the scale factors for registration because the scale factors have been made uniform. Instead, we define new signals $f_S, f_T : S^2 \rightarrow \mathbb{R}$ by computing Heat Kernel Signatures [SOG09] on S and T , and pulling them back to the sphere. Figure 4 shows examples of these functions, pulled back to the sphere from two models in TOSCA’s *Victoria* category.

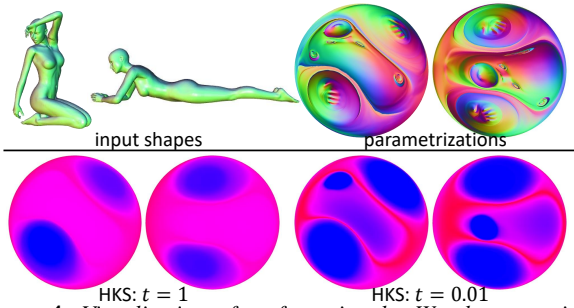


Figure 4: Visualization of surface signals: We show two input shapes from TOSCA’s *Victoria* category and their parametrization in the top row. In the bottom row we show the mapping of the HKS signals to the sphere for two different timescales, with smaller values drawn in red and larger values drawn in blue.

Rotational alignment

As in the work of Baden *et al* [BCK18], we find the rotation that best aligns the two parametrizations. Given functions $f_S : S^2 \rightarrow \mathbb{R}$ and $f_T : S^2 \rightarrow \mathbb{R}$, the optimal rotation is obtained by finding the rotation R^* maximizing the correlation:

$$R^* = \arg \max_{R \in SO(3)} \langle R(f_S), f_T \rangle_{S^2}$$

where the inner-product $\langle \cdot, \cdot \rangle_{S^2}$ is obtained by integrating the point-wise product of the functions over the sphere and $R(f) = f \circ R^{-1}$ is the rotation of the function f by R .

Optical flow Refinement

To refine the registration, we compute a vector field \vec{W} that advects two (rotationally-aligned) signals to each other. Using the approach of Prada *et al.* [PKC*16], we find \vec{W} hierarchically by using the Lie-derivative on 0-forms to linearize advection:

$$\mathcal{L}_{\vec{W}}f = -\langle \vec{W}, \nabla f \rangle.$$

As we expect the aligning deformation to be isometric and we start with authalic parametrizations of the surfaces over a sphere, we constrain \vec{W} to be divergence-free.

Figure 5 visualizes the contribution of the individual steps. Only using Möbius Registration to align the two spherical parametrizations (top row), extremities like the right thumb are severely compressed and mapped to the wrong location. Using authalic evolution (middle row), extremities are expanded and the resulting correspondence is improved. However, rotational alignment has too few degrees of freedom to resolve the fine details and local misregistration persists. Incorporating optical flow (bottom row) significantly improves the point-to-point correspondences.

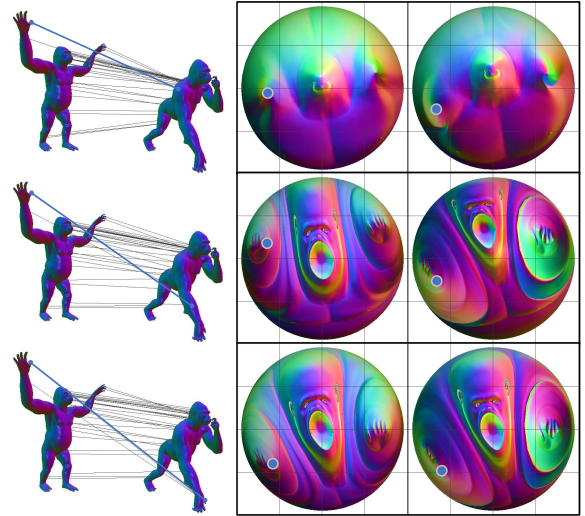


Figure 5: Contribution of the registration steps: Results obtained only using Möbius registration (top), adding authalic evolution (middle), and incorporating optical flow (bottom).

Putting these together, we get a composition of maps that are either computed per shape, using conformalized mean curvature flow and authalic evolution, or between pairs of shapes, using rotation alignment and optical flow. Since both maps between shapes are symmetric, our final map is as well:

$$\Phi_{S \rightarrow T}^{-1} = \Phi_{T \rightarrow S}.$$

4. Implementation

For the system to work well in practice, a number of design decisions need to be considered. These include efficient and robust implementation of authalic evolution and optical flow, choosing the signals used for guiding the alignment of the spherical parametrizations, and adaptation of rotational alignment to handle nearly symmetric surfaces. We discuss the implementation details below but defer discussion of individual parameter choices to the appendix. Prior to processing, we rescale all models to have unit area so that parameter settings are consistent.

4.1. Authalic evolution

To perform authalic evolution, we iteratively rasterize the scale factors of the current spherical parametrization, smooth the logarithm of the rasterized signal, compute its gradients, and use those to advect the vertices of the spherical triangulation.

Rasterizing the scale factors To compute the flow field, we rasterize the scale factors from the spherical triangulation onto an equirectangular $N \times N$ spherical grid. This makes it easy and efficient to trace geodesics for signal advection, a step that would otherwise become the bottleneck of the implementation.

Following the approach of [BCK18], we tessellate the spherical triangles using regularly spaced lines of longitude and latitude, and set the value in a spherical cell to the logarithm of the ratio of the area of the surface parametrized over the cell and the area of the cell. We define a vertex-based representation of the signal using area-weighted averaging from the faces of the spherical grid.

Smoothing In practice, the gradient of the log scale factors is high-frequency, limiting the step sizes we can take in performing advection. We address this by diffusing the signal in the spectral domain. That is, given the spherical harmonic decomposition of the log scale factors, $\log(\lambda(\theta, \phi)) = \sum_{l=0}^{N/2} \sum_{m=-l}^l a_{lm} Y_l^m(\theta, \phi)$, we obtain the diffused signal by dampening the spherical harmonic coefficients:

$$a_{lm} \mapsto e^{-l \cdot (l+1) \cdot \tau} a_{lm}$$

with τ the timescale of diffusion.

Computing the gradient We compute a face-based gradient field by solving for the tangent vector that best explains the differences along the edges of the face. As above, we sample the face-based gradient field back to the vertices using area-weighted averaging to obtain a piecewise bilinear flow-field \vec{V} .

Advecting the vertices Given the flow field \vec{V} , we advect each vertex $p \in S^2$ along \vec{V} to get its new position $q \in S^2$. We do this by taking a geodesic step from p along direction $\vec{v} \in T_p S^2$. That is, given the tangent vector \vec{v} , we exponentiate to get the position q :

$$q = \exp_p(\vec{v}) = p \cos(|\vec{v}|) + \vec{v}/|\vec{v}| \sin(|\vec{v}|).$$

We use a Runge-Kutta like approach, defining the tangent vector \vec{v} by constructing the geodesic curve from p in direction $\vec{V}(p)$ and averaging the values of \vec{V} at its endpoints. Formally, if we define $\gamma_p(t) = \exp_p(t \cdot \vec{V}(p))$ to be the geodesic starting at p and travelling in direction $\vec{V}(p)$ for a unit time, we define $\gamma_p^{-1}(t) = \gamma_p(1-t)$ to be the geodesic in the opposite direction, and we denote by Γ_γ the parallel transport operator from $T_{\gamma(0)} S^2$ to $T_{\gamma(1)} S^2$ along the curve γ , then the flow direction is expressed as:

$$\vec{v} = \frac{\vec{V}(\gamma_p(0)) + \Gamma_{\gamma_p^{-1}}(\vec{V}(\gamma_p(1)))}{2}.$$

Parameters We set the grid resolution to $N = 256$, the timescale of diffusion to $\tau = \frac{1}{4} \times 10^{-2}$, and run 100 time-steps of authalic flow. Surprisingly, despite the relatively coarse sampling of the sphere, the authalic flow resolves small features at the extremities (like fingers) that are initially mapped to regions of the sphere that are smaller than the size of a cell in a 256×256 grid.

4.2. Surface signals

To compute the Heat Kernel Signatures for a surface S , we compute the first 200 eigenvectors of the generalized eigenvalue problem $Sv_i = \lambda M v_i$, with M and S the mass and stiffness matrices and set the HKS at vertex v and time t to:

$$\text{HKS}_t(v) = \sum_i e^{-\lambda_i t} \phi_i^2(v).$$

4.3. Rotational alignment

In practice, many of the shapes we consider have near 180° rotational symmetry. For example, humans are nearly symmetric with respect to rotation about the vertical axis. As we depend on rotational alignment to bootstrap registration, a bad choice of initial rotation can result in poor point-to-point correspondences. To address this, we compute the, at most, K best aligning rotations, $R_1, \dots, R_K \in SO(3)$, subject to the constraints that the rotations are separated, $\|R_j^{-1} \cdot R_k - \text{Id}\|_F^2 > \epsilon_1$, and their correlation values are relatively large, $\langle R_j(f_S), f_T \rangle > \langle R_1(f_S), f_T \rangle \cdot (100\% - \epsilon_2)$. Then, we perform optical flow registration using all candidate rotations and return the best one.

Parameters We set the maximum number of candidates to $K = 4$, the separation distance to $\epsilon_1 = 3$, and the correlation quality threshold to $\epsilon_2 = 10\%$. As we do not expect the initial rotation to register high-frequency detail, we perform the rotational alignment using a smooth signal, setting the timescale for the HKS to $t_R = 1.0$. In computing the correlation, we also shift the signals f_S and f_T so that they are zero-mean and the correlation values are not dominated by the rotation-invariant DC terms.

4.4. Optical flow refinement

To compute the optical flow field, we follow the approach of Prada *et al.*, solving for the vector field \vec{W} minimizing a combination of a symmetric fitting energy and a smoothness energy:

$$E(\vec{W}) = \int_{S^2} \sum_{i \in \{S, T\}} \left(\langle \nabla f_i, \vec{W} \rangle - (f_S - f_T) \right)^2 + \alpha \|\nabla \vec{W}\|^2.$$

That approach begins by choosing a (hat) basis $\mathcal{B} = \{b_1, \dots, b_n\}$ for representing signals and then discretizes the space of vector fields using the gradients $\mathcal{V}^\nabla = \{\nabla b_1, \dots, \nabla b_n\}$ and 90° -rotated gradients $\mathcal{V}^{\nabla} = \{J\nabla b_1, \dots, J\nabla b_n\}$ of the scalar basis functions, with J the operator performing a 90° counter-clockwise rotation in each tangent space.

Using this basis, the authors use a hierarchical (coarse-to-fine) approach to solve the optical flow problem. As meshes do not generally come with a multiresolution structure, the solver is implemented by smoothing the signals f_S and f_T in a level-adapted manner (with more smoothing at coarser resolutions) and adapting the smoothness weight α to the level (using larger weights at coarser resolutions). We modify this implementation in several ways.

Discretization We perform the computation using regular $N \times N$ spherical grids and rasterizing the signals f_S and f_T as we do for

authalic evolution. Using this discretization we obtain a basis \mathcal{B} consisting of piecewise bilinear functions on the quad elements.

Expressing f_S and f_T in this basis, it is trivial to discretize the gradients of f_S and f_T in terms of \mathcal{V}^∇ . Noting that the 90° -rotated gradients of the basis functions are divergence-free, we discretize the optical flow field using just \mathcal{V}^∇ .

Hierarchy As we use $N \times N$ grids, defining a (power-of-two) discretization hierarchy is straight-forward. However, to avoid aliasing artifacts due to the rasterization of the signals f_S and f_T , we use smoother signals at coarser levels of the hierarchy.

Rotation candidate selection Given the, at most, K optical flow fields corresponding to the different candidate rotations, we select the best one by choosing the flow that best registers two high-frequency HKSs pulled back to the sphere.

Parameters We use a four-level hierarchy and perform the optical flow using Heat Kernel Signatures with timescales $t_0 = 1/2$, $t_1 = 1/4$, $t_2 = 1/8$, and $t_3 = 1/16$ at resolutions 16×16 , 32×32 , 64×64 , and 128×128 respectively. We use a fixed smoothness weight of $\alpha = 0.05$ for all levels of the hierarchy. We select the best candidate flow using an HKS with timescale $t_S = 0.01$ and measuring the distance between registered signals using the L_1 -norm.

5. Results

We tested our approach on the TOSCA benchmark [BBK08], which contains 80 models across 9 categories. Within each category, models are consistently tessellated, giving ground truth vertex-to-vertex correspondences. We pre-process the database by triangulating holes to make the meshes water-tight and keeping the largest connected component when a model is disconnected.

We evaluate our approach in defining correspondences between shapes in the same category and in detecting intrinsic bilateral symmetries. We conclude by analyzing the performance of our method and discussing limitations.

5.1. Intra-class correspondences

We begin by considering the problem of computing correspondences between a single shape in different poses. For this evaluation we compare our method, both with (Ours+) and without (Ours-) optical flow refinement, to Blended Intrinsic Maps (BIM) [KLF11] and Möbius Registration (MR) [BCK18].

Given models S and T , from the same category we compare a candidate correspondence $\Phi_{S \rightarrow T} : S \rightarrow T$ to the ground-truth correspondence $\Phi_{S \rightarrow T}^* : S \rightarrow T$ by measuring the distance between the images of points $p \in S$ under the two correspondences:

$$D_{S \rightarrow T}(p) = d_T(\Phi_{S \rightarrow T}(p), \Phi_{S \rightarrow T}^*(p)),$$

with $d_T : T \times T \rightarrow \mathbb{R}$ a measure of proximity on T . As we do not expect the correspondences to be rigid, we cannot use Euclidean distance for proximity. Instead, we use (intrinsic) biharmonic distances [LRF10] as an efficient proxy for geodesic distances:

$$d_T(p, q) = \sqrt{\sum_{i=1}^{\infty} \frac{(\phi_i(p) - \phi_i(q))^2}{\lambda_i^2}},$$

with (λ_i, ϕ_i) the (ordered) eigenvalues/eigenvectors of the Laplace-Beltrami operator.




Figure 6: Distribution of biharmonic distances across the computed intra-category correspondences (left) and distribution of geodesic distances across the computed bilateral symmetries (right).

We make the measure symmetric in S and T by integrating the pointwise distances in both directions:


$$D(S, T) = \sqrt{\int_S D_{S \rightarrow T}^2(p) dp + \int_T D_{T \rightarrow S}^2(p) dp}.$$

Figure 6 (left) shows the cumulative distribution of biharmonic distances, averaged over all intra-category pairs of models. As expected, Möbius Registration (MR) exhibits poor performance due to the significant area compression resulting from the conformal parametrization. This is alleviated somewhat by performing authalic evolution (Ours-). Incorporating optical flow (Ours+) we obtain results on par with those of Blended Intrinsic Maps (BIM), which avoids the problem of compression by giving precedence to conformal maps which are locally area preserving.

Figure 7 shows several example correspondences found using BIM and our method. We found that both approaches can get confused by intrinsic symmetry. This appears to be worse for BIM because the computed mapping is not always continuous, so even partial symmetries can be problematic. We also found that BIM's approximation of the geodesic centroid by projection of the Euclidean centroid causes artifacts when a mesh self-intersects. For our approach, the initial spherical parametrizations varies with non-isometric deformations and significant anisotropic scaling after authalic evolution can imply that flow fields that appear smooth on the sphere may be high-frequency when pulled back to the surface.

Authalic evolution We validate the effectiveness of authalic evolution in defining shape correspondences, noting significant improvement both with and without optical flow.

Representative results are shown in the inset which plots the distance, $D(S, T)$ for correspondences computed between a pair of models in the *Gorilla* category. As expected, refining the registration using optical flow provides significantly better results, and the average biharmonic distance decreases as more steps of authalic evolution are performed.



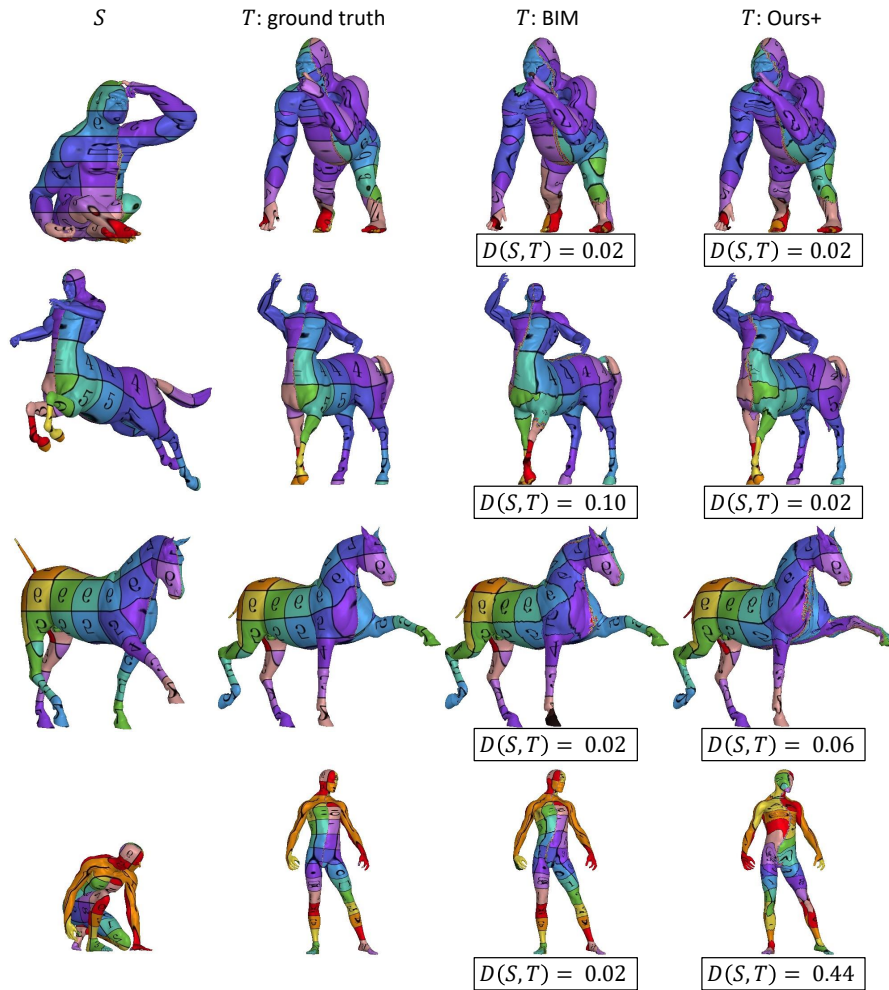


Figure 7: Example correspondences found using BIM and our method. Mapping quality is measured using biharmonic distances. The texture mapping is made discontinuous across the bilateral symmetry so that symmetric regions are texture mapped differently.

5.2. Intrinsic bilateral symmetry

We also consider the problem of intrinsic bilateral symmetry detection. For this evaluation we compare our method (Ours+) to Blended Intrinsic Maps (BIM) [KLF11], bijective and continuous ICP (BCICP) [RPWO18], and Reversible Harmonic Maps (RHM) [ESBC19].

As RHM and BCICP do not scale well to meshes at the resolution of the TOSCA dataset, we down-sample the meshes to 5K vertices. In our experiments, when registering meshes from the *David* category with 53K vertices, we found that RHM took almost 3.5 hours, and BCICP ran out of memory before it could complete.

We obtain ground-truth correspondences by observing that, with the exception of *Gorilla*, the other eight TOSCA categories each contain a representative having perfect extrinsic bilateral symmetry with respect to the yz -plane. Using this representative we randomly sample 100 vertices and, for each vertex, find the closest vertex on the reflection of the mesh through the yz -plane. Since models within each category are consistently triangulated, this also gives ground truth correspondences for the other models in the category.

To measure bilateral symmetry, we take every model in the database and generate its antipode. That is, we replace each vertex v with $-v$ and reverse the ordering of vertices in the triangles so that the new mesh is properly oriented. We then compute the correspondence between the initial mesh and its antipode and measure how close the correspondence is to a bilateral symmetry. As the antipodal mapping is orientation-reversing and BIM, BCPIP, and our method are orientation-preserving, we expect the best correspondence to be the map taking a point to its (intrinsic) reflection. For RHM, which requires initial correspondences, we sample additional bilaterally symmetric correspondences on the representative mesh and provide those as input.

Figure 6 (right) plots the distribution of the measures of bilateral symmetry across the shapes in the database. As the evaluation is performed using only the 100 selected correspondences, we use geodesic distances [QHY*16] to measure point proximity. For RHM, we compute results using 10, 20, 40, and 80 point pairs (chosen independently of the 100 correspondences used for testing).

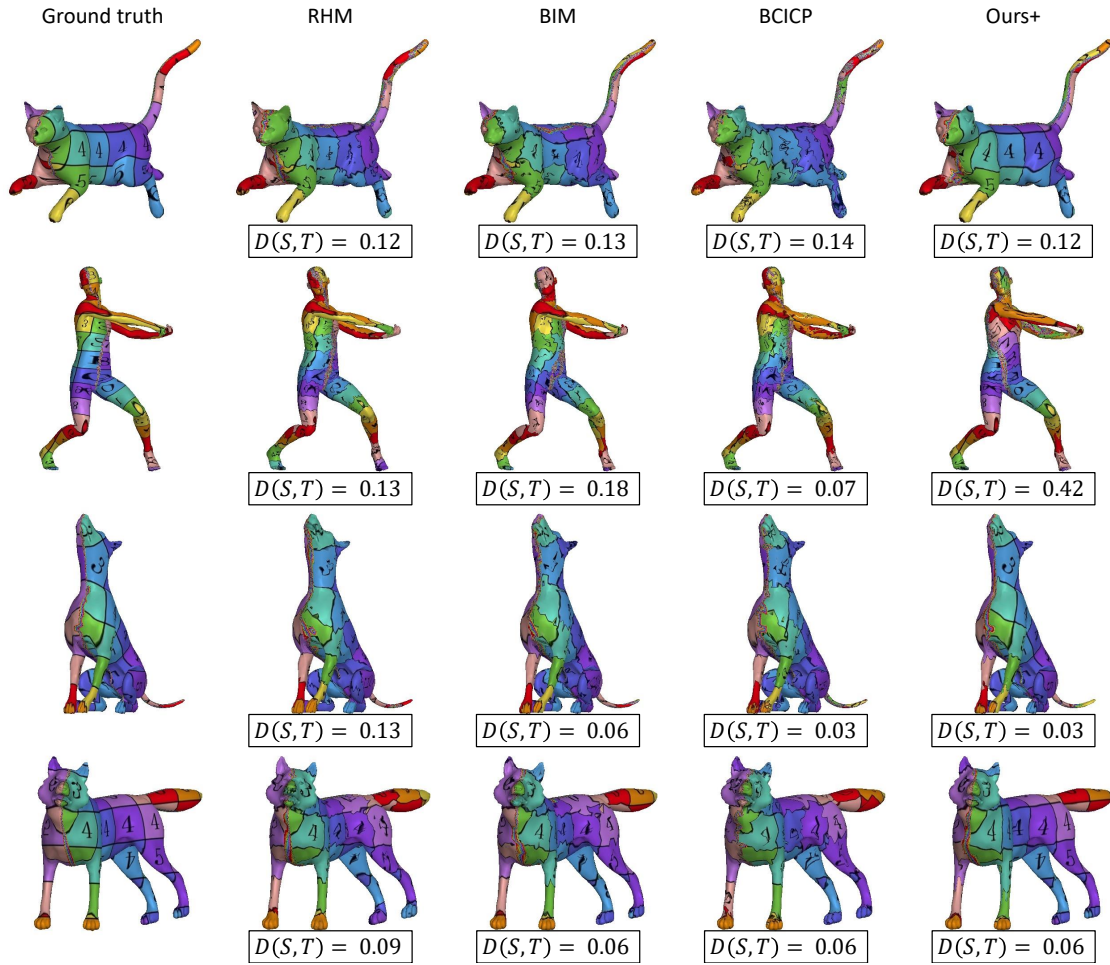


Figure 8: Example bilateral symmetries found using RHM, BIM, BCICP, and our method. Mapping quality is computed using geodesic distances from the 100 landmark pairs. The texture mapping is made discontinuous across the bilateral symmetry so that symmetric regions are texture mapped differently.

The results are quite similar and the figure only shows the (best) results using 80 point-pair correspondences.

As the figure shows, our approach remains comparable to BIM which is better than RHM though not as good as BCICP. Figure 8 shows several example bilateral symmetries found using the four methods. We found that RHM, BIM, and BCICP tend to generate mappings that are more uniformly “jaggy” while our approach generates mappings that are smoother in regions where the spherical parametrization exhibits less shearing and more jaggy in regions of significant anisotropic scaling.

5.3. Performance

Our approach can be decomposed into a pre-processing step, consisting of computation that can be performed per model and a run-time step, consisting of computation that is performed per pair of models. We evaluate the performance of our method by measuring running times on an Intel Core i7-6700HQ 2.6GHz processor. We use the Eigen [JG08] and Spectra [Qiu15] libraries for solving sparse linear systems and generalized eigenvalue problems.

As all models within a single category of the TOSCA dataset have the same size, Table 1 gives running times for each category. The table decomposes the running times into the separate phases and provides the running times for Blended Intrinsic Maps. Comparing, we find that it would take under 4 hours to compute pairwise correspondences across all pairs within all categories of the TOSCA dataset using our method and roughly 31 hours using BIM.

The table also highlights the dependency of our running time on both the mesh resolution and the resolution of the sphere. For steps like spherical parametrization (SP) and function computation (FC), run-time is roughly linear in the number of vertices. For steps like authalic evolution (AE), optical flow (OF), and rotational alignment (RA), which are performed over a regular spherical grid, the vertex count is less significant, only playing a role in the rasterization of the spherical mesh onto the equirectangular grid.

For bilateral symmetry detection, which was run on lower-resolution models, the running times are more comparable. In this experiment RHM, BIM, BCICP and our method take, on average, 170, 270, 110 and 45 seconds to find the bilateral symmetry.

Model	Vertices	SP	AE	FC	RA	OF	BIM
Wolf	4K	0.65	15.2	2.9	0.6	5.4($\times 1.0$)	184
Centaur	16K	2.07	21.8	14.1	0.7	6.0($\times 1.0$)	1129
Horse	19K	2.63	22.0	18.3	0.7	6.0($\times 1.8$)	311
Dog	25K	3.64	25.1	26.2	1.0	6.2($\times 1.6$)	883
Gorilla	25K	3.35	24.0	24.7	0.6	6.1($\times 1.5$)	94
Cat	28K	4.54	26.7	31.4	0.6	6.4($\times 4.0$)	294
Victoria	46K	6.78	34.3	50.9	0.9	7.3($\times 2.0$)	130
David	53K	7.97	38.7	59.3	1.0	7.2($\times 2.0$)	148
Michael	53K	7.96	35.7	59.3	0.9	7.3($\times 2.0$)	133


Table 1: Per model performance, given in seconds, for the different steps of registration: The individual steps include the per models steps of spherical parametrization (SP), authalic evolution (AE), and function computation (FC), as well as the per pair of models steps of rotational alignment (RA) and optical flow (OF). We also provide the running times for Blended Intrinsic Maps (BIM). The optical flow time represents the time for a single candidate rotation, with the number in parentheses giving the average number of different rotations tried for that category.

5.4. Limitations

While our approach is efficient and gives a symmetric and orientation-preserving mapping between surfaces, it suffers from several limitations.

Poor alignment The optical flow step is dependent on the initial choice of rotation. This can fail if we do not select the correct candidate. An example of this can be seen in Figure 7 (bottom row) which shows two models from the *Michael* category. In this case, the L_1 -distance between HKSs for the selected rotation was 8.16 while the L_1 -distance for the rotation minimizing the biharmonic distance was 8.18. Had the algorithm chosen the correct rotation, the biharmonic distance would have gone down from 0.44 to 0.03.

Another issue can be seen in the bilateral symmetry example visualized in Figure 8 (second row). Because of the non-isometric twist in the pose of the model, the spherical parametrization (inset) twists the torso by an angle larger than 45° . Thus, when the initial rotation aligns the left leg of the model to the right leg of the antipode (bottom), the right arm of the model is at an angle greater than 90° from the left arm of the antipode (top). As a result, it is easier for optical flow to continue twisting the torso by an angle less than 90° , incorrectly matching the right arm on the model to the right arm on the antipode, than it is to (un)rotate by an angle greater than 90° .



Theoretical limitations The dependence of our method on the initial parametrization to the sphere restricts our method to genus-zero surfaces. Additionally, in performing the optical flow over the sphere, we ignore the metric distortion of the parametrization. Although authalic evolution mitigates the effects of this problem by encouraging the parametrization to be area-preserving, the mapping can still have significant anisotropic stretching. As mentioned above, in such regions the flow-field may appear smooth with re-

spect to the metric on the sphere but may be high-frequency when pulled back to the original surface.

Discretization concerns Finally, though the expression of the correspondence map as a combination of rotation and authalic/optical flows implies a continuous map in theory, in practice artifacts due to discretization of advection can result in a map that is not.

6. Conclusion and Future Work

In this work we propose a novel and fully automatic system for finding point-to-point correspondences between genus-zero surfaces using a combination of conformal parameterization, Möbius registration, authalic evolution, rotational correlation, and optical flow. We compare to state-of-the-art methods and show that our approach provides good correspondences in significantly less time.

In the future, we would like to consider several extensions. We would like to consider adaptations of authalic evolution that can provide still more robust parameterizations in the case of significant non-isometric deformation and we would also like to explore performing optical flow on the sphere using the pull-back metric from the parametrized surface. In addition, while our current implementation only considers HKS signals, we would like to explore using other signals as well.

References

- [BBK06] BRONSTEIN A., BRONSTEIN M., KIMMEL R.: Generalized multidimensional scaling: A framework for isometry-invariant partial surface matching. *Proc. National Academy of Sciences* 103, 5 (2006), 1168–1172. 2
- [BBK08] BRONSTEIN A., BRONSTEIN M., KIMMEL R.: *Numerical Geometry of Non-Rigid Shapes*, 1 ed. Springer Publishing Company, Incorporated, 2008. 6
- [BCBB16] BIASOTTI S., CERRI A., BRONSTEIN A., BRONSTEIN M.: Recent trends, applications, and perspectives in 3D shape similarity assessment. *Comp. Graph. Forum* 35, 6 (2016), 87–119. 2
- [BCK18] BADEN A., CRANE K., KAZHDAN M.: Möbius Registration. *Comp. Graph. Forum* (2018). 2, 3, 4, 5, 6
- [BM92] BESL P., MCKAY N.: A method for registration of 3D shapes. *IEEE Trans. Pattern Anal. and Machine Intelligence* 14, 2 (1992), 239–256. 2
- [BRLB14] BOGO F., ROMERO J., LOPER M., BLACK M.: FAUST: Dataset and evaluation for 3D mesh registration. In *Proc. IEEE Conf. on Comp. Vis. and Pattern Recognition* (2014). 2
- [CPR*18] COSMO L., PANINE M., RAMPINI A., OVSIANIKOV M., BRONSTEIN M., RODOLÁ E.: Isospectralization, or how to hear shape, style, and correspondence. [arXiv:1811.11465](https://arxiv.org/abs/1811.11465). 2
- [ESBC19] EZUZ D., SOLOMON J., BEN-CHEN M.: Reversible harmonic maps between discrete surfaces. *ACM Trans. Graph.* 38, 2 (2019), 15:1–15:12. 2, 7
- [GFK*18] GROUEIX T., FISHER M., KIM V. G., RUSSELL B., AUBRY M.: 3D-CODED: 3D correspondences by deep deformation. In *European Conf. Comp. Vis.* (2018). 2
- [HG13] HUANG Q.-X., GUIBAS L.: Consistent shape maps via semidefinite programming. *Comp. Graph. Forum* 32, 5 (2013), 177–186. 2
- [HZG*12] HUANG Q.-X., ZHANG G.-X., GAO L., HU S.-M., BUTSCHER A., GUIBAS L.: An optimization approach for extracting and encoding consistent maps in a shape collection. *ACM Trans. Graph.* 31, 6 (2012), 167:1–167:11. 2

- [JG08] JACOB B., GUENNEBAUD G.: Eigen. <http://eigen.tuxfamily.org/>, 2008. 8
- [KLF11] KIM V. G., LIPMAN Y., FUNKHOUSER T.: Blended intrinsic maps. *ACM Trans. Graph.* 30, 4 (2011), 79:1–79:12. 2, 6, 7
- [KLM*12] KIM V. G., LI W., MITRA N. J., DIVERDI S., FUNKHOUSER T.: Exploring collections of 3D models using fuzzy correspondences. *ACM Trans. Graph.* 31, 4 (2012), 54:1–54:11. 2
- [KSBC12] KAZHDAN M., SOLOMON J., BEN-CHEN M.: Can mean-curvature flow be modified to be non-singular? *Comp. Graph. Forum* 31, 5 (2012), 1745–1754. 2, 3
- [LF09] LIPMAN Y., FUNKHOUSER T.: Möbius voting for surface correspondence. *ACM Trans. Graph.* 28, 3 (2009), 72:1–72:12. 2
- [LK81] LUCAS B. D., KANADE T.: An iterative image registration technique with an application to stereo vision. In *Proc. Int. Joint Conf. on Artificial Intelligence* (1981), pp. 674–679. 2
- [LRB*16] LITANY O., RODOLÁ E., BRONSTEIN A. M., BRONSTEIN M. M., CREMERS D.: Non-rigid puzzles. *Comp. Graph. Forum* 35, 5 (2016), 135–143. 2
- [LRF10] LIPMAN Y., RUSTAMOV R. M., FUNKHOUSER T. A.: Biharmonic distance. *ACM Trans. Graph.* 29, 3 (2010), 27:1–27:11. 6
- [LRR*17] LITANY O., REMEZ T., RODOLÁ E., BRONSTEIN A., BRONSTEIN M.: Deep functional maps: Structured prediction for dense shape correspondence. In *IEEE Int. Conf. Comp. Vis.* (2017), pp. 5660–5668. 2
- [MCSK*17] MANDAD M., COHEN-STEINER D., KOBELT L., ALLIEZ P., DESBRUN M.: Variance-minimizing transport plans for inter-surface mapping. *ACM Trans. Graph.* 36, 4 (2017), 39:1–39:14. 2
- [NBCW*] NGUYEN A., BEN-CHEN M., WELNICKA K., YE Y., GUIBAS L.: An optimization approach to improving collections of shape maps. *Comp. Graph. Forum* 30, 5, 1481–1491. 2
- [OBCS*12] OVSJANIKOV M., BEN-CHEN M., SOLOMON J., BUTSCHER A., GUIBAS L.: Functional maps: A flexible representation of maps between shapes. *ACM Trans. Graph.* 31, 4 (2012), 30:1–30:11. 2
- [OCB*17] OVSJANIKOV M., CORMAN E., BRONSTEIN M., RODOLÁ E., BEN-CHEN M., GUIBAS L., CHAZAL F., BRONSTEIN A.: Computing and processing correspondences with functional maps. In *ACM SIGGRAPH 2017 Courses* (2017), SIGGRAPH '17, pp. 5:1–5:62. 2
- [PBDSh13] PANOZZO D., BARAN I., DIAMANTI O., SORKINE-HORNUNG O.: Weighted averages on surfaces. *ACM Trans. Graph.* 32, 4 (2013), 60:1–60:12. 2
- [PKC*16] PRADA F., KAZHDAN M., CHUANG M., COLLET A., HOPPE H.: Motion graphs for unstructured textured meshes. *ACM Trans. Graph.* 35, 4 (2016), 108:1–108:14. 2, 4, 11
- [QHY*16] QIN Y., HAN X., YU H., YU Y., ZHANG J.: Fast and exact discrete geodesic computation based on triangle-oriented wavefront propagation. *ACM Trans. Graph.* 35, 4 (2016), 125:1–125:13. 7
- [Qiu15] QIU Y.: Spectra. <https://spectralib.org/>, 2015. 8
- [RL01] RUSINKIEWICZ S., LEVOY M.: Efficient variants of the ICP algorithm. In *Proc. Conf. 3D Digital Imag. and Modeling* (2001), pp. 145–152. 2
- [RPWO18] REN J., POULENARD A., WONKA P., OVSJANIKOV M.: Continuous and orientation-preserving correspondences via functional maps. *ACM Trans. Graph.* 37, 6 (2018), 248:1–248:16. 2, 7
- [Rus07] RUSTAMOV R. M.: Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In *Proc. Symposium on Geometry Processing* (2007), pp. 225–233. 2
- [SOG09] SUN J., OVSJANIKOV M., GUIBAS L.: A concise and provably informative multi-scale signature based on heat diffusion. In *Proc. Symposium on Geometry Processing* (2009), pp. 1383–1392. 4
- [TCL*13] TAM G. K. L., CHENG Z., LAI Y., LANGBEIN F. C., LIU Y., MARSHALL D., MARTIN R. R., SUN X., ROSIN P. L.: Registration of 3D point clouds and meshes: A survey from rigid to nonrigid. *IEEE Tran. Vis. Comp. Graph.* 19, 7 (2013), 1199–1217. 2
- [vKZHC011] VAN KAICK O., ZHANG H., HAMARNEH G., COHEN-OR D.: A survey on shape correspondence. *Comp. Graph. Forum* 30, 6 (2011), 1681–1707. 2
- [YSV*10] YEO B. T. T., SABUNCU M. R., VERCAUTEREN T., AYACHE N., FISCHL B., GOLLAND P.: Spherical demons: Fast diffeomorphic landmark-free surface registration. *IEEE Trans. Medical Imag.* 29, 3 (2010), 650–668. 2
- [ZHG11] ZOU G., HU J., GU X.: Authalic parameterization of general surfaces using lie advection. In *IEEE Tran. Vis. Comp. Graph.* (2011), pp. 2004–2014. 2, 3

Appendix: Parameter Selection

There are a number of parameters that need to be set in our implementation. These include the authalic diffusion (τ), rotational HKS timescale (t_R), maximum number of rotation candidates (K), rotation separation distance (ϵ_1), correlation threshold (ϵ_2), optical flow HKS timescales ($\{t_0, t_1, t_2, t_3\}$), optical flow smoothing weight (α), and candidate selection HKS timescale (t_S).

A. Authalic evolution

We choose the value for the diffusion timescale τ by considering the deviation of the log scale factors from zero and the stability of the resulting flow. Intuitively, we would like to perform as little smoothing as possible while maintaining a robust implementation.

Figure 9 plots the average deviation across the 80 models in the TOSCA dataset (left) and the fraction of the sphere that is flipped (right) as a function of the number of authalic iterations, for different timescales τ . Initially, as the flow-field is smoothed less, the deviation in scale factors reduces more quickly. However, for too small timescales ($\tau < \frac{1}{8} \times 10^{-2}$ in our experiments), the flow-field is too high-frequency resulting in unstable advection, as evidenced by the large percentage of the sphere that becomes flipped.

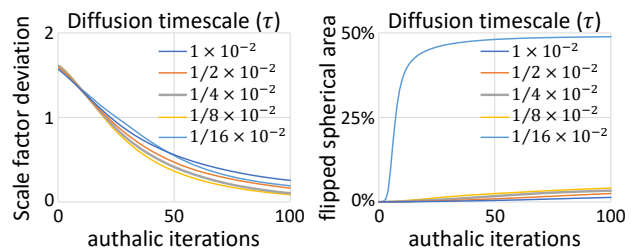


Figure 9: Average deviation of log scale factors (left) and average fraction of flipped triangles (right) as a function of the number of authalic iterations, for different timescales of diffusion.

B. Surface Signals

Given the relatively low-frequency timescales of the HKSs we use for alignment, we find that using the first 200 eigenvalues/eigenvectors is sufficient to accurately represent the signals.

C. Rotational alignment

To define candidate initial rotations, we prescribe the timescale of the HKS used for rotational alignment t_R , as well as the maximum number of candidates K , the separation distance ϵ_1 , and the correlation quality threshold ϵ_2 . We do this by fixing the optical flow timescales, $\{t_0, t_1, t_2, t_3\}$ and independently “training” the parameters by computing correspondences within the first three TOSCA categories (*Cat*, *Centaur*, and *David*). For our loss function we use the median biharmonic distance, for each pair of shapes returning the minimum biharmonic distance over all rotation candidates.

Using a single rotation candidate, $K = 1$, and trying timescales $t_R \in \{4, 2, 1, 0.5, 0.25\}$ we get corresponding median values 0.0250, 0.0236, 0.0235, 0.0235, and 0.0238, indicating a local minimum near $t_R = 1.0$. Fixing $t_R = 1.0$, we search the space of rotations by generating a large number of candidates for each pair of shapes, setting $K = \infty$, $\epsilon_1 = 2$, and $\epsilon_2 = 100\%$.

Analyzing the results obtained using the numerous candidates, we find that setting the correlation threshold $\epsilon_2 = 5\%$, we get a biharmonic distance that is within 0.003 of the biharmonic distance obtained using a threshold of $\epsilon_2 = 100\%$. Furthermore, the best rotation is always one of the top 3 candidates. To provide some leeway, we set $\epsilon_2 = 10\%$ and $K = 4$. Finally, trying different separation distances, we find little difference between the results using $\epsilon_1 = 3$ and $\epsilon_1 = 2$. Results with $\epsilon_1 = 4$ are slightly worse.

D. Optical flow refinement

As above we set the optical flow parameters by training on correspondences from the first three categories, using the median biharmonic distance as a loss function. Given the large parameter space, we follow the intuition of Prada *et al.* [PKC*16], using smoother signals for lower levels of the hierarchy. Specifically, given a choice of timescale t_0 for the coarsest level, we set the timescale at subsequent levels by dividing by two, $t_k = t_0/2^k$.

Fixing the optical flow smoothing weight $\alpha = 0.05$ and trying coarse optical flow timescales $t_0 \in \{2, 1, 0.5, 0.25, 0.125\}$ we get corresponding median values 0.0361, 0.0240, 0.0227, 0.0251, and 0.0306, indicating a local minimum near $t_0 = 0.5$. Fixing the coarse optical flow timescale to $t_0 = 0.5$ and trying smoothing weights $\alpha \in \{0.2, 0.1, 0.05, 0.025, 0.0125\}$ we get median values 0.0234, 0.0227, 0.0227, 0.0230, and 0.0236, indicating a local minimum near $\alpha = 0.05$.

Finally, fixing both the coarse timescale $t_0 = 0.5$ and smoothing weight $\alpha = 0.05$ and setting the candidate selection HKS timescales $t_S = 0.01$, we find that the selected candidate is almost always the one minimizing the biharmonic distance. Specifically, using the training set, we find that the difference between the biharmonic distance given by the selected rotation candidate and the rotation candidate minimizing the biharmonic distance is never larger than 0.002.